
Replication and MongoDB

Release 3.2.3

MongoDB, Inc.

February 17, 2016

© MongoDB, Inc. 2008 - 2016 This work is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 3.0 United States License](#)

1	Replication Introduction	3
1.1	Redundancy and Data Availability	3
1.2	Replication in MongoDB	3
1.3	Additional Resources	7
2	Replication Concepts	9
2.1	Replica Set Members	9
2.2	Replica Set Deployment Architectures	19
2.3	Replica Set High Availability	27
2.4	Replica Set Read and Write Semantics	32
2.5	Replication Processes	39
2.6	Master Slave Replication	43
3	Replica Set Tutorials	51
3.1	Replica Set Deployment Tutorials	52
3.2	Member Configuration Tutorials	71
3.3	Replica Set Maintenance Tutorials	80
3.4	Troubleshoot Replica Sets	101
4	Replication Reference	107
4.1	Replication Methods in the <code>mongo</code> Shell	107
4.2	Replication Database Commands	108
4.3	Replica Set Reference Documentation	108

A *replica set* in MongoDB is a group of `mongod` processes that maintain the same data set. Replica sets provide redundancy and high availability, and are the basis for all production deployments. This section introduces replication in MongoDB as well as the components and architecture of replica sets. The section also provides tutorials for common tasks related to replica sets.

***Replication Introduction* (page 3)** An introduction to replica sets, their behavior, operation, and use.

***Replication Concepts* (page 9)** The core documentation of replica set operations, configurations, architectures and behaviors.

***Replica Set Members* (page 9)** Introduces the components of replica sets.

***Replica Set Deployment Architectures* (page 19)** Introduces architectural considerations related to replica sets deployment planning.

***Replica Set High Availability* (page 27)** Presents the details of the automatic failover and recovery process with replica sets.

***Replica Set Read and Write Semantics* (page 32)** Presents the semantics for targeting read and write operations to the replica set, with an awareness of location and set configuration.

***Replica Set Tutorials* (page 51)** Tutorials for common tasks related to the use and maintenance of replica sets.

***Replication Reference* (page 107)** Reference for functions and operations related to replica sets.

Replication Introduction

On this page

- [Redundancy and Data Availability](#) (page 3)
- [Replication in MongoDB](#) (page 3)
- [Additional Resources](#) (page 7)

Replication is the process of synchronizing data across multiple servers.

1.1 Redundancy and Data Availability

Replication provides redundancy and increases data availability. With multiple copies of data on different database servers, replication provides a level of fault tolerance against the loss of a single database server.

In some cases, replication can provide increased read capacity as clients can send read operations to different servers. Maintaining copies of data in different data centers can increase data locality and availability for distributed applications. You can also maintain additional copies for dedicated purposes, such as disaster recovery, reporting, or backup.

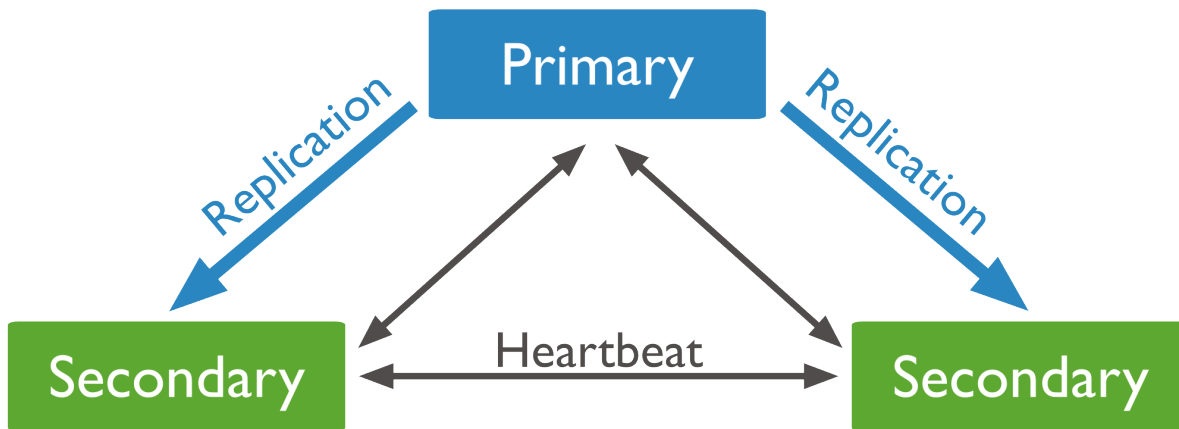
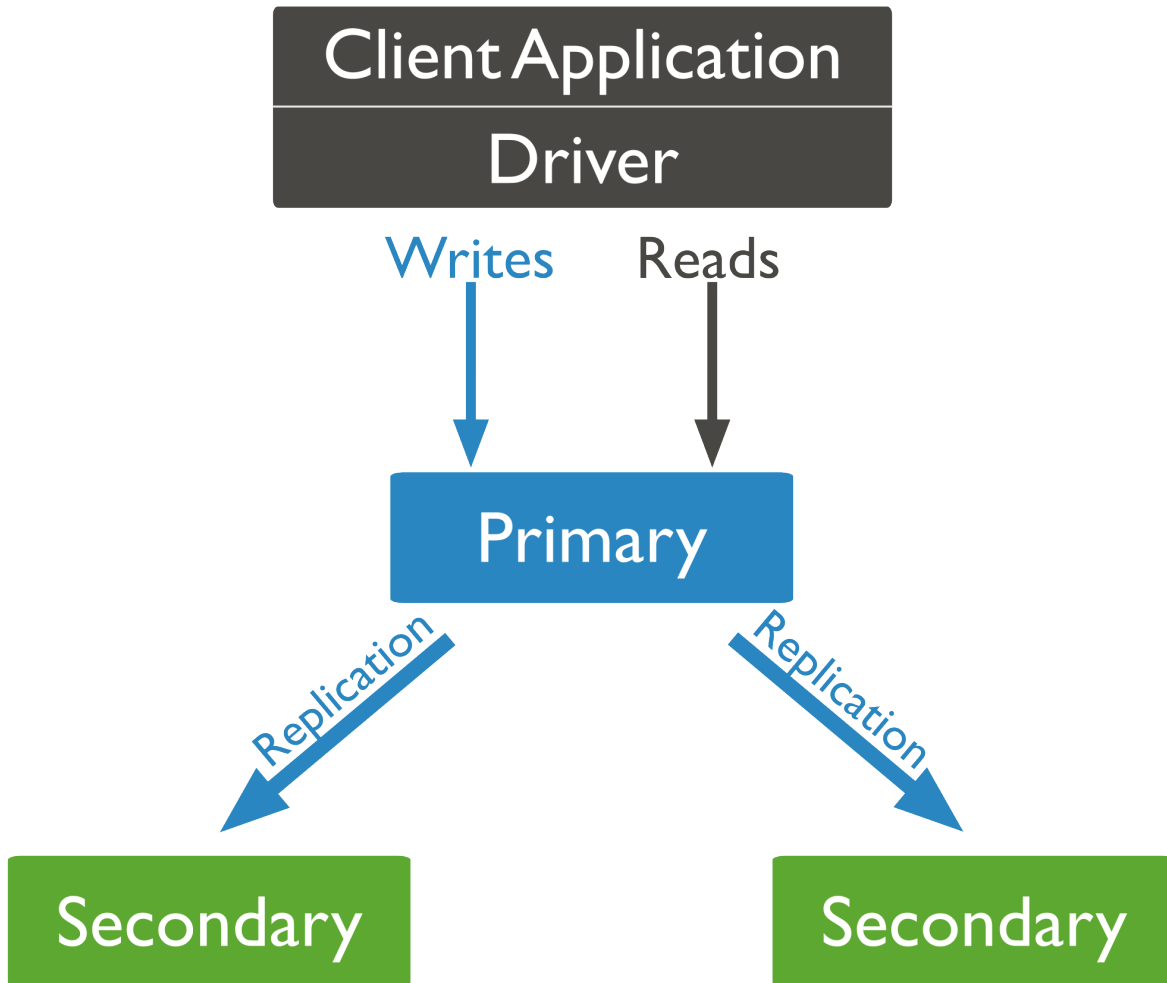
1.2 Replication in MongoDB

A replica set is a group of `mongod` instances that maintain the same data set. A replica set contains several data bearing nodes and optionally one arbiter node. Of the data bearing nodes, one and only one member is deemed the primary node, while the other nodes are deemed secondary nodes.

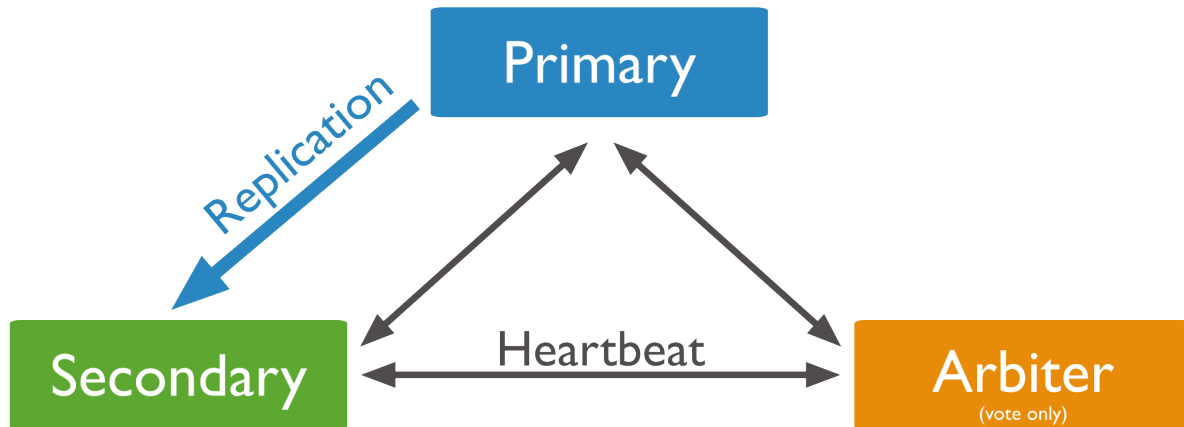
The *primary node* (page 10) receives all write operations. A replica set can have only one primary capable of confirming writes with `{ w: "majority" }` write concern; although in some circumstances, another `mongod` instance may transiently believe itself to also be primary.¹ The primary records all changes to its data sets in its operation log, i.e. *oplog* (page 40). For more information on primary node operation, see *Replica Set Primary* (page 10).

The *secondaries* (page 10) replicate the primary's *oplog* and apply the operations to their data sets such that the secondaries' data sets reflect the primary's data set. If the primary is unavailable, an eligible secondary will hold an election to elect itself the new primary. For more information on secondary members, see *Replica Set Secondary Members* (page 10).

¹ In *some circumstances* (page 121), two nodes in a replica set may *transiently* believe that they are the primary, but at most, one of them will be able to complete writes with `{ w: "majority" }` write concern. The node that can complete `{ w: "majority" }` writes is the current primary, and the other node is a former primary that has not yet recognized its demotion, typically due to a *network partition*. When this occurs, clients that connect to the former primary may observe stale data despite having requested read preference `primary` (page 120), and new writes to the former primary will eventually roll back.



You may add an extra `mongod` instance to a replica set as an *arbiter* (page 17). Arbiters do not maintain a data set. The purpose of an arbiter is to maintain a quorum in a replica set by responding to heartbeat and election requests by other replica set members. Because they do not store a data set, arbiters can be a good way to provide replica set quorum functionality with a cheaper resource cost than a fully functional replica set member with a data set. If your replica set has an even number of members, add an arbiter to obtain a majority of votes in an election for primary. Arbiters do not require dedicated hardware. For more information on arbiters, see *Replica Set Arbiter* (page 17).



An *arbiter* (page 17) will always be an arbiter whereas a *primary* (page 10) may step down and become a *secondary* (page 10) and a *secondary* (page 10) may become the primary during an election.

1.2.1 Asynchronous Replication

Secondaries apply operations from the primary asynchronously. By applying operations after the primary, sets can continue to function despite the failure of one or more members. For more information on replication mechanics, see *Replica Set Oplog* (page 40) and *Replica Set Data Synchronization* (page 41).

1.2.2 Automatic Failover

When a primary does not communicate with the other members of the set for more than 10 seconds, an eligible secondary will hold an election to elect itself the new primary. The first secondary to hold an election and receive a majority of the members' votes becomes primary.


New in version 3.2: MongoDB introduces a version 1 of the replication protocol (`protocolVersion: 1` (page 109)) to reduce replica set failover time and accelerates the detection of multiple simultaneous primaries. New replica sets will, by default, use `protocolVersion: 1` (page 109). Previous versions of MongoDB use version 0 of the protocol.

See *Replica Set Elections* (page 28) and *Rollbacks During Replica Set Failover* (page 31) for more information.

1.2.3 Read Operations

By default, clients read from the primary¹; however, clients can specify a *read preference* (page 34) to send read operations to secondaries. *Asynchronous replication* (page 5) to secondaries means that reads from secondaries may return data that does not reflect the state of the data on the primary. For information on reading from replica sets, see *Read Preference* (page 34).

In MongoDB, clients can see the results of writes before the writes are *durable*:



- Regardless of `write concern`, other clients using "local" (i.e. the default) `readConcern` can see the result of a write operation before the write operation is acknowledged to the issuing client.
- Clients using "local" (i.e. the default) `readConcern` can read data which may be subsequently *rolled back* (page 31).

For more information on read isolations, consistency and recency for MongoDB, see <https://docs.mongodb.org/manual/core/read-isolation-consistency-recency>.

1.2.4 Additional Features

Replica sets provide a number of options to support application needs. For example, you may deploy a replica set with *members in multiple data centers* (page 26), or control the outcome of elections by adjusting the `members[n].priority` (page 111) of some members. Replica sets also support dedicated members for reporting, disaster recovery, or backup functions.

See *Priority 0 Replica Set Members* (page 14), *Hidden Replica Set Members* (page 15) and *Delayed Replica Set Members* (page 16) for more information.

1.3 Additional Resources

- Quick Reference Cards²
- Webinar: Managing Your Mission Critical App - Ensuring Zero Downtime³

²<https://www.mongodb.com/lp/misc/quick-reference-cards?jmp=docs>

³<http://www.mongodb.com/webinar/managing-mission-critical-app-downtime?jmp=docs>

Replication Concepts

These documents describe and provide examples of replica set operation, configuration, and behavior. For an overview of replication, see *Replication Introduction* (page 3). For documentation of the administration of replica sets, see *Replica Set Tutorials* (page 51). The *Replication Reference* (page 107) documents commands and operations specific to replica sets.

***Replica Set Members* (page 9)** Introduces the components of replica sets.

***Replica Set Primary* (page 10)** The primary is the only member of a replica set that accepts write operations.

***Replica Set Secondary Members* (page 10)** Secondary members replicate the primary's data set and accept read operations. If the set has no primary, a secondary can become primary.

***Priority 0 Replica Set Members* (page 14)** Priority 0 members are secondaries that cannot become the primary.

***Hidden Replica Set Members* (page 15)** Hidden members are secondaries that are invisible to applications. These members support dedicated workloads, such as reporting or backup.

***Replica Set Arbiter* (page 17)** An arbiter does not maintain a copy of the data set but participate in elections.

***Replica Set Deployment Architectures* (page 19)** Introduces architectural considerations related to replica sets deployment planning.

***Replica Set High Availability* (page 27)** Presents the details of the automatic failover and recovery process with replica sets.

***Replica Set Elections* (page 28)** Elections occur when the primary becomes unavailable and the replica set members autonomously select a new primary.

***Read Preference* (page 34)** Read preference specifies where (i.e. which members of the replica set) the drivers should direct the read operations.

***Replication Processes* (page 39)** Mechanics of the replication process and related topics.

***Master Slave Replication* (page 43)** Master-slave replication provided redundancy in early versions of MongoDB. Replica sets replace master-slave for most use cases.

2.1 Replica Set Members

A *replica set* in MongoDB is a group of `mongod` processes that provide redundancy and high availability. The members of a replica set are:

***Primary* (page ??)**. The primary receives all write operations.

Secondaries (page ??). Secondaries replicate operations from the primary to maintain an identical data set. Secondaries may have additional configurations for special usage profiles. For example, secondaries may be *non-voting* (page 30) or *priority 0* (page 14).

You can also maintain an *arbiter* (page ??) as part of a replica set. Arbiters do not keep a copy of the data. However, arbiters play a role in the elections that select a primary if the current primary is unavailable.

The minimum requirements for a replica set are: A *primary* (page ??), a *secondary* (page ??), and an *arbiter* (page ??). Most deployments, however, will keep three members that store data: A *primary* (page ??) and two *secondary members* (page ??).

Changed in version 3.0.0: A replica set can have up to 50 members but only 7 voting members.¹ In previous versions, replica sets can have up to 12 members.

2.1.1 Replica Set Primary

The primary is the only member in the replica set that receives write operations. MongoDB applies write operations on the *primary* and then records the operations on the primary's *oplog* (page 40). *Secondary* (page ??) members replicate this log and apply the operations to their data sets.

In the following three-member replica set, the primary accepts all write operations. Then the secondaries replicate the oplog to apply to their data sets.

All members of the replica set can accept read operations. However, by default, an application directs its read operations to the primary member. See *Read Preference* (page 34) for details on changing the default read behavior.

The replica set can have at most one primary.² If the current primary becomes unavailable, an election determines the new primary. See *Replica Set Elections* (page 28) for more details.

In the following 3-member replica set, the primary becomes unavailable. This triggers an election which selects one of the remaining secondaries as the new primary.

2.1.2 Replica Set Secondary Members

A secondary maintains a copy of the *primary's* data set. To replicate data, a secondary applies operations from the primary's *oplog* (page 40) to its own data set in an asynchronous process. A replica set can have one or more secondaries.

The following three-member replica set has two secondary members. The secondaries replicate the primary's oplog and apply the operations to their data sets.

Although clients cannot write data to secondaries, clients can read data from secondary members. See *Read Preference* (page 34) for more information on how clients direct read operations to replica sets.

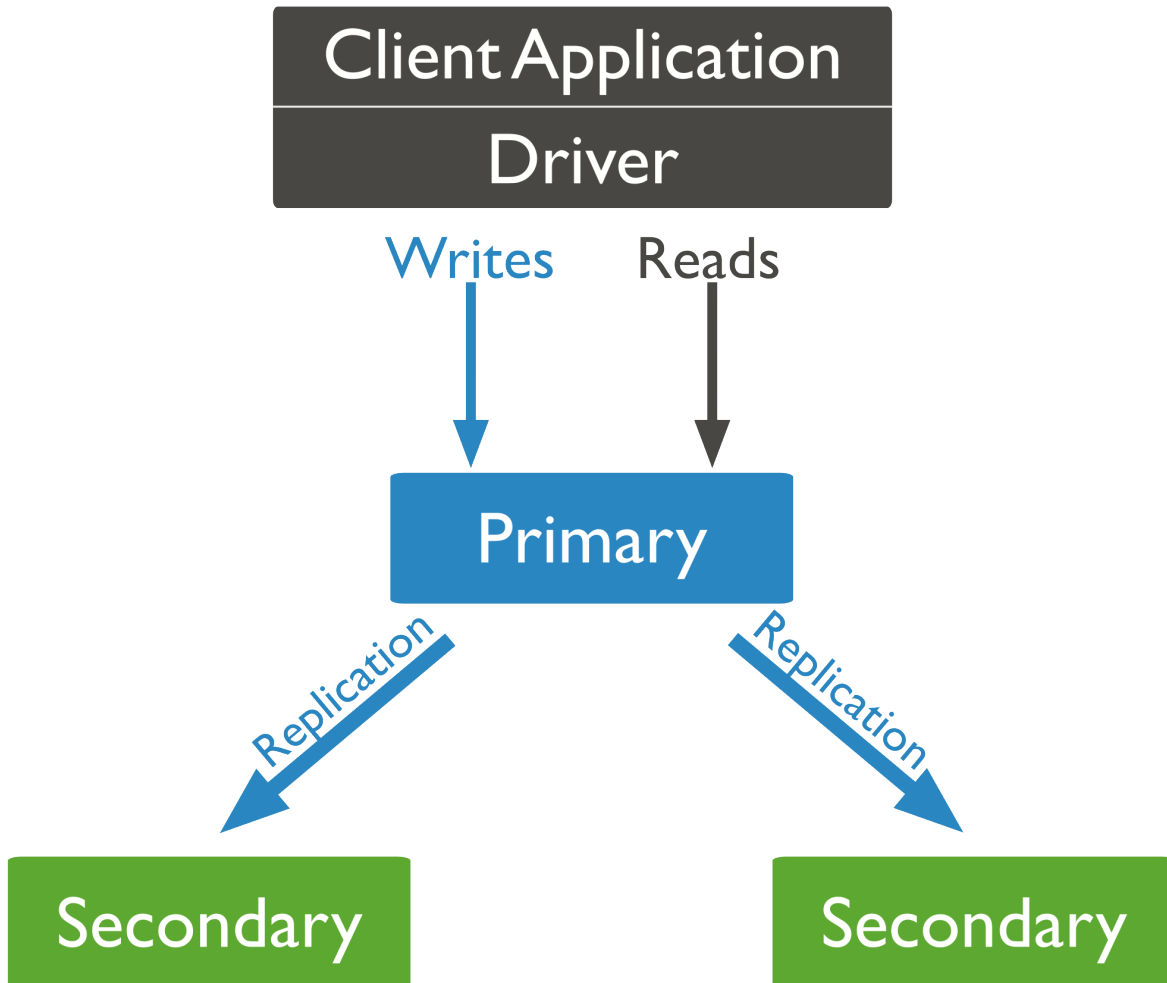
A secondary can become a primary. If the current primary becomes unavailable, the replica set holds an *election* to choose which of the secondaries becomes the new primary.


In the following three-member replica set, the primary becomes unavailable. This triggers an election where one of the remaining secondaries becomes the new primary.

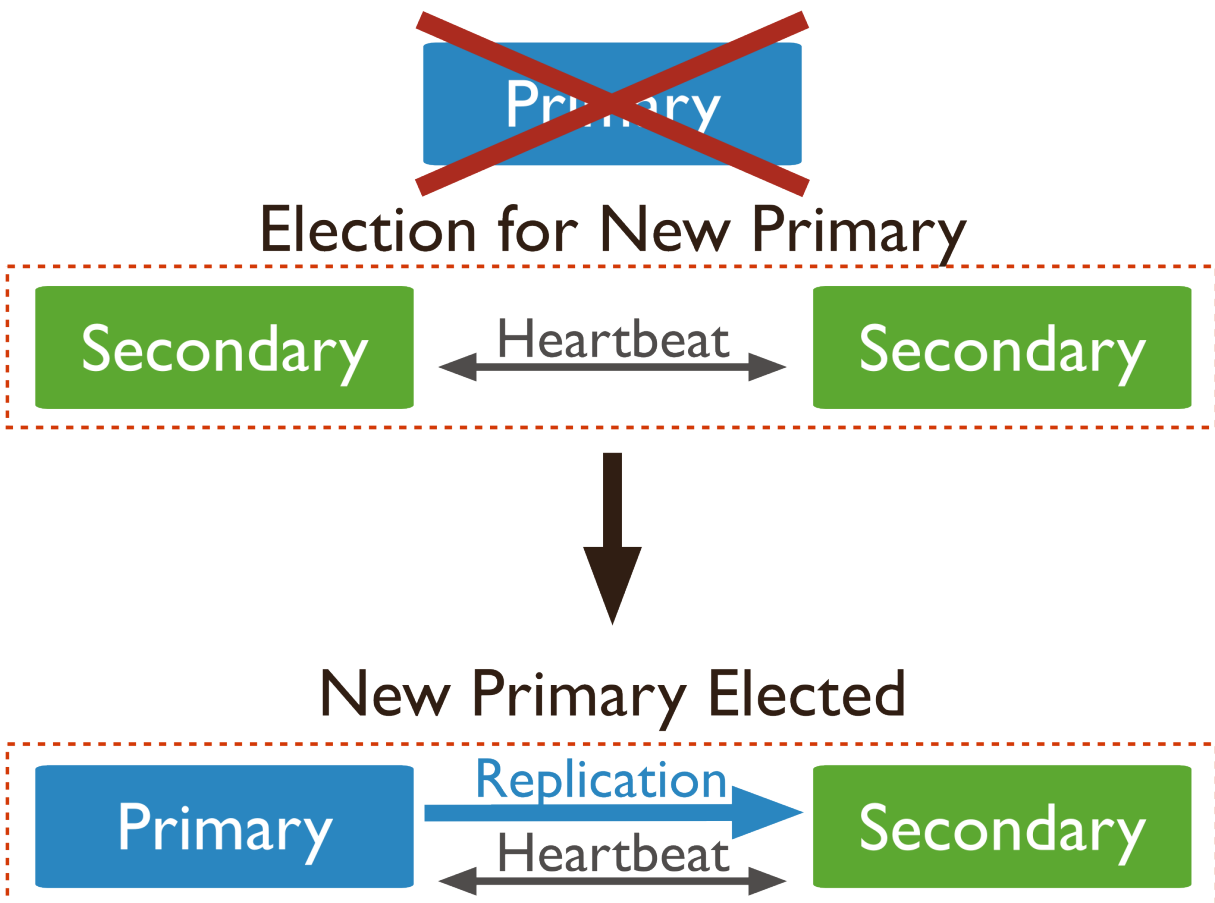
See *Replica Set Elections* (page 28) for more details.

¹ While replica sets are the recommended solution for production, a replica set can support up to 50 members in total. If your deployment requires more than 50 members, you'll need to use *master-slave* (page 43) replication. However, master-slave replication lacks the automatic failover capabilities.

² In *some circumstances* (page 121), two nodes in a replica set may *transiently* believe that they are the primary, but at most, one of them will be able to complete writes with { w: "majority" } write concern. The node that can complete { w: "majority" } writes is the current primary, and the other node is a former primary that has not yet recognized its demotion, typically due to a *network partition*. When this occurs, clients that connect to the former primary may observe stale data despite having requested read preference *primary* (page 120), and new writes to the former primary will eventually roll back.







You can configure a secondary member for a specific purpose. You can configure a secondary to:

- Prevent it from becoming a primary in an election, which allows it to reside in a secondary data center or to serve as a cold standby. See *Priority 0 Replica Set Members* (page 14).
- Prevent applications from reading from it, which allows it to run applications that require separation from normal traffic. See *Hidden Replica Set Members* (page 15).
- Keep a running “historical” snapshot for use in recovery from certain errors, such as unintentionally deleted databases. See *Delayed Replica Set Members* (page 16).


Priority 0 Replica Set Members

On this page

- [Priority 0 Members as Standbys](#) (page 14)
- [Priority 0 Members and Failover](#) (page 15)
- [Configuration](#) (page 15)

A *priority 0* member is a secondary that **cannot** become *primary*. *Priority 0* members cannot *trigger elections*. Otherwise these members function as normal secondaries. A *priority 0* member maintains a copy of the data set, accepts read operations, and votes in elections. Configure a *priority 0* member to prevent *secondaries* from becoming primary, which is particularly useful in multi-data center deployments.

In a three-member replica set, in one data center hosts the primary and a secondary. A second data center hosts one *priority 0* member that cannot become primary.



Priority 0 Members as Standbys

A *priority 0* member can function as a standby. In some replica sets, it might not be possible to add a new member in a reasonable amount of time. A standby member keeps a current copy of the data to be able to replace an unavailable member.

In many cases, you need not set standby to *priority 0*. However, in sets with varied hardware or *geographic distribution* (page 26), a *priority 0* standby ensures that only qualified members become primary.

A *priority 0* standby may also be valuable for some members of a set with different hardware or workload profiles. In these cases, deploy a member with *priority 0* so it can't become primary. Also consider using an *hidden member* (page 15) for this purpose.

If your set already has seven voting members, also configure the member as *non-voting* (page 30).

Priority 0 Members and Failover

When configuring a *priority 0* member, consider potential failover patterns, including all possible network partitions. Always ensure that your main data center contains both a quorum of voting members and contains members that are eligible to be primary.

Configuration

To configure a *priority 0* member, see *Prevent Secondary from Becoming Primary* (page 73).

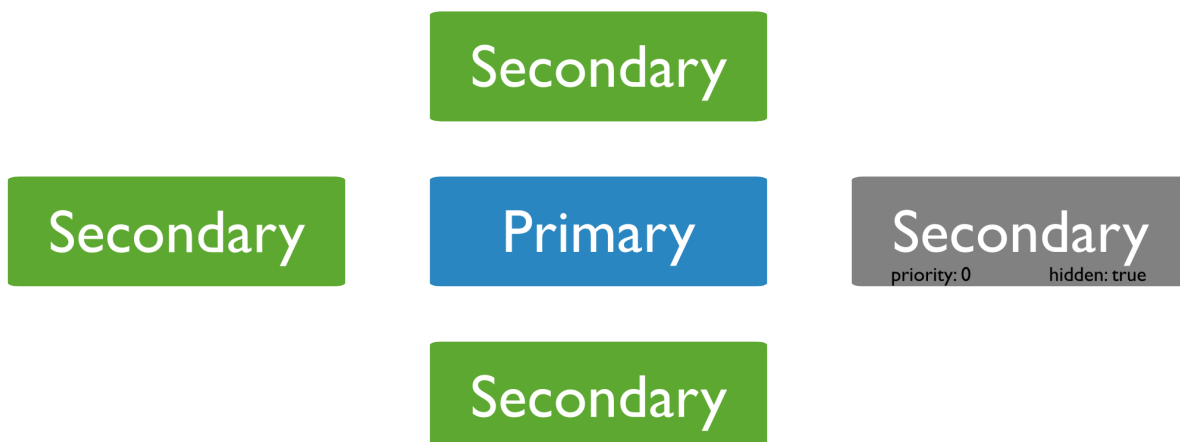
Hidden Replica Set Members

On this page

- [Behavior](#) (page 15)
- [Further Reading](#) (page 16)

A hidden member maintains a copy of the *primary's* data set but is **invisible** to client applications. Hidden members are good for workloads with different usage patterns from the other members in the *replica set*. Hidden members must always be *priority 0 members* (page 14) and so **cannot become primary**. The `db.isMaster()` method does not display hidden members. Hidden members, however, **may vote** in *elections* (page 28).

In the following five-member replica set, all four secondary members have copies of the primary's data set, but one of the secondary members is hidden.



Behavior

Read Operations Clients will not distribute reads with the appropriate *read preference* (page 34) to hidden members. As a result, these members receive no traffic other than basic replication. Use hidden members for dedicated tasks such as reporting and backups. *Delayed members* (page 16) should be hidden.

In a sharded cluster, `mongos` do not interact with hidden members.

Voting Hidden members *may* vote in replica set elections. If you stop a voting hidden member, ensure that the set has an active majority or the *primary* will step down.

For the purposes of backups,

- If using the MMAPv1 storage engine, you can avoid stopping a hidden member with the `db.fsyncLock()` and `db.fsyncUnlock()` operations to flush all writes and lock the `mongod` instance for the duration of the backup operation.
- Changed in version 3.2: Starting in MongoDB 3.2, `db.fsyncLock()` can ensure that the data files do not change for MongoDB instances using either the MMAPv1 or the WiredTiger storage engine, thus providing consistency for the purposes of creating backups.

In previous MongoDB version, `db.fsyncLock()` *cannot* guarantee a consistent set of files for low-level backups (e.g. via file copy `cp`, `scp`, `tar`) for WiredTiger.

Further Reading

For more information about backing up MongoDB databases, see <https://docs.mongodb.org/manual/core/backups>. To configure a hidden member, see *Configure a Hidden Replica Set Member* (page 74).

Delayed Replica Set Members

On this page

- [Considerations](#) (page 16)
- [Example](#) (page 17)
- [Configuration](#) (page 17)

Delayed members contain copies of a *replica set's* data set. However, a delayed member's data set reflects an earlier, or delayed, state of the set. For example, if the current time is 09:52 and a member has a delay of an hour, the delayed member has no operation more recent than 08:52.

Because delayed members are a “rolling backup” or a running “historical” snapshot of the data set, they may help you recover from various kinds of human error. For example, a delayed member can make it possible to recover from unsuccessful application upgrades and operator errors including dropped databases and collections.

Considerations

Requirements Delayed members:

- **Must be** *priority 0* (page 14) members. Set the priority to 0 to prevent a delayed member from becoming primary.
- **Should be** *hidden* (page 15) members. Always prevent applications from seeing and querying delayed members.
- *do* vote in *elections* for primary.

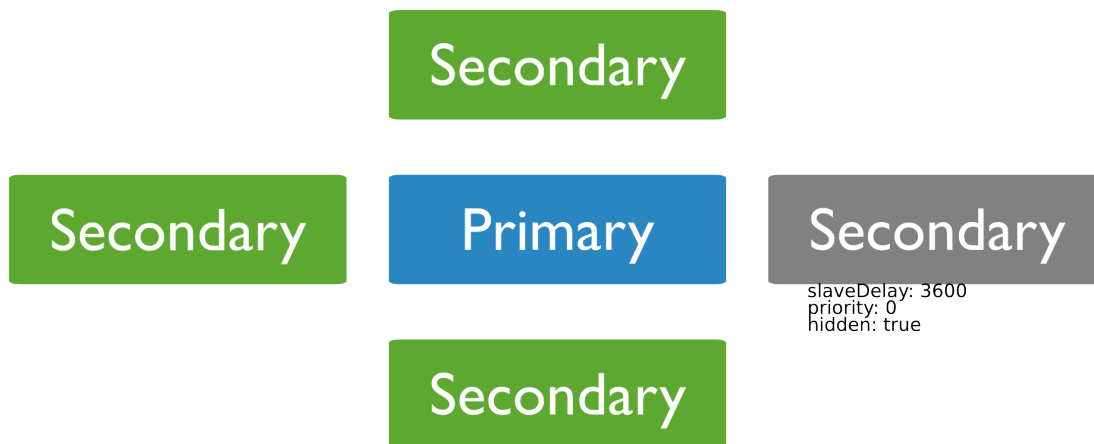
Behavior Delayed members apply operations from the *oplog* on a delay. When choosing the amount of delay, consider that the amount of delay:

- must be is equal to or greater than your maintenance windows.
- must be *smaller* than the capacity of the oplog. For more information on oplog size, see *Oplog Size* (page 40).

Sharding In sharded clusters, delayed members have limited utility when the *balancer* is enabled. Because delayed members replicate chunk migrations with a delay, the state of delayed members in a sharded cluster are not useful for recovering to a previous state of the sharded cluster if any migrations occur during the delay window.

Example

In the following 5-member replica set, the primary and all secondaries have copies of the data set. One member applies operations with a delay of 3600 seconds, or an hour. This delayed member is also *hidden* and is a *priority 0 member*.



Configuration

A delayed member has its `members[n].priority` (page 111) equal to 0, `members[n].hidden` (page 111) equal to `true`, and its `members[n].slaveDelay` (page 112) equal to the number of seconds of delay:

```
{
  "_id" : <num>,
  "host" : <hostname:port>,
  "priority" : 0,
  "slaveDelay" : <seconds>,
  "hidden" : true
}
```

To configure a delayed member, see *Configure a Delayed Replica Set Member* (page 76).

2.1.3 Replica Set Arbiter

On this page

- [Example](#) (page 18)
- [Security](#) (page 18)

An arbiter does **not** have a copy of data set and **cannot** become a primary. Replica sets may have arbiters to add a vote in *elections of for primary* (page 28). Arbiters *always* have exactly 1 vote election, and thus allow replica sets to have an uneven number of members, without the overhead of a member that replicates data.

Important: Do not run an arbiter on systems that also host the primary or the secondary members of the replica set.

Only add an arbiter to sets with even numbers of members. If you add an arbiter to a set with an odd number of members, the set may suffer from tied *elections*. To add an arbiter, see *Add an Arbiter to Replica Set* (page 64).

Example

For example, in the following replica set, an arbiter allows the set to have an odd number of votes for elections:



Security

Authentication

When running with `authorization`, arbiters exchange credentials with other members of the set to authenticate. MongoDB encrypts the authentication process. The MongoDB authentication exchange is cryptographically secure.

Arbiters use `keyfiles` to authenticate to the replica set.

Communication

The only communication between arbiters and other set members are: votes during elections, heartbeats, and configuration data. These exchanges are not encrypted.

However, if your MongoDB deployment uses TLS/SSL, MongoDB will encrypt *all* communication between replica set members. See <https://docs.mongodb.org/manual/tutorial/configure-ssl> for more information.

As with all MongoDB components, run arbiters in trusted network environments.

2.2 Replica Set Deployment Architectures

On this page

- [Strategies](#) (page 19)
- [Replica Set Naming](#) (page 20)
- [Deployment Patterns](#) (page 21)

The architecture of a *replica set* affects the set's capacity and capability. This document provides strategies for replica set deployments and describes common architectures.

The standard replica set deployment for production system is a three-member replica set. These sets provide redundancy and fault tolerance. Avoid complexity when possible, but let your application requirements dictate the architecture.

2.2.1 Strategies

Determine the Number of Members

Add members in a replica set according to these strategies.

Deploy an Odd Number of Members

An odd number of members ensures that the replica set is always able to elect a primary. If you have an even number of members, add an arbiter to get an odd number. *Arbiters* do not store a copy of the data and require fewer resources. As a result, you may run an arbiter on an application server or other shared process.

Consider Fault Tolerance

Fault tolerance for a replica set is the number of members that can become unavailable and still leave enough members in the set to elect a primary. In other words, it is the difference between the number of members in the set and the majority needed to elect a primary. Without a primary, a replica set cannot accept write operations. Fault tolerance is an effect of replica set size, but the relationship is not direct. See the following table:

Number of Members.	Majority Required to Elect a New Primary.	Fault Tolerance.
3	2	1
4	3	1
5	3	2
6	4	2

Adding a member to the replica set does not *always* increase the fault tolerance. However, in these cases, additional members can provide support for dedicated functions, such as backups or reporting.

Use Hidden and Delayed Members for Dedicated Functions

Add *hidden* (page 15) or *delayed* (page 16) members to support dedicated functions, such as backup or reporting.

Load Balance on Read-Heavy Deployments

In a deployment with *very* high read traffic, you can improve read throughput by distributing reads to secondary members. As your deployment grows, add or move members to alternate data centers to improve redundancy and availability.

Always ensure that the main facility is able to elect a primary.

Add Capacity Ahead of Demand

The existing members of a replica set must have spare capacity to support adding a new member. Always add new members before the current demand saturates the capacity of the set.

Determine the Distribution of Members

Distribute Members Geographically

To protect your data if your main data center fails, keep at least one member in an alternate data center. Set these members' `members[n].priority` (page 111) to 0 to prevent them from becoming primary.

Keep a Majority of Members in One Location

When a replica set has members in multiple data centers, network partitions can prevent communication between data centers. To replicate data, members must be able to communicate to other members.

In an election, members must see each other to create a majority. To ensure that the replica set members can confirm a majority and elect a primary, keep a majority of the set's members in one location.

Target Operations with Tag Sets

Use *replica set tag sets* (page 88) to ensure that operations replicate to specific data centers. Tag sets also allow the routing of read operations to specific machines.

See also:

[/data-center-awareness](#) and <https://docs.mongodb.org/manual/core/operational-segregation>.

Use Journaling to Protect Against Power Failures

Enable journaling to protect data against service interruptions. Without journaling MongoDB cannot recover data after unexpected shutdowns, including power failures and unexpected reboots.

All 64-bit versions of MongoDB after version 2.0 have journaling enabled by default.

2.2.2 Replica Set Naming

If your application connects to more than one replica set, each set should have a distinct name. Some drivers group replica set connections by replica set name.

2.2.3 Deployment Patterns

The following documents describe common replica set deployment patterns. Other patterns are possible and effective depending on the application's requirements. If needed, combine features of each architecture in your own deployment:

Three Member Replica Sets (page 21) Three-member replica sets provide the minimum recommended architecture for a replica set.

Replica Sets with Four or More Members (page 23) Four or more member replica sets provide greater redundancy and can support greater distribution of read operations and dedicated functionality.

Geographically Distributed Replica Sets (page 26) Geographically distributed sets include members in multiple locations to protect against facility-specific failures, such as power outages.

Three Member Replica Sets

On this page

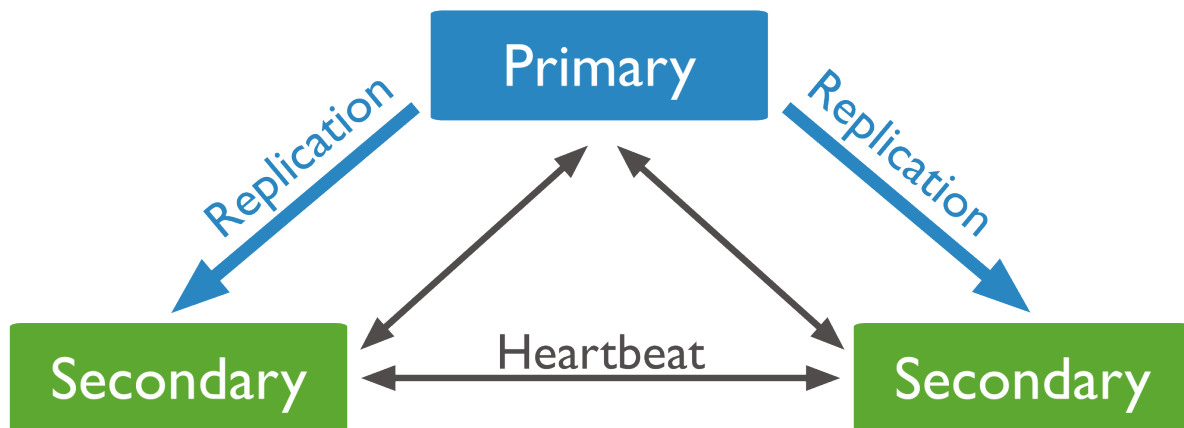
- [Primary with Two Secondary Members \(page 21\)](#)
- [Primary with a Secondary and an Arbiter \(page 23\)](#)

The minimum architecture of a replica set has three members. A three member replica set can have either three members that hold data, or two members that hold data and an arbiter.


Primary with Two Secondary Members

A replica set with three members that store data has:

- One *primary* (page 10).
- Two *secondary* (page 10) members. Both secondaries can become the primary in an *election* (page 28).



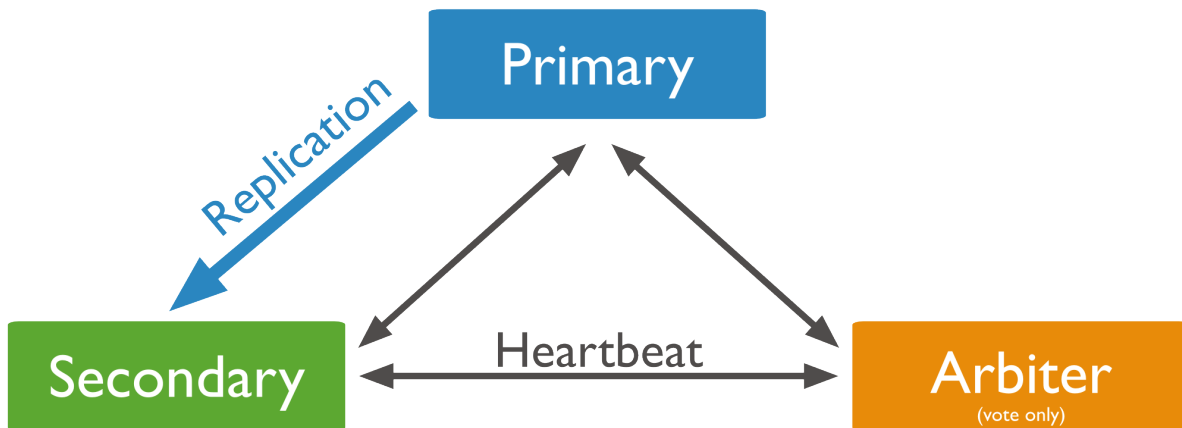
These deployments provide two complete copies of the data set at all times in addition to the primary. These replica sets provide additional fault tolerance and *high availability* (page 27). If the primary is unavailable, the replica set elects a secondary to be primary and continues normal operation. The old primary rejoins the set when available.



Primary with a Secondary and an Arbiter

A three member replica set with a two members that store data has:

- One *primary* (page 10).
- One *secondary* (page 10) member. The secondary can become primary in an *election* (page 28).
- One *arbiter* (page 17). The arbiter only votes in elections.



Since the arbiter does not hold a copy of the data, these deployments provides only one complete copy of the data. Arbiters require fewer resources, at the expense of more limited redundancy and fault tolerance.

However, a deployment with a primary, secondary, and an arbiter ensures that a replica set remains available if the primary *or* the secondary is unavailable. If the primary is unavailable, the replica set will elect the secondary to be primary.

See also:

[Deploy a Replica Set](#) (page 52).

Replica Sets with Four or More Members

On this page


- [Overview](#) (page 23)
- [Considerations](#) (page 23)

Overview

Although the standard replica set configuration has three members, you can deploy larger sets. Add additional members to a set to increase redundancy or to add capacity for distributing secondary read operations.

Considerations

As you add new members to a replica set, consider the following:



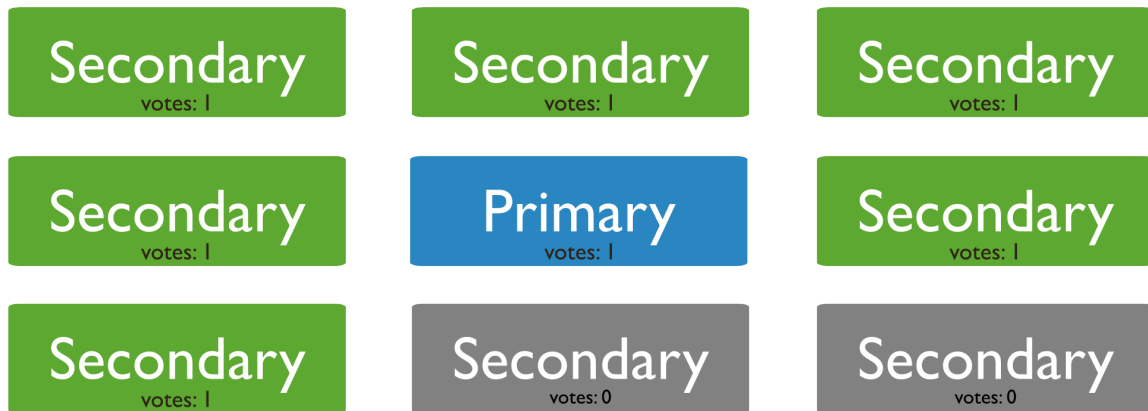
Odd Number of Voting Members Ensure that the replica set has an odd number of voting members. If you have an *even* number of voting members, deploy an *arbiter* (page ??) so that the set has an odd number.

For example, the following replica set includes an arbiter to ensure an odd number of voting members.



Maximum Number of Voting Members A replica set can have up to 50 members, but only 7 voting members.³ If the replica set already has 7 voting members, additional members must be *non-voting members* (page 30).


For example, the following 9 member replica set has 7 voting members and 2 non-voting members.



See *Non-Voting Members* (page 30) for more information.


Location of the Members A majority of the replica set’s members should be in your application’s main data center. For example, the following 5 member replica set has the majority, 3, of its members in its main data center, Data Center 1.

³ While replica sets are the recommended solution for production, a replica set can support up to 50 members in total. If your deployment requires more than 50 members, you’ll need to use *master-slave* (page 43) replication. However, master-slave replication lacks the automatic failover capabilities.



Electability of Members Some members of the replica set, such as members that have networking restraint or limited resources, should not be able to become primary in a *failover*. Configure members that should not become primary to have *priority 0* (page 14).

For example, the secondary member in the third data center with a priority of 0 cannot become primary:



See also:

Deploy a Replica Set (page 52), *Add an Arbiter to Replica Set* (page 64), and *Add Members to a Replica Set* (page 66).

Geographically Distributed Replica Sets

On this page

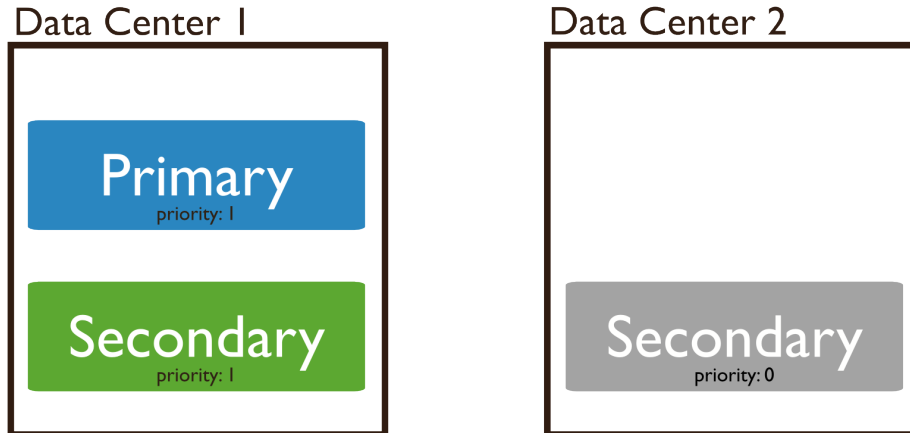
- [Additional Resource](#) (page 27)

Adding members to a replica set in multiple data centers adds redundancy and provides fault tolerance if one data center is unavailable. Members in additional data centers should have a *priority of 0* (page 14) to prevent them from becoming primary.

For example: the architecture of a geographically distributed replica set may be:

- One *primary* in the main data center.
- One *secondary* member in the main data center. This member can become primary at any time.
- One *priority 0* (page 14) member in a second data center. This member cannot become primary.

In the following replica set, the primary and one secondary are in *Data Center 1*, while *Data Center 2* has a *priority 0* (page 14) secondary that cannot become a primary.



If the primary is unavailable, the replica set will elect a new primary from *Data Center 1*. If the data centers cannot connect to each other, the member in *Data Center 2* will not become the primary.

If *Data Center 1* becomes unavailable, you can manually recover the data set from *Data Center 2* with minimal downtime.

To facilitate elections, the main data center should hold a majority of members. Also ensure that the set has an odd number of members. If adding a member in another data center results in a set with an even number of members, deploy an *arbiter* (page ??). For more information on elections, see *Replica Set Elections* (page 28).

See also:

Deploy a Geographically Redundant Replica Set (page 58).

Additional Resource

- Whitepaper: MongoDB Multi-Data Center Deployments⁴
- Webinar: Multi-Data Center Deployment⁵

2.3 Replica Set High Availability

Replica sets provide high availability using automatic *failover*. Failover allows a *secondary* member to become *primary* if the current primary becomes unavailable.

Changed in version 3.2: MongoDB introduces a version 1 of the replication protocol (`protocolVersion: 1` (page 109)) to reduce replica set failover time and accelerates the detection of multiple simultaneous primaries. New replica sets will, by default, use `protocolVersion: 1` (page 109). Previous versions of MongoDB use version

⁴<http://www.mongodb.com/lp/white-paper/multi-de?jmp=docs>

⁵<https://www.mongodb.com/presentations/webinar-multi-data-center-deployment?jmp=docs>

0 of the protocol. To upgrade existing replica sets to use `protocolVersion: 1` (page 109), see *3.2-upgrade-replica-set*.

Replica set members keep the same data set but are otherwise independent. If the primary becomes unavailable, an eligible secondary holds an *election* (page 28) to elect itself as a new primary. In some situations, the failover process may undertake a *rollback* (page 31).⁶

The deployment of a replica set affects the outcome of failover situations. To support effective failover, ensure that one facility can elect a primary if needed. Choose the facility that hosts the core application systems to host the majority of the replica set. Place a majority of voting members and all the members that can become primary in this facility. Otherwise, network partitions could prevent the set from being able to form a majority.

2.3.1 Replica Set Elections

On this page

- [Factors and Conditions that Affect Elections](#) (page 28)
- [Non-Voting Members](#) (page 30)

Replica sets use elections to determine which set member will become *primary*. Elections occur after initiating a replica set, and also any time the primary becomes unavailable. The primary is the only member in the set that can accept write operations. If a primary becomes unavailable, elections allow the set to recover normal operations without manual intervention. Elections are part of the *failover process* (page 27).

In the following three-member replica set, the primary is unavailable. One of the remaining secondaries holds an election to elect itself as a new primary.

Elections are essential for independent operation of a replica set; however, elections take time to complete. While an election is in process, the replica set has no primary and cannot accept writes and all remaining members become read-only. MongoDB avoids elections unless necessary.

If a majority of the replica set is inaccessible or unavailable to the current primary, the primary will step down and become a secondary. The replica set cannot accept writes after this occurs, but remaining members can continue to serve read queries if such queries are configured to run on secondaries.

Factors and Conditions that Affect Elections

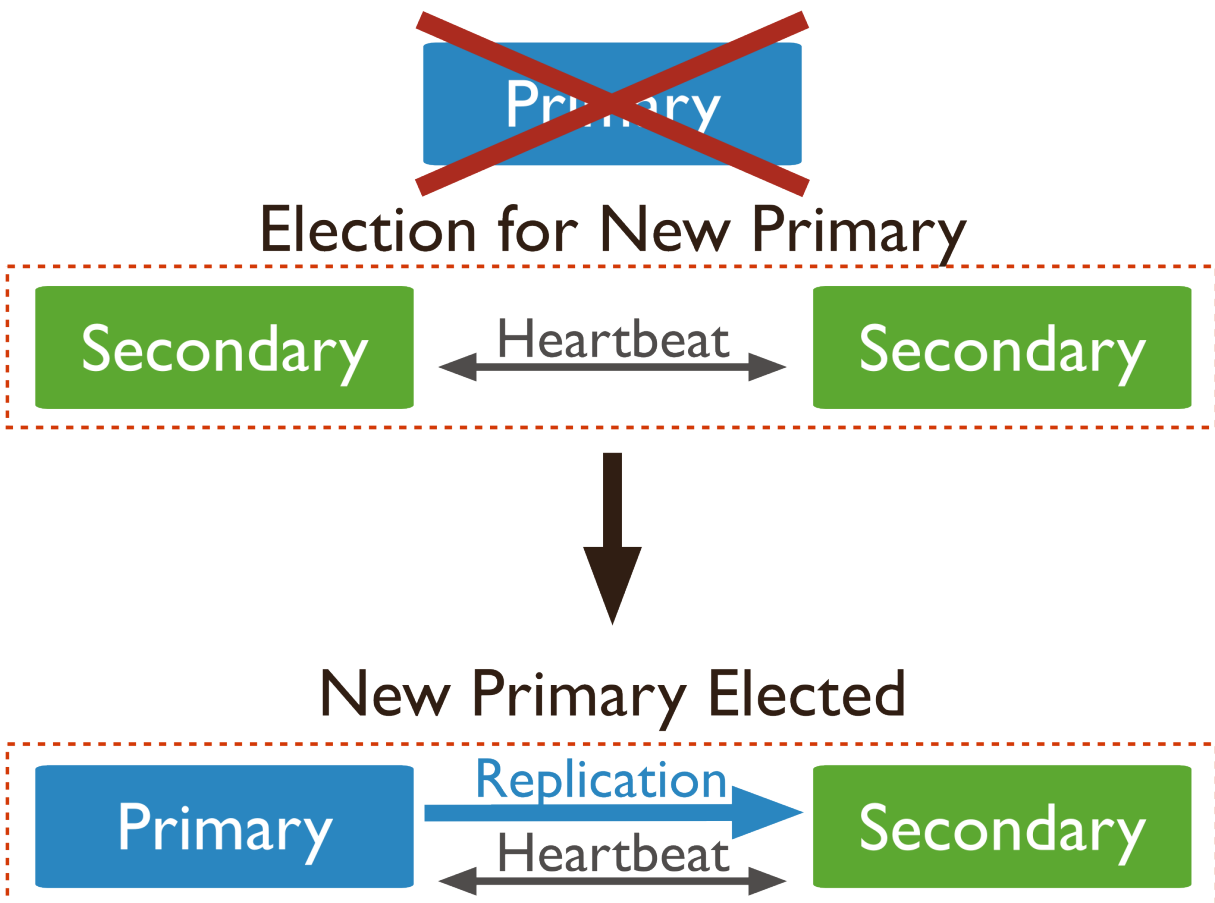
Replication Election Protocol

New in version 3.2: MongoDB introduces a version 1 of the replication protocol (`protocolVersion: 1` (page 109)) to reduce replica set failover time and accelerates the detection of multiple simultaneous primaries. New replica sets will, by default, use `protocolVersion: 1` (page 109). Previous versions of MongoDB use version 0 of the protocol.

Heartbeats

Replica set members send heartbeats (pings) to each other every two seconds. If a heartbeat does not return within 10 seconds, the other members mark the delinquent member as inaccessible.

⁶ Replica sets remove “rollback” data when needed without intervention. Administrators must apply or discard rollback data manually.



Member Priority

After a replica set has a stable primary, the election algorithm will make a “best-effort” attempt to have the secondary with the highest `priority` (page 111) available call an election. Higher priority secondaries call elections relatively sooner than lower priority secondaries; however, a lower priority node can still be elected as primary for brief periods of time, even if a higher priority secondary is available. Replica set members will continue to call elections until the highest priority available member becomes primary.

Members with a priority value of 0 cannot become primary and do not seek election. For details, see *Priority 0 Replica Set Members* (page 14).

Network Partitions

Network partitions affect the formation of a majority for an election. If a primary steps down and neither portion of the replica set has a majority the set will **not** elect a new primary. The replica set becomes read-only.

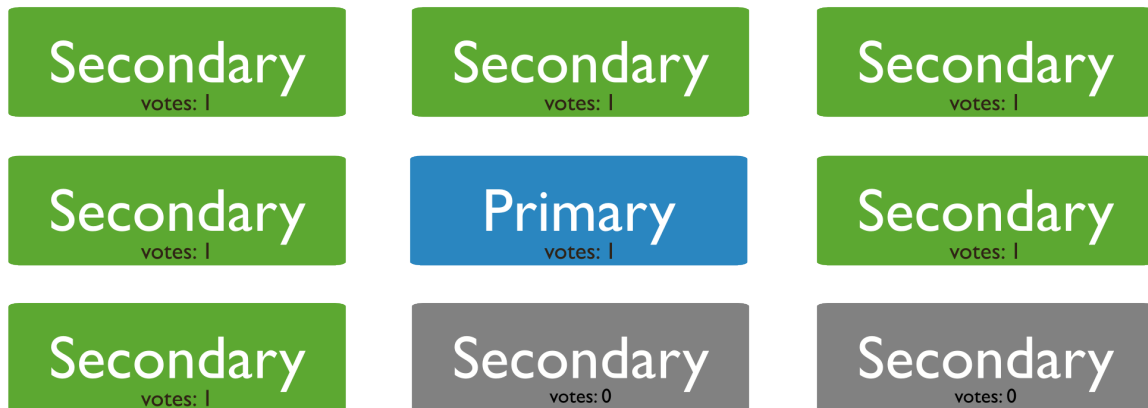
To avoid this situation, place a majority of instances in one data center and a minority of instances in any other data centers combined.

Non-Voting Members

Although non-voting members do not vote in elections, these members hold copies of the replica set’s data and can accept read operations from client applications.

Because a replica set can have up to 50 members, but only 7 voting members, non-voting members allow a replica set to have more than seven members.

For instance, the following nine-member replica set has seven voting members and two non-voting members.



A non-voting member has a `members[n].votes` (page 112) setting equal to 0 in its member configuration:

```
{
  "_id" : <num>
  "host" : <hostname:port>,
  "votes" : 0
}
```

Important: Do **not** alter the number of votes to control which members will become primary. Instead, modify the

`members[n].priority` (page 111) option. *Only* alter the number of votes in exceptional cases. For example, to permit more than seven members.

To configure a non-voting member, see *Configure Non-Voting Replica Set Member* (page 77).

2.3.2 Rollbacks During Replica Set Failover

On this page

- [Collect Rollback Data](#) (page 31)
- [Avoid Replica Set Rollbacks](#) (page 31)
- [Rollback Limitations](#) (page 32)

A rollback reverts write operations on a former *primary* when the member rejoins its *replica set* after a *failover*. A rollback is necessary only if the primary had accepted write operations that the *secondaries* had **not** successfully replicated before the primary stepped down. When the primary rejoins the set as a secondary, it reverts, or “rolls back,” its write operations to maintain database consistency with the other members.

MongoDB attempts to avoid rollbacks, which should be rare. When a rollback does occur, it is often the result of a network partition. Secondaries that can not keep up with the throughput of operations on the former primary, increase the size and impact of the rollback.

A rollback does *not* occur if the write operations replicate to another member of the replica set before the primary steps down *and* if that member remains available and accessible to a majority of the replica set.

Collect Rollback Data

When a rollback does occur, administrators must decide whether to apply or ignore the rollback data. MongoDB writes the rollback data to *BSON* files in the `rollback/` folder under the database’s `dbPath` directory. The names of rollback files have the following form:

```
<database>.<collection>.<timestamp>.bson
```

For example:

```
records.accounts.2011-05-09T18-10-04.0.bson
```

Administrators must apply rollback data manually after the member completes the rollback and returns to secondary status. Use `bsondump` to read the contents of the rollback files. Then use `mongorestore` to apply the changes to the new primary.

Avoid Replica Set Rollbacks

For replica sets, the default `write concern {w: 1}` only provides acknowledgement of write operations on the primary. With the default write concern, data may be rolled back if the primary steps down before the write operations have replicated to any of the secondaries.

To prevent rollbacks of data that have been acknowledged to the client, use `w: majority write concern` to guarantee that the write operations propagate to a majority of the replica set nodes before returning with acknowledgement to the issuing client.

Note:

- Regardless of `write concern`, other clients using "local" (i.e. the default) `readConcern` can see the result of a write operation before the write operation is acknowledged to the issuing client.
 - Clients using "local" (i.e. the default) `readConcern` can read data which may be subsequently *rolled back* (page 31).
-

Rollback Limitations

A `mongod` instance will not rollback more than 300 megabytes of data. If your system must rollback more than 300 megabytes, you must manually intervene to recover the data. If this is the case, the following line will appear in your `mongod` log:

```
[replica set sync] replSet syncThread: 13410 replSet too much data to roll back
```

In this situation, save the data directly or force the member to perform an initial sync. To force initial sync, sync from a “current” member of the set by deleting the content of the `dbPath` directory for the member that requires a larger rollback.

See also:

Replica Set High Availability (page 27) and *Replica Set Elections* (page 28).

2.4 Replica Set Read and Write Semantics

From the perspective of a client application, whether a MongoDB instance is running as a single server (i.e. “standalone”) or a *replica set* is transparent. However, MongoDB provides additional read and write configurations for replica sets.

Note: *Sharded clusters* where the shards are also replica sets provide the same operational semantics with regards to write and read operations.

Write Concern for Replica Sets (page 32) Write concern describes the level of acknowledgement requested from MongoDB for write operations.

Read Preference (page 34) Read preference specifies where (i.e. which members of the replica set) the drivers should direct the read operations.

Read Preference Processes (page 37) Describes the mechanics of read preference.

2.4.1 Write Concern for Replica Sets

On this page

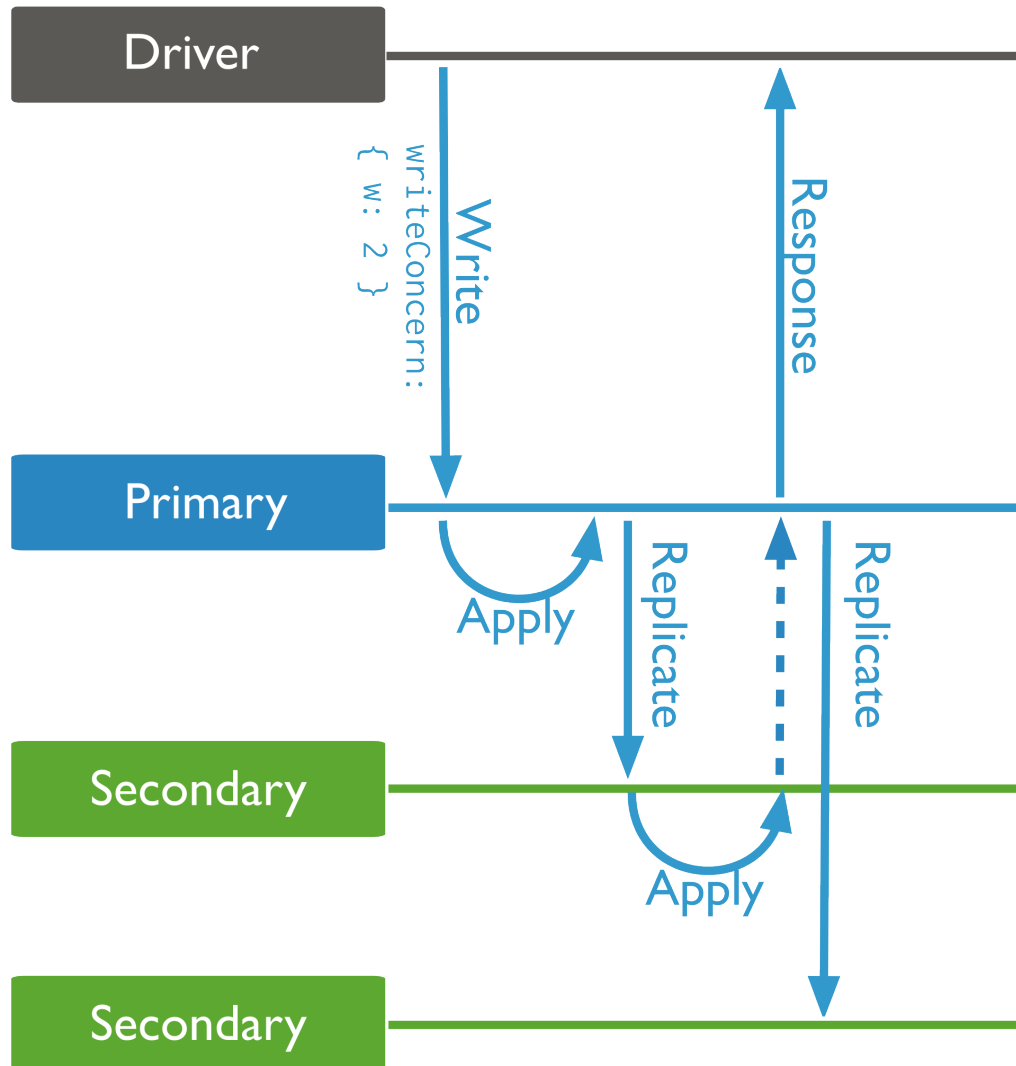
- [Verify Write Operations to Replica Sets](#) (page 33)
- [Modify Default Write Concern](#) (page 34)
- [Custom Write Concerns](#) (page 34)

From the perspective of a client application, whether a MongoDB instance is running as a single server (i.e. “standalone”) or a *replica set* is transparent. However, replica sets offer some configuration options for write.⁷

⁷ *Sharded clusters* where the shards are also replica sets provide the same configuration options with regards to write and read operations.

Verify Write Operations to Replica Sets

For a replica set, the default `write concern` requests acknowledgement only from the primary. You can, however, override this default write concern, such as to confirm write operations on a specified number of the replica set members.



To override the default write concern, specify a write concern with each write operation. For example, the following method includes a write concern that specifies that the method return only after the write propagates to the primary and at least one secondary or the method times out after 5 seconds.

```
db.products.insert(
  { item: "envelopes", qty : 100, type: "Clasp" },
  { writeConcern: { w: 2, wtimeout: 5000 } }
)
```

You can include a timeout threshold for a write concern. This prevents write operations from blocking indefinitely if the write concern is unachievable. For example, if the write concern requires acknowledgement from 4 members

of the replica set and the replica set has only available 3 members, the operation blocks until those members become available. See *wc-wtimeout*.

See also:

write-methods-incompatibility

Modify Default Write Concern

You can modify the default write concern for a replica set by setting the `settings.getLastErrorDefaults` (page 113) setting in the *replica set configuration* (page 108). The following sequence of commands creates a configuration that waits for the write operation to complete on a majority of the voting members before returning:

```
cfg = rs.conf()
cfg.settings = {}
cfg.settings.getLastErrorDefaults = { w: "majority", wtimeout: 5000 }
rs.reconfig(cfg)
```

If you issue a write operation with a specific write concern, the write operation uses its own write concern instead of the default.

See also:

<https://docs.mongodb.org/manual/reference/write-concern>

Custom Write Concerns

You can *tag* (page 88) the members of replica sets and use the resulting tag sets to create custom write concerns. See *Configure Replica Set Tag Sets* (page 88) for information on configuring custom write concerns using tag sets.

2.4.2 Read Preference

On this page

- [Use Cases](#) (page 35)
- [Read Preference Modes](#) (page 36)
- [Tag Sets](#) (page 37)

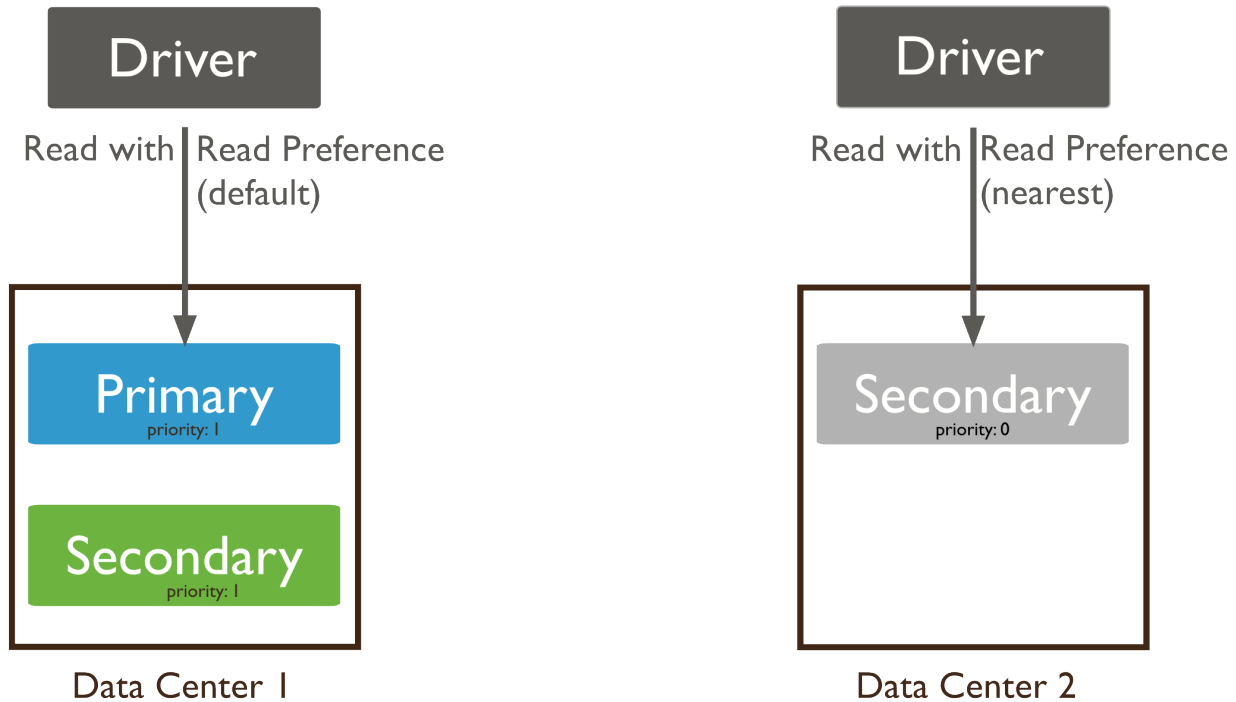
Read preference describes how MongoDB clients route read operations to the members of a *replica set*.

By default, an application directs its read operations to the *primary* member in a *replica set*.

In MongoDB, in a replica set with one primary member⁸,

- With "local" readConcern, reads from the primary reflect the latest writes in absence of a failover;
- With "majority" readConcern, read operations from the primary or the secondaries have *eventual consistency*.

⁸ In *some circumstances* (page 121), two nodes in a replica set may *transiently* believe that they are the primary, but at most, one of them will be able to complete writes with { w: "majority" } write concern. The node that can complete { w: "majority" } writes is the current primary, and the other node is a former primary that has not yet recognized its demotion, typically due to a *network partition*. When this occurs, clients that connect to the former primary may observe stale data despite having requested read preference `primary` (page 120), and new writes to the former primary will eventually roll back.



Important: Exercise care when specifying read preferences: Modes other than `primary` (page 120) may return stale data because with *asynchronous replication* (page 5), data in the secondary may not reflect the most recent write operations. ¹

Note: The read preference does not affect the visibility of data; i.e, clients can see the results of writes before they are acknowledged or have propagated to a majority of replica set members:

- Regardless of `write concern`, other clients using "local" (i.e. the default) `readConcern` can see the result of a write operation before the write operation is acknowledged to the issuing client.
 - Clients using "local" (i.e. the default) `readConcern` can read data which may be subsequently *rolled back* (page 31).
-

Use Cases

Indications

The following are common use cases for using non-`primary` (page 120) read preference modes:

- Running systems operations that do not affect the front-end application.

Note: Read preferences aren't relevant to direct connections to a single `mongod` instance. However, in order to perform read operations on a direct connection to a secondary member of a replica set, you must set a read preference, such as `secondary`.

- Providing local reads for geographically distributed applications.

If you have application servers in multiple data centers, you may consider having a *geographically distributed replica set* (page 26) and using a non primary read preference or the *nearest* (page 120). This allows the client to read from the lowest-latency members, rather than always reading from the primary.

- Maintaining availability during a failover.

Use *primaryPreferred* (page 120) if you want an application to read from the primary under normal circumstances, but to allow stale reads from secondaries when the primary is unavailable. This provides a “read-only mode” for your application during a failover.

Counter-Indications

In general, do *not* use *secondary* (page 120) and *secondaryPreferred* (page 120) to provide extra capacity for reads, because:

- All members of a replica have roughly equivalent write traffic; as a result, secondaries will service reads at roughly the same rate as the primary.
- Replication is asynchronous and there is some amount of delay between a successful write operation and its replication to secondaries. Reading from a secondary can return out-of-date data; reading from different secondaries may result in non-monotonic reads.
- Distributing read operations to secondaries can compromise availability if *any* members of the set become unavailable because the remaining members of the set will need to be able to handle all application requests.
- For queries of sharded collections, for clusters with the *balancer* active, secondaries may return stale results with missing or duplicated data because of incomplete or terminated chunk migrations.

Sharding increases read and write capacity by distributing read and write operations across a group of machines, and is often a better strategy for adding capacity.

See *Read Preference Processes* (page 37) for more information about the internal application of read preferences.

Read Preference Modes

Important: All read preference modes except *primary* (page 120) may return stale data because *secondaries* replicate operations from the primary with some delay.¹ Ensure that your application can tolerate stale data if you choose to use a non-*primary* (page 120) mode.

MongoDB *drivers* support five read preference modes.

Read Preference Mode	Description
<i>primary</i> (page 120)	Default mode. All operations read from the current replica set <i>primary</i> .
<i>primaryPreferred</i> (page 120)	In most situations, operations read from the <i>primary</i> but if it is unavailable, operations read from <i>secondary</i> members.
<i>secondary</i> (page 120)	All operations read from the <i>secondary</i> members of the replica set.
<i>secondaryPreferred</i> (page 120)	In most situations, operations read from <i>secondary</i> members but if no <i>secondary</i> members are available, operations read from the <i>primary</i> .
<i>nearest</i> (page 120)	Operations read from member of the <i>replica set</i> with the least network latency, irrespective of the member’s type.

The syntax for specifying the read preference mode is *specific to the driver and to the idioms of the host language*⁹.

⁹<https://api.mongodb.org/>

Read preference modes are also available to clients connecting to a *sharded cluster* through a `mongos`. The `mongos` instance obeys specified read preferences when connecting to the *replica set* that provides each *shard* in the cluster.

In the `mongo` shell, the `readPref()` cursor method provides access to read preferences.

For more information, see [read preference background](#) (page 34) and [read preference behavior](#) (page 37). See also the [documentation for your driver](#)¹⁰.

Tag Sets

Tag sets allow you to target read operations to specific members of a replica set.

Custom read preferences and write concerns evaluate tag sets in different ways. Read preferences consider the value of a tag when selecting a member to read from. Write concerns ignore the value of a tag to when selecting a member, *except* to consider whether or not the value is unique.

You can specify tag sets with the following read preference modes:

- `primaryPreferred` (page 120)
- `secondary` (page 120)
- `secondaryPreferred` (page 120)
- `nearest` (page 120)

Tags are not compatible with mode `primary` (page 120) and, in general, only apply when *selecting* (page 38) a *secondary* member of a set for a read operation. However, the `nearest` (page 120) read mode, when combined with a tag set, selects the matching member with the lowest network latency. This member may be a primary or secondary.

All interfaces use the same *member selection logic* (page 38) to choose the member to which to direct read operations, basing the choice on read preference mode and tag sets.

For information on configuring tag sets, see the [Configure Replica Set Tag Sets](#) (page 88) tutorial.

For more information on how read preference *modes* (page 120) interact with tag sets, see the [documentation for each read preference mode](#) (page 118).

2.4.3 Read Preference Processes

On this page

- [Member Selection](#) (page 38)
- [Request Association](#) (page 38)
- [Auto-Retry](#) (page 38)
- [Read Preference in Sharded Clusters](#) (page 39)

Changed in version 2.2.

MongoDB drivers use the following procedures to direct operations to replica sets and sharded clusters. To determine how to route their operations, applications periodically update their view of the replica set's state, identifying which members are up or down, which member is *primary*, and verifying the latency to each `mongod` instance.

¹⁰<https://api.mongodb.org/>

Member Selection

Clients, by way of their drivers, and `mongos` instances for sharded clusters, periodically update their view of the replica set's state.

When you select non-`primary` (page 120) read preference, the driver will determine which member to target using the following process:

1. Assembles a list of suitable members, taking into account member type (i.e. secondary, primary, or all members).
2. Excludes members not matching the tag sets, if specified.
3. Determines which suitable member is the closest to the client in absolute terms.
4. Builds a list of members that are within a defined ping distance (in milliseconds) of the “absolute nearest” member.

Applications can configure the threshold used in this stage. The default “acceptable latency” is 15 milliseconds, which you can override in the drivers with their own `secondaryAcceptableLatencyMS` option. For `mongos` you can use the `--localThreshold` or `localPingThresholdMs` runtime options to set this value.

5. Selects a member from these hosts at random. The member receives the read operation.

Drivers can then associate the thread or connection with the selected member. This *request association* (page 38) is configurable by the application. See your `driver` documentation about request association configuration and default behavior.

Request Association

Important: *Request association* is configurable by the application. See your `driver` documentation about request association configuration and default behavior.

Because *secondary* members of a *replica set* may lag behind the current *primary* by different amounts, reads for *secondary* members may reflect data at different points in time. To prevent sequential reads from jumping around in time, the driver **can** associate application threads to a specific member of the set after the first read, thereby preventing reads from other members. The thread will continue to read from the same member until:

- The application performs a read with a different read preference,
- The thread terminates, or
- The client receives a socket exception, as is the case when there's a network error or when the `mongod` closes connections during a *failover*. This triggers a *retry* (page 38), which may be transparent to the application.

When using request association, if the client detects that the set has elected a new *primary*, the driver will discard all associations between threads and members.

Auto-Retry

Connections between MongoDB drivers and `mongod` instances in a *replica set* must balance two concerns:

1. The client should attempt to prefer current results, and any connection should read from the same member of the replica set as much as possible. Requests should prefer *request association* (page 38) (e.g. *pinning*).
2. The client should minimize the amount of time that the database is inaccessible as the result of a connection issue, networking problem, or *failover* in a replica set.

As a result, MongoDB drivers:

- Reuse a connection to a specific `mongod` for as long as possible after establishing a connection to that instance. This connection is *pinned* to this `mongod`.
- Attempt to reconnect to a new member, obeying existing *read preference modes* (page 120), if the connection to `mongod` is lost.

Reconnections are transparent to the application itself. If the connection permits reads from *secondary* members, after reconnecting, the application can receive two sequential reads returning from different secondaries. Depending on the state of the individual secondary member's replication, the documents can reflect the state of your database at different moments.

- Return an error *only* after attempting to connect to three members of the set that match the *read preference mode* (page 120) and *tag set* (page 37). If there are fewer than three members of the set, the client will error after connecting to all existing members of the set.

After this error, the driver selects a new member using the specified read preference mode. In the absence of a specified read preference, the driver uses `primary` (page 120).

- After detecting a failover situation,¹¹ the driver attempts to refresh the state of the replica set as quickly as possible.

Changed in version 3.0.0: `mongos` instances take a slightly different approach. `mongos` instances return connections to secondaries to the connection pool after every request. As a result, the `mongos` reevaluates read preference for every operation.

Read Preference in Sharded Clusters

Changed in version 2.2: Before version 2.2, `mongos` did not support the *read preference mode semantics* (page 120).

In most *sharded clusters*, each shard consists of a *replica set*. As such, read preferences are also applicable. With regard to read preference, read operations in a sharded cluster are identical to unsharded replica sets.

Unlike simple replica sets, in sharded clusters, all interactions with the shards pass from the clients to the `mongos` instances that are actually connected to the set members. `mongos` is then responsible for the application of read preferences, which is transparent to applications.

There are no configuration changes required for full support of read preference modes in sharded environments, as long as the `mongos` is at least version 2.2. All `mongos` maintain their own connection pool to the replica set members. As a result:

- A request without a specified preference has `primary` (page 120), the default, unless, the `mongos` reuses an existing connection that has a different mode set.

To prevent confusion, always explicitly set your read preference mode.

- All `nearest` (page 120) and latency calculations reflect the connection between the `mongos` and the `mongod` instances, not the client and the `mongod` instances.

This produces the desired result, because all results must pass through the `mongos` before returning to the client.

2.5 Replication Processes

Members of a *replica set* replicate data continuously. First, a member uses *initial sync* to capture the data set. Then the member continuously records and applies every operation that modifies the data set. Every member records operations in its *oplog* (page 40), which is a *capped collection*.

¹¹ When a *failover* occurs, all members of the set close all client connections that produce a socket error in the driver. This behavior prevents or minimizes *rollback*.

Replica Set Oplog (page 40) The oplog records all operations that modify the data in the replica set.

Replica Set Data Synchronization (page 41) Secondaries must replicate all changes accepted by the primary. This process is the basis of replica set operations.

2.5.1 Replica Set Oplog

On this page

- [Oplog Size \(page 40\)](#)
- [Workloads that Might Require a Larger Oplog Size \(page 41\)](#)
- [Oplog Status \(page 41\)](#)

The *oplog* (operations log) is a special *capped collection* that keeps a rolling record of all operations that modify the data stored in your databases. MongoDB applies database operations on the *primary* and then records the operations on the primary's oplog. The *secondary* members then copy and apply these operations in an asynchronous process. All replica set members contain a copy of the oplog, in the `local.oplog.rs` (page 115) collection, which allows them to maintain the current state of the database.

To facilitate replication, all replica set members send heartbeats (pings) to all other members. Any member can import oplog entries from any other member.

Whether applied once or multiple times to the target dataset, each operation in the oplog produces the same results, i.e. each operation in the oplog is *idempotent*. For proper replication operations, entries in the oplog must be idempotent:

- initial sync
- post-rollback catch-up
- sharding chunk migrations

Oplog Size

When you start a replica set member for the first time, MongoDB creates an oplog of a default size. The size depends on the architectural details of your operating system.

In most cases, the default oplog size is sufficient. For example, if an oplog is 5% of free disk space and fills up in 24 hours of operations, then secondaries can stop copying entries from the oplog for up to 24 hours without becoming too stale to continue replicating. However, most replica sets have much lower operation volumes, and their oplogs can hold much higher numbers of operations.

Before `mongod` creates an oplog, you can specify its size with the `oplogSizeMB` option. However, after you have started a replica set member for the first time, you can only change the size of the oplog using the [Change the Size of the Oplog \(page 80\)](#) procedure.

By default, the size of the oplog is as follows:

- For 64-bit Linux, Solaris, FreeBSD, and Windows systems, MongoDB allocates 5% of the available free disk space, but will always allocate at least 1 gigabyte and never more than 50 gigabytes.
- For 64-bit OS X systems, MongoDB allocates 183 megabytes of space to the oplog.
- For 32-bit systems, MongoDB allocates about 48 megabytes of space to the oplog.

Workloads that Might Require a Larger Oplog Size

If you can predict your replica set's workload to resemble one of the following patterns, then you might want to create an oplog that is larger than the default. Conversely, if your application predominantly performs reads with a minimal amount of write operations, a smaller oplog may be sufficient.

The following workloads might require a larger oplog size.

Updates to Multiple Documents at Once

The oplog must translate multi-updates into individual operations in order to maintain *idempotency*. This can use a great deal of oplog space without a corresponding increase in data size or disk use.

Deletions Equal the Same Amount of Data as Inserts

If you delete roughly the same amount of data as you insert, the database will not grow significantly in disk use, but the size of the operation log can be quite large.

Significant Number of In-Place Updates

If a significant portion of the workload is updates that do not increase the size of the documents, the database records a large number of operations but does not change the quantity of data on disk.

Oplog Status

To view oplog status, including the size and the time range of operations, issue the `rs.printReplicationInfo()` method. For more information on oplog status, see *Check the Size of the Oplog* (page 104).

Under various exceptional situations, updates to a *secondary's* oplog might lag behind the desired performance time. Use `db.getReplicationInfo()` from a secondary member and the `replication` status output to assess the current state of replication and determine if there is any unintended replication delay.

See *Replication Lag* (page 101) for more information.

2.5.2 Replica Set Data Synchronization

On this page

- [Initial Sync](#) (page 42)
- [Replication](#) (page 42)

In order to maintain up-to-date copies of the shared data set, secondary members of a replica set *sync* or replicate data from other members. MongoDB uses two forms of data synchronization: *initial sync* (page 42) to populate new members with the full data set, and replication to apply ongoing changes to the entire data set.

Initial Sync

Initial sync copies all the data from one member of the replica set to another member. A member uses initial sync when the member has no data, such as when the member is new, or when the member has data but is missing a history of the set's replication.

When you perform an initial sync, MongoDB:

1. Clones all databases. To clone, the `mongod` queries every collection in each source database and inserts all data into its own copies of these collections. At this time, `_id` indexes are also built. The clone process only copies valid data, omitting invalid documents.
2. Applies all changes to the data set. Using the oplog from the source, the `mongod` updates its data set to reflect the current state of the replica set.
3. Builds all indexes on all collections (except `_id` indexes, which were already completed).

When the `mongod` finishes building all index builds, the member can transition to a normal state, i.e. *secondary*.

Changed in version 3.0: When the clone process omits an invalid document from the sync, MongoDB writes a message to the logs that begins with `Cloner: found corrupt document in <collection>`.

To perform an initial sync, see *Resync a Member of a Replica Set* (page 87).

Replication

Secondary members replicate data continuously after the initial sync. Secondary members copy the *oplog* (page 40) from their *sync from* source and apply these operations in an asynchronous process.

In most cases, secondaries sync from the primary. Secondaries may automatically change their *sync from* source if needed based on changes in the ping time and state of other members' replication.

For a member to sync from another, both members must have the same value for the `members[n].buildIndexes` (page 110) setting.

Secondaries avoid syncing from *delayed members* (page 16) and *hidden members* (page 15).

Multithreaded Replication

MongoDB applies write operations in batches using multiple threads to improve concurrency. MongoDB groups batches by namespace (*MMAPv1*) or by document id (*WiredTiger*) and simultaneously applies each group of operations using a different thread. MongoDB always applies write operations to a given document in their original write order.

While applying a batch, MongoDB blocks all read operations. As a result, secondary read queries can never return data that reflect a state that never existed on the primary.

Pre-Fetching Indexes to Improve Replication Throughput

Note: Applies to MMAPv1 only.

With the *MMAPv1* storage engine, MongoDB fetches memory pages that hold affected data and indexes to help improve the performance of applying oplog entries. This *pre-fetch* stage minimizes the amount of time MongoDB holds write locks while applying oplog entries. By default, secondaries will pre-fetch all *indexes*.

Optionally, you can disable all pre-fetching or only pre-fetch the index on the `_id` field. See the `secondaryIndexPrefetch` setting for more information.

2.6 Master Slave Replication

On this page

- [Fundamental Operations](#) (page 43)
- [Run time Master-Slave Configuration](#) (page 44)
- [Security](#) (page 45)
- [Ongoing Administration and Operation of Master-Slave Deployments](#) (page 45)

Important: *Replica sets* (page 9) replace *master-slave* replication for most use cases. If possible, use replica sets rather than master-slave replication for all new production deployments. This documentation remains to support legacy deployments and for archival purposes only.

In addition to providing all the functionality of master-slave deployments, replica sets are also more robust for production use. Master-slave replication preceded replica sets and made it possible to have a large number of non-master (i.e. slave) nodes, as well as to restrict replicated operations to only a single database; however, master-slave replication provides less redundancy and does not automate failover. See *Deploy Master-Slave Equivalent using Replica Sets* (page 45) for a replica set configuration that is equivalent to master-slave replication. If you wish to convert an existing master-slave deployment to a replica set, see *Convert a Master-Slave Deployment to a Replica Set* (page 46).

2.6.1 Fundamental Operations

Initial Deployment

To configure a *master-slave* deployment, start two `mongod` instances: one in master mode, and the other in slave mode.

To start a `mongod` instance in master mode, invoke `mongod` as follows:

```
mongod --master --dbpath /data/masterdb/
```

With the `--master` option, the `mongod` will create a `local.oplog.$main` (page 116) collection, which the “operation log” that queues operations that the slaves will apply to replicate operations from the master. The `--dbpath` is optional.

To start a `mongod` instance in slave mode, invoke `mongod` as follows:

```
mongod --slave --source <masterhostname><:<port>> --dbpath /data/slavedb/
```

Specify the hostname and port of the master instance to the `--source` argument. The `--dbpath` is optional.

For slave instances, MongoDB stores data about the source server in the `local.sources` (page 116) collection.

Configuration Options for Master-Slave Deployments

As an alternative to specifying the `--source` run-time option, can add a document to `local.sources` (page 116) specifying the master instance, as in the following operation in the `mongo` shell:

```
use local
db.sources.find()
db.sources.insert( { host: <masterhostname> <,only: <databasename>> } );
```


In line 1, you switch context to the `local` database. In line 2, the `find()` operation should return no documents, to ensure that there are no documents in the `sources` collection. Finally, line 3 uses `db.collection.insert()` to insert the source document into the `local.sources` (page 116) collection. The model of the `local.sources` (page 116) document is as follows:

host

The `host` field specifies the master `mongod` instance, and holds a resolvable hostname, i.e. IP address, or a name from a `host` file, or preferably a fully qualified domain name.

You can append `<:port>` to the `host` name if the `mongod` is not running on the default 27017 port.

only

Optional. Specify a name of a database. When specified, MongoDB will only replicate the indicated database.

Operational Considerations for Replication with Master Slave Deployments

Master instances store operations in an *oplog* which is a capped collection. As a result, if a slave falls too far behind the state of the master, it cannot “catchup” and must re-sync from scratch. Slave may become out of sync with a master if:

- The slave falls far behind the data updates available from that master.
- The slave stops (i.e. shuts down) and restarts later after the master has overwritten the relevant operations from the master.

When slaves are out of sync, replication stops. Administrators must intervene manually to restart replication. Use the `resync` command. Alternatively, the `--autoresync` allows a slave to restart replication automatically, after ten second pause, when the slave falls out of sync with the master. With `--autoresync` specified, the slave will only attempt to re-sync once in a ten minute period.

To prevent these situations you should specify a larger oplog when you start the *master* instance, by adding the `--oplogSize` option when starting `mongod`. If you do not specify `--oplogSize`, `mongod` will allocate 5% of available disk space on start up to the oplog, with a minimum of 1 GB for 64-bit machines and 50 MB for 32-bit machines.

2.6.2 Run time Master-Slave Configuration

MongoDB provides a number of command line options for `mongod` instances in *master-slave* deployments. See the *Master-Slave Replication Command Line Options* for options.

Diagnosics

On a *master* instance, issue the following operation in the `mongo` shell to return replication status from the perspective of the master:

```
rs.printReplicationInfo()
```

New in version 2.6: `rs.printReplicationInfo()`. For previous versions, use `db.printReplicationInfo()`.

On a *slave* instance, use the following operation in the `mongo` shell to return the replication status from the perspective of the slave:

```
rs.printSlaveReplicationInfo()
```

New in version 2.6: `rs.printSlaveReplicationInfo()`. For previous versions, use `db.printSlaveReplicationInfo()`.

Use the `serverStatus` as in the following operation, to return status of the replication:

```
db.serverStatus( { repl: 1 } )
```

See *server status repl fields* for documentation of the relevant section of output.

2.6.3 Security

When running with authorization enabled, in *master-slave* deployments configure a `keyFile` so that slave `mongod` instances can authenticate and communicate with the master `mongod` instance.

To enable authentication and configure the `keyFile` add the following option to your configuration file:

```
keyFile = /srv/mongodb/keyfile
```

Note: You may chose to set these run-time configuration options using the `--keyFile` option on the command line.

Setting `keyFile` enables authentication and specifies a key file for the `mongod` instances to use when authenticating to each other. The content of the key file is arbitrary but must be the same on all members of the deployment can connect to each other.

The key file must be less one kilobyte in size and may only contain characters in the base64 set. The key file must not have group or “world” permissions on UNIX systems. Use the following command to use the OpenSSL package to generate “random” content for use in a key file:

```
openssl rand -base64 741
```

See also:

<https://docs.mongodb.org/manual/security> for more information about security in MongoDB

2.6.4 Ongoing Administration and Operation of Master-Slave Deployments

Deploy Master-Slave Equivalent using Replica Sets

If you want a replication configuration that resembles *master-slave* replication, using *replica sets* replica sets, consider the following replica configuration document. In this deployment hosts `<master>` and `<slave>` ¹² provide replication that is roughly equivalent to a two-instance master-slave deployment:

```
{
  _id : 'setName',
  members : [
    { _id : 0, host : "<master>", priority : 1 },
    { _id : 1, host : "<slave>", priority : 0, votes : 0 }
  ]
}
```

See *Replica Set Configuration* (page 108) for more information about replica set configurations.

¹² In replica set configurations, the `members[n].host` (page 110) field must hold a resolvable hostname.

Convert a Master-Slave Deployment to a Replica Set

To convert a master-slave deployment to a replica set, restart the current master as a one-member replica set. Then remove the data directories from previous secondaries and add them as new secondaries to the new replica set.

1. To confirm that the current instance is master, run:

```
db.isMaster()
```

This should return a document that resembles the following:

```
{
  "ismaster" : true,
  "maxBsonObjectSize" : 16777216,
  "maxMessageSizeBytes" : 48000000,
  "localTime" : ISODate("2013-07-08T20:15:13.664Z"),
  "ok" : 1
}
```

2. Shut down the `mongod` processes on the master and all slave(s), using the following command while connected to each instance:

```
db.adminCommand({shutdown : 1, force : true})
```

3. Back up your `/data/db` directories, in case you need to revert to the master-slave deployment.
4. Start the former master with the `--replSet` option, as in the following:

```
mongod --replSet <setname>
```

5. Connect to the `mongod` with the `mongo` shell, and initiate the replica set with the following command:

```
rs.initiate()
```

When the command returns, you will have successfully deployed a one-member replica set. You can check the status of your replica set at any time by running the following command:

```
rs.status()
```

You can now follow the [convert a standalone to a replica set](#) (page 65) tutorial to deploy your replica set, picking up from the [Expand the Replica Set](#) (page 65) section.

Failing over to a Slave (Promotion)

To permanently failover from an unavailable or damaged *master* (A in the following example) to a *slave* (B):

1. Shut down A.
2. Stop `mongod` on B.
3. Back up and move all data files that begin with `local` on B from the `dbPath`.

Warning: Removing `local.*` is irrevocable and cannot be undone. Perform this step with extreme caution.

4. Restart `mongod` on B with the `--master` option.

Note: This is a one time operation, and is not reversible. A cannot become a slave of B until it completes a full resync.

Inverting Master and Slave

If you have a *master* (A) and a *slave* (B) and you would like to reverse their roles, follow this procedure. The procedure assumes A is healthy, up-to-date and available.

If A is not healthy but the hardware is okay (power outage, server crash, etc.), skip steps 1 and 2 and in step 8 replace all of A's files with B's files in step 8.

If A is not healthy and the hardware is not okay, replace A with a new machine. Also follow the instructions in the previous paragraph.

To invert the master and slave in a deployment:

1. Halt writes on A using the *fsync* command.
2. Make sure B is up to date with the state of A.
3. Shut down B.
4. Back up and move all data files that begin with `local` on B from the `dbPath` to remove the existing `local.sources` data.

Warning: Removing `local.*` is irrevocable and cannot be undone. Perform this step with extreme caution.

5. Start B with the `--master` option.
6. Do a write on B, which primes the *oplog* to provide a new sync start point.
7. Shut down B. B will now have a new set of data files that start with `local`.
8. Shut down A and replace all files in the `dbPath` of A that start with `local` with a copy of the files in the `dbPath` of B that begin with `local`.
Considering compressing the `local` files from B while you copy them, as they may be quite large.
9. Start B with the `--master` option.
10. Start A with all the usual slave options, but include *fastsync*.

Creating a Slave from an Existing Master's Disk Image

If you can stop write operations to the *master* for an indefinite period, you can copy the data files from the master to the new *slave* and then start the slave with `--fastsync`.

Warning: Be careful with `--fastsync`. If the data on both instances is **not** identical, a discrepancy will exist forever.

fastsync is a way to start a slave by starting with an existing master disk image/backup. This option declares that the administrator guarantees the image is correct and completely up-to-date with that of the master. If you have a full and complete copy of data from a master you can use this option to avoid a full synchronization upon starting the slave.

Creating a Slave from an Existing Slave's Disk Image

You can just copy the other *slave's* data file snapshot without any special options. Only take data snapshots when:

- a `mongod` process is down, or
- when the `mongod` is locked using `db.fsyncLock()` for MMAPv1 or WiredTiger storage engine.

Changed in version 3.2: Starting in MongoDB 3.2, `db.fsyncLock()` can ensure that the data files do not change for MongoDB instances using either the MMAPv1 or the WiredTiger storage engine, thus providing consistency for the purposes of creating backups.

In previous MongoDB version, `db.fsyncLock()` *cannot* guarantee a consistent set of files for low-level backups (e.g. via file copy `cp`, `scp`, `tar`) for WiredTiger.

Resyncing a Slave that is too Stale to Recover

Slaves asynchronously apply write operations from the *master* that the slaves poll from the master's *oplog*. The *oplog* is finite in length, and if a slave is too far behind, a full resync will be necessary. To resync the slave, connect to a slave using the `mongo` and issue the `resync` command:

```
use admin
db.runCommand( { resync: 1 } )
```

This forces a full resync of all data (which will be very slow on a large database). You can achieve the same effect by stopping `mongod` on the slave, deleting the entire content of the `dbPath` on the slave, and restarting the `mongod`.

Slave Chaining

Slaves cannot be “chained.” They must all connect to the *master* directly.

If a slave attempts “slave from” another slave you will see the following line in the `mongod` log of the shell:

```
assertion 13051 tailable cursor requested on non capped collection ns:local.oplog.$main
```

Correcting a Slave's Source

To change a *slave's* source, manually modify the slave's `local.sources` (page 116) collection.

Example

Consider the following: If you accidentally set an incorrect hostname for the slave's *source*, as in the following example:

```
mongod --slave --source prod.mississippi
```

You can correct this, by restarting the slave without the `--slave` and `--source` arguments:

```
mongod
```

Connect to this `mongod` instance using the `mongo` shell and update the `local.sources` (page 116) collection, with the following operation sequence:

```
use local

db.sources.update( { host : "prod.mississippi" },
                  { $set : { host : "prod.mississippi.example.net" } } )
```

Restart the slave with the correct command line arguments or with no `--source` option. After configuring `local.sources` (page 116) the first time, the `--source` will have no subsequent effect. Therefore, both of the following invocations are correct:

```
mongod --slave --source prod.mississippi.example.net
```

or

```
mongod --slave
```

The slave now polls data from the correct *master*.

Replica Set Tutorials

The administration of *replica sets* includes the initial deployment of the set, adding and removing members to a set, and configuring the operational parameters and properties of the set. Administrators generally need not intervene in failover or replication processes as MongoDB automates these functions. In the exceptional situations that require manual interventions, the tutorials in these sections describe processes such as resyncing a member. The tutorials in this section form the basis for all replica set administration.

***Replica Set Deployment Tutorials* (page 52)** Instructions for deploying replica sets, as well as adding and removing members from an existing replica set.

***Deploy a Replica Set* (page 52)** Configure a three-member replica set for production systems.

***Convert a Standalone to a Replica Set* (page 65)** Convert an existing standalone `mongod` instance into a three-member replica set.

***Add Members to a Replica Set* (page 66)** Add a new member to an existing replica set.

***Remove Members from Replica Set* (page 69)** Remove a member from a replica set.

Continue reading from *Replica Set Deployment Tutorials* (page 52) for additional tutorials of related to setting up replica set deployments.

***Member Configuration Tutorials* (page 71)** Tutorials that describe the process for configuring replica set members.

***Adjust Priority for Replica Set Member* (page 72)** Change the precedence given to a replica set members in an election for primary.

***Prevent Secondary from Becoming Primary* (page 73)** Make a secondary member ineligible for election as primary.

***Configure a Hidden Replica Set Member* (page 74)** Configure a secondary member to be invisible to applications in order to support significantly different usage, such as a dedicated backups.

Continue reading from *Member Configuration Tutorials* (page 71) for more tutorials that describe replica set configuration.

***Replica Set Maintenance Tutorials* (page 80)** Procedures and tasks for common operations on active replica set deployments.

***Change the Size of the Oplog* (page 80)** Increase the size of the *oplog* which logs operations. In most cases, the default oplog size is sufficient.

***Resync a Member of a Replica Set* (page 87)** Sync the data on a member. Either perform initial sync on a new member or resync the data on an existing member that has fallen too far behind to catch up by way of normal replication.

***Force a Member to Become Primary* (page 85)** Force a replica set member to become primary.

Change Hostnames in a Replica Set (page 96) Update the replica set configuration to reflect changes in members' hostnames.

Continue reading from *Replica Set Maintenance Tutorials* (page 80) for descriptions of additional replica set maintenance procedures.

Troubleshoot Replica Sets (page 101) Describes common issues and operational challenges for replica sets. For additional diagnostic information, see <https://docs.mongodb.org/manual/faq/diagnostics>.

3.1 Replica Set Deployment Tutorials

The following tutorials provide information in deploying replica sets.

Deploy a Replica Set (page 52) Configure a three-member replica set for production systems.

Deploy a Replica Set for Testing and Development (page 55) Configure a three-member replica set for either development or testing systems.

Deploy a Geographically Redundant Replica Set (page 58) Create a geographically redundant replica set to protect against location-centered availability limitations (e.g. network and power interruptions).

Add an Arbiter to Replica Set (page 64) Add an arbiter give a replica set an odd number of voting members to prevent election ties.

Convert a Standalone to a Replica Set (page 65) Convert an existing standalone `mongod` instance into a three-member replica set.

Add Members to a Replica Set (page 66) Add a new member to an existing replica set.

Remove Members from Replica Set (page 69) Remove a member from a replica set.

Replace a Replica Set Member (page 70) Update the replica set configuration when the hostname of a member's corresponding `mongod` instance has changed.

3.1.1 Deploy a Replica Set

On this page

- [Overview](#) (page 52)
- [Requirements](#) (page 53)
- [Considerations When Deploying a Replica Set](#) (page 53)
- [Procedure](#) (page 54)

This tutorial describes how to create a three-member *replica set* from three existing `mongod` instances running with `access control` disabled.

To deploy a replica set with enabled `access control`, see *deploy-repl-set-with-auth*. If you wish to deploy a replica set from a single MongoDB instance, see *Convert a Standalone to a Replica Set* (page 65). For more information on replica set deployments, see the *Replication* (page 1) and *Replica Set Deployment Architectures* (page 19) documentation.

Overview

Three member *replica sets* provide enough redundancy to survive most network partitions and other system failures. These sets also have sufficient capacity for many distributed read operations. Replica sets should always have an odd

number of members. This ensures that *elections* (page 28) will proceed smoothly. For more about designing replica sets, see *the Replication overview* (page 3).

The basic procedure is to start the `mongod` instances that will become members of the replica set, configure the replica set itself, and then add the `mongod` instances to it.

Requirements

For production deployments, you should maintain as much separation between members as possible by hosting the `mongod` instances on separate machines. When using virtual machines for production deployments, you should place each `mongod` instance on a separate host server serviced by redundant power circuits and redundant network paths.

Before you can deploy a replica set, you must install MongoDB on each system that will be part of your *replica set*. If you have not already installed MongoDB, see the *installation tutorials*.

Before creating your replica set, you should verify that your network configuration allows all possible connections between each member. For a successful replica set deployment, every member must be able to connect to every other member. For instructions on how to check your connection, see *Test Connections Between all Members* (page 103).

Considerations When Deploying a Replica Set

Architecture

In a production, deploy each member of the replica set to its own machine and if possible bind to the standard MongoDB port of 27017. Use the `bind_ip` option to ensure that MongoDB listens for connections from applications on configured addresses.

For a geographically distributed replica sets, ensure that the majority of the set's `mongod` instances reside in the primary site.

See *Replica Set Deployment Architectures* (page 19) for more information.

Connectivity

Ensure that network traffic can pass between all members of the set and all clients in the network securely and efficiently. Consider the following:

- Establish a virtual private network. Ensure that your network topology routes all traffic between members within a single site over the local area network.
- Configure access control to prevent connections from unknown clients to the replica set.
- Configure networking and firewall rules so that incoming and outgoing packets are permitted only on the default MongoDB port and only from within your deployment.

Finally ensure that each member of a replica set is accessible by way of resolvable DNS or hostnames. You should either configure your DNS names appropriately or set up your systems' `/etc/hosts` file to reflect this configuration.

Configuration

Specify the run time configuration on each system in a `configuration` file stored in `/etc/mongod.conf` or a related location. Create the directory where MongoDB stores data files before deploying MongoDB.

For more information about the run time options used above and other configuration options, see <https://docs.mongodb.org/manual/reference/configuration-options>.

Procedure

The following procedure outlines the steps to deploy a replica set when access control is disabled.

Step 1: Start each member of the replica set with the appropriate options.

For each member, start a `mongod` and specify the replica set name through the `replSet` option. Specify any other parameters specific to your deployment. For replication-specific parameters, see *cli-mongod-replica-set*.

If your application connects to more than one replica set, each set should have a distinct name. Some drivers group replica set connections by replica set name.

The following example specifies the replica set name through the `--replSet` command-line option:

```
mongod --replSet "rs0"
```

You can also specify the replica set name in the configuration file. To start `mongod` with a configuration file, specify the file with the `--config` option:

```
mongod --config $HOME/.mongodb/config
```

In production deployments, you can configure a *init script* to manage this process. Init scripts are beyond the scope of this document.

Step 2: Connect a mongo shell to a replica set member.

For example, to connect to a `mongod` running on localhost on the default port of 27017, simply issue:

```
mongo
```

Step 3: Initiate the replica set.

Use `rs.initiate()` on *one and only one* member of the replica set:

```
rs.initiate()
```

MongoDB initiates a set that consists of the current member and that uses the default replica set configuration.

Step 4: Verify the initial replica set configuration.

Use `rs.conf()` to display the *replica set configuration object* (page 108):

```
rs.conf()
```

The replica set configuration object resembles the following:

```
{
  "_id" : "rs0",
  "version" : 1,
  "members" : [
    {
      "_id" : 1,
      "host" : "mongodb0.example.net:27017"
    }
  ]
}
```

```
    ]
}
```

Step 5: Add the remaining members to the replica set.

Add the remaining members with the `rs.add()` method. You must be connected to the *primary* to add members to a replica set.

`rs.add()` can, in some cases, trigger an election. If the *mongod* you are connected to becomes a *secondary*, you need to connect the *mongo* shell to the new primary to continue adding new replica set members. Use `rs.status()` to identify the primary in the replica set.

The following example adds two members:

```
rs.add("mongodb1.example.net")
rs.add("mongodb2.example.net")
```

When complete, you have a fully functional replica set. The new replica set will elect a *primary*.

Step 6: Check the status of the replica set.

Use the `rs.status()` operation:

```
rs.status()
```

See also:

[deploy-repl-set-with-auth](#)

3.1.2 Deploy a Replica Set for Testing and Development

On this page

- [Overview](#) (page 55)
- [Requirements](#) (page 56)
- [Considerations](#) (page 56)
- [Procedure](#) (page 56)

This procedure describes deploying a replica set in a development or test environment. For a production deployment, refer to the [Deploy a Replica Set](#) (page 52) tutorial.

This tutorial describes how to create a three-member *replica set* from three existing `mongod` instances running with `access control` disabled.

To deploy a replica set with enabled `access control`, see [deploy-repl-set-with-auth](#). If you wish to deploy a replica set from a single MongoDB instance, see [Convert a Standalone to a Replica Set](#) (page 65). For more information on replica set deployments, see the [Replication](#) (page 1) and [Replica Set Deployment Architectures](#) (page 19) documentation.

Overview

Three member *replica sets* provide enough redundancy to survive most network partitions and other system failures. These sets also have sufficient capacity for many distributed read operations. Replica sets should always have an odd

number of members. This ensures that *elections* (page 28) will proceed smoothly. For more about designing replica sets, see *the Replication overview* (page 3).

The basic procedure is to start the `mongod` instances that will become members of the replica set, configure the replica set itself, and then add the `mongod` instances to it.

Requirements

For test and development systems, you can run your `mongod` instances on a local system, or within a virtual instance.

Before you can deploy a replica set, you must install MongoDB on each system that will be part of your *replica set*. If you have not already installed MongoDB, see the *installation tutorials*.

Before creating your replica set, you should verify that your network configuration allows all possible connections between each member. For a successful replica set deployment, every member must be able to connect to every other member. For instructions on how to check your connection, see *Test Connections Between all Members* (page 103).

Considerations

Replica Set Naming

Important: These instructions should only be used for test or development deployments.

The examples in this procedure create a new replica set named `rs0`.

If your application connects to more than one replica set, each set should have a distinct name. Some drivers group replica set connections by replica set name.

You will begin by starting three `mongod` instances as members of a replica set named `rs0`.

Procedure

1. Create the necessary data directories for each member by issuing a command similar to the following:

```
mkdir -p /srv/mongodb/rs0-0 /srv/mongodb/rs0-1 /srv/mongodb/rs0-2
```

This will create directories called “rs0-0”, “rs0-1”, and “rs0-2”, which will contain the instances’ database files.

2. Start your `mongod` instances in their own shell windows by issuing the following commands:

First member:

```
mongod --port 27017 --dbpath /srv/mongodb/rs0-0 --replSet rs0 --smallfiles --oplogSize 128
```

Second member:

```
mongod --port 27018 --dbpath /srv/mongodb/rs0-1 --replSet rs0 --smallfiles --oplogSize 128
```

Third member:

```
mongod --port 27019 --dbpath /srv/mongodb/rs0-2 --replSet rs0 --smallfiles --oplogSize 128
```

This starts each instance as a member of a replica set named `rs0`, each running on a distinct port, and specifies the path to your data directory with the `--dbpath` setting. If you are already using the suggested ports, select different ports.

The `--smallfiles` and `--oplogSize` settings reduce the disk space that each `mongod` instance uses. This is ideal for testing and development deployments as it prevents overloading your machine. For more information on these and other configuration options, see <https://docs.mongodb.org/manual/reference/configuration-options>.

3. Connect to one of your `mongod` instances through the `mongo` shell. You will need to indicate which instance by specifying its port number. For the sake of simplicity and clarity, you may want to choose the first one, as in the following command;

```
mongo --port 27017
```

4. In the `mongo` shell, use `rs.initiate()` to initiate the replica set. You can create a replica set configuration object in the `mongo` shell environment, as in the following example:

```
rsconf = {
  _id: "rs0",
  members: [
    {
      _id: 0,
      host: "<hostname>:27017"
    }
  ]
}
```

replacing `<hostname>` with your system's hostname, and then pass the `rsconf` file to `rs.initiate()` as follows:

```
rs.initiate( rsconf )
```

5. Display the current *replica configuration* (page 108) by issuing the following command:

```
rs.conf()
```

The replica set configuration object resembles the following

```
{
  "_id" : "rs0",
  "version" : 4,
  "members" : [
    {
      "_id" : 1,
      "host" : "localhost:27017"
    }
  ]
}
```

6. In the `mongo` shell connected to the *primary*, add the second and third `mongod` instances to the replica set using the `rs.add()` method. Replace `<hostname>` with your system's hostname in the following examples:

```
rs.add("<hostname>:27018")
rs.add("<hostname>:27019")
```

When complete, you should have a fully functional replica set. The new replica set will elect a *primary*.

Check the status of your replica set at any time with the `rs.status()` operation.

See also:

The documentation of the following shell functions for more information:

- `rs.initiate()`

- `rs.conf()`
- `rs.reconfig()`
- `rs.add()`

You may also consider the [simple setup script](#)¹ as an example of a basic automatically-configured replica set.

Refer to *Replica Set Read and Write Semantics* (page 32) for a detailed explanation of read and write semantics in MongoDB.

3.1.3 Deploy a Geographically Redundant Replica Set

On this page

- [Overview](#) (page 58)
- [Considerations](#) (page 58)
- [Prerequisites](#) (page 58)
- [Procedures](#) (page 59)

Overview

This tutorial outlines the process for deploying a *replica set* with members in multiple locations. The tutorial addresses three-member sets, four-member sets, and sets with more than four members.

For appropriate background, see *Replication* (page 1) and *Replica Set Deployment Architectures* (page 19). For related tutorials, see *Deploy a Replica Set* (page 52) and *Add Members to a Replica Set* (page 66).

Considerations

While *replica sets* provide basic protection against single-instance failure, replica sets whose members are all located in a single facility are susceptible to errors in that facility. Power outages, network interruptions, and natural disasters are all issues that can affect replica sets whose members are collocated. To protect against these classes of failures, deploy a replica set with one or more members in a geographically distinct facility or data center to provide redundancy.

Prerequisites

In general, the requirements for any geographically redundant replica set are as follows:

- Ensure that a majority of the *voting members* (page 30) are within a primary facility, “Site A”. This includes *priority 0 members* (page 14) and *arbiters* (page 17). Deploy other members in secondary facilities, “Site B”, “Site C”, etc., to provide additional copies of the data. See *Determine the Distribution of Members* (page 20) for more information on the voting requirements for geographically redundant replica sets.
- If you deploy a replica set with an even number of members, deploy an *arbiter* (page 17) on Site A. The arbiter must be on site A to keep the majority there.

For instance, for a three-member replica set you need two instances in a Site A, and one member in a secondary facility, Site B. Site A should be the same facility or very close to your primary application infrastructure (i.e. application servers, caching layer, users, etc.)

¹<https://github.com/mongodb/mongo-snippets/blob/master/replication/simple-setup.py>

A four-member replica set should have at least two members in Site A, with the remaining members in one or more secondary sites, as well as a single *arbiter* in Site A.

For all configurations in this tutorial, deploy each replica set member on a separate system. Although you may deploy more than one replica set member on a single system, doing so reduces the redundancy and capacity of the replica set. Such deployments are typically for testing purposes and beyond the scope of this tutorial.

This tutorial assumes you have installed MongoDB on each system that will be part of your replica set. If you have not already installed MongoDB, see the *installation tutorials*.

Procedures

General Considerations

Architecture

In a production, deploy each member of the replica set to its own machine and if possible bind to the standard MongoDB port of 27017. Use the `bind_ip` option to ensure that MongoDB listens for connections from applications on configured addresses.

For a geographically distributed replica sets, ensure that the majority of the set's `mongod` instances reside in the primary site.

See *Replica Set Deployment Architectures* (page 19) for more information.

Connectivity

Ensure that network traffic can pass between all members of the set and all clients in the network securely and efficiently. Consider the following:

- Establish a virtual private network. Ensure that your network topology routes all traffic between members within a single site over the local area network.
- Configure access control to prevent connections from unknown clients to the replica set.
- Configure networking and firewall rules so that incoming and outgoing packets are permitted only on the default MongoDB port and only from within your deployment.

Finally ensure that each member of a replica set is accessible by way of resolvable DNS or hostnames. You should either configure your DNS names appropriately or set up your systems' `/etc/hosts` file to reflect this configuration.


Configuration

Specify the run time configuration on each system in a `configuration` file stored in `/etc/mongod.conf` or a related location. Create the directory where MongoDB stores data files before deploying MongoDB.

For more information about the run time options used above and other configuration options, see <https://docs.mongodb.org/manual/reference/configuration-options>.

Deploy a Geographically Redundant Three-Member Replica Set

Step 1: Start each member of the replica set with the appropriate options. For each member, start a `mongod` and specify the replica set name through the `replSet` option. Specify any other parameters specific to your deployment. For replication-specific parameters, see *cli-mongod-replica-set*.



If your application connects to more than one replica set, each set should have a distinct name. Some drivers group replica set connections by replica set name.

The following example specifies the replica set name through the `--replSet` command-line option:

```
mongod --replSet "rs0"
```

You can also specify the replica set name in the configuration file. To start `mongod` with a configuration file, specify the file with the `--config` option:

```
mongod --config $HOME/.mongodb/config
```

In production deployments, you can configure a *init script* to manage this process. Init scripts are beyond the scope of this document.

Step 2: Connect a mongo shell to a replica set member. For example, to connect to a `mongod` running on localhost on the default port of 27017, simply issue:

```
mongo
```

Step 3: Initiate the replica set. Use `rs.initiate()` on *one and only one* member of the replica set:

```
rs.initiate()
```

MongoDB initiates a set that consists of the current member and that uses the default replica set configuration.

Step 4: Verify the initial replica set configuration. Use `rs.conf()` to display the *replica set configuration object* (page 108):

```
rs.conf()
```

The replica set configuration object resembles the following:

```
{
  "_id" : "rs0",
  "version" : 1,
  "members" : [
    {
      "_id" : 1,
```

```

        "host" : "mongodb0.example.net:27017"
    }
]
}

```

Step 5: Add the remaining members to the replica set. Add the remaining members with the `rs.add()` method. You must be connected to the *primary* to add members to a replica set.

`rs.add()` can, in some cases, trigger an election. If the *mongod* you are connected to becomes a *secondary*, you need to connect the *mongo* shell to the new primary to continue adding new replica set members. Use `rs.status()` to identify the primary in the replica set.

The following example adds two members:

```

rs.add("mongodb1.example.net")
rs.add("mongodb2.example.net")

```

When complete, you have a fully functional replica set. The new replica set will elect a *primary*.

Step 6: Configure the outside member as *priority 0* members. Configure the member located in Site B (in this example, `mongodb2.example.net`) as a *priority 0* member (page 14).

1. View the replica set configuration to determine the `members` (page 110) array position for the member. Keep in mind the array position is not the same as the `_id`:

```
rs.conf()
```

2. Copy the replica set configuration object to a variable (to `cfg` in the example below). Then, in the variable, set the correct priority for the member. Then pass the variable to `rs.reconfig()` to update the replica set configuration.

For example, to set priority for the third member in the array (i.e., the member at position 2), issue the following sequence of commands:

```

cfg = rs.conf()
cfg.members[2].priority = 0
rs.reconfig(cfg)

```

Note: The `rs.reconfig()` shell method can force the current primary to step down, causing an election. When the primary steps down, all clients will disconnect. This is the intended behavior. While most elections complete within a minute, always make sure any replica configuration changes occur during scheduled maintenance periods.

After these commands return, you have a geographically redundant three-member replica set.

Step 7: Check the status of the replica set. Use the `rs.status()` operation:

```
rs.status()
```

Deploy a Geographically Redundant Four-Member Replica Set

A geographically redundant four-member deployment has two additional considerations:

- One host (e.g. `mongodb4.example.net`) must be an *arbiter*. This host can run on a system that is also used for an application server or on the same machine as another MongoDB process.

- You must decide how to distribute your systems. There are three possible architectures for the four-member replica set:
 - Three members in Site A, one *priority 0 member* (page 14) in Site B, and an arbiter in Site A.
 - Two members in Site A, two *priority 0 members* (page 14) in Site B, and an arbiter in Site A.
 - Two members in Site A, one priority 0 member in Site B, one priority 0 member in Site C, and an arbiter in site A.

In most cases, the first architecture is preferable because it is the least complex.

To deploy a geographically redundant four-member set:

Step 1: Start each member of the replica set with the appropriate options. For each member, start a `mongod` and specify the replica set name through the `replSet` option. Specify any other parameters specific to your deployment. For replication-specific parameters, see *cli-mongod-replica-set*.

If your application connects to more than one replica set, each set should have a distinct name. Some drivers group replica set connections by replica set name.

The following example specifies the replica set name through the `--replSet` command-line option:

```
mongod --replSet "rs0"
```

You can also specify the replica set name in the configuration file. To start `mongod` with a configuration file, specify the file with the `--config` option:

```
mongod --config $HOME/.mongodb/config
```

In production deployments, you can configure a *init script* to manage this process. Init scripts are beyond the scope of this document.

Step 2: Connect a mongo shell to a replica set member. For example, to connect to a `mongod` running on localhost on the default port of 27017, simply issue:

```
mongo
```

Step 3: Initiate the replica set. Use `rs.initiate()` on *one and only one* member of the replica set:

```
rs.initiate()
```

MongoDB initiates a set that consists of the current member and that uses the default replica set configuration.

Step 4: Verify the initial replica set configuration. Use `rs.conf()` to display the *replica set configuration object* (page 108):

```
rs.conf()
```

The replica set configuration object resembles the following:

```
{
  "_id" : "rs0",
  "version" : 1,
  "members" : [
    {
      "_id" : 1,
```

```

    "host" : "mongodb0.example.net:27017"
  }
]
}

```

Step 5: Add the remaining members to the replica set. Use `rs.add()` in a mongo shell connected to the current primary. The commands should resemble the following:

```

rs.add("mongodb1.example.net")
rs.add("mongodb2.example.net")
rs.add("mongodb3.example.net")

```

When complete, you should have a fully functional replica set. The new replica set will elect a *primary*.

Step 6: Add the arbiter. In the same shell session, issue the following command to add the arbiter (e.g. `mongodb4.example.net`):

```
rs.addArb("mongodb4.example.net")
```

Step 7: Configure outside members as *priority 0 members*. Configure each member located outside of Site A (e.g. `mongodb3.example.net`) as a *priority 0 member* (page 14).

1. View the replica set configuration to determine the `members` (page 110) array position for the member. Keep in mind the array position is not the same as the `_id`:

```
rs.conf()
```

2. Copy the replica set configuration object to a variable (to `cfg` in the example below). Then, in the variable, set the correct priority for the member. Then pass the variable to `rs.reconfig()` to update the replica set configuration.

For example, to set priority for the third member in the array (i.e., the member at position 2), issue the following sequence of commands:

```

cfg = rs.conf()
cfg.members[2].priority = 0
rs.reconfig(cfg)

```

Note: The `rs.reconfig()` shell method can force the current primary to step down, causing an election. When the primary steps down, all clients will disconnect. This is the intended behavior. While most elections complete within a minute, always make sure any replica configuration changes occur during scheduled maintenance periods.

After these commands return, you have a geographically redundant four-member replica set.

Step 8: Check the status of the replica set. Use the `rs.status()` operation:

```
rs.status()
```

Deploy a Geographically Redundant Set with More than Four Members

The above procedures detail the steps necessary for deploying a geographically redundant replica set. Larger replica set deployments follow the same steps, but have additional considerations:

- Never deploy more than seven voting members.
- If you have an even number of members, use *the procedure for a four-member set* (page 61)). Ensure that a single facility, “Site A”, always has a majority of the members by deploying the *arbiter* in that site. For example, if a set has six members, deploy at least three voting members in addition to the arbiter in Site A, and the remaining members in alternate sites.
- If you have an odd number of members, use *the procedure for a three-member set* (page 59). Ensure that a single facility, “Site A” always has a majority of the members of the set. For example, if a set has five members, deploy three members within Site A and two members in other facilities.
- If you have a majority of the members of the set *outside* of Site A and the network partitions to prevent communication between sites, the current primary in Site A will step down, even if none of the members outside of Site A are eligible to become primary.

3.1.4 Add an Arbiter to Replica Set

On this page

- [Considerations](#) (page 64)
- [Add an Arbiter](#) (page 64)

Arbiters are `mongod` instances that are part of a *replica set* but do not hold data. Arbiters participate in *elections* (page 28) in order to break ties. If a replica set has an even number of members, add an arbiter.

Arbiters have minimal resource requirements and do not require dedicated hardware. You can deploy an arbiter on an application server or a monitoring host.

Important: Do not run an arbiter on the same system as a member of the replica set.

Considerations

An arbiter does not store data, but until the arbiter’s `mongod` process is added to the replica set, the arbiter will act like any other `mongod` process and start up with a set of data files and with a full-sized *journal*.

To minimize the default creation of data, set the following in the arbiter’s configuration file:

- `journal.enabled` to `false`

Warning: Never set `journal.enabled` to `false` on a data-bearing node.

- `smallFiles` to `true`

These settings are specific to arbiters. Do not set `journal.enabled` to `false` on a data-bearing node. Similarly, do not set `smallFiles` unless specifically indicated.

Add an Arbiter

1. Create a data directory (e.g. `dbPath`) for the arbiter. The `mongod` instance uses the directory for configuration data. The directory *will not* hold the data set. For example, create the `/data/arb` directory:

```
mkdir /data/arb
```

2. Start the arbiter. Specify the data directory and the replica set name. The following, starts an arbiter using the `/data/arb` dbPath for the `rs` replica set:

```
mongod --port 30000 --dbpath /data/arb --replSet rs
```

3. Connect to the primary and add the arbiter to the replica set. Use the `rs.addArb()` method, as in the following example:

```
rs.addArb("m1.example.net:30000")
```

This operation adds the arbiter running on port 30000 on the `m1.example.net` host.

3.1.5 Convert a Standalone to a Replica Set

On this page

- [Procedure](#) (page 65)

This tutorial describes the process for converting a *standalone* `mongod` instance into a three-member *replica set*. Use standalone instances for testing and development, but always use replica sets in production. To install a standalone instance, see the *installation tutorials*.

To deploy a replica set without using a pre-existing `mongod` instance, see *Deploy a Replica Set* (page 52).

Procedure

1. Shut down the *standalone* `mongod` instance.
2. Restart the instance. Use the `--replSet` option to specify the name of the new replica set.

For example, the following command starts a standalone instance as a member of a new replica set named `rs0`. The command uses the standalone's existing database path of `/srv/mongodb/db0`:

```
mongod --port 27017 --dbpath /srv/mongodb/db0 --replSet rs0
```

If your application connects to more than one replica set, each set should have a distinct name. Some drivers group replica set connections by replica set name.

For more information on configuration options, see <https://docs.mongodb.org/manual/reference/configuration/> and the `mongod` manual page.

3. Connect to the `mongod` instance.
4. Use `rs.initiate()` to initiate the new replica set:

```
rs.initiate()
```

The replica set is now operational.

To view the replica set configuration, use `rs.conf()`. To check the status of the replica set, use `rs.status()`.

Expand the Replica Set

Add additional replica set members by doing the following:

1. On two distinct systems, start two new standalone `mongod` instances. For information on starting a standalone instance, see the *installation tutorial* specific to your environment.
2. On your connection to the original `mongod` instance (the former standalone instance), issue a command in the following form for each new instance to add to the replica set:

```
rs.add("<hostname>:<port>")
```

Replace `<hostname>` and `<port>` with the resolvable hostname and port of the `mongod` instance to add to the set. For more information on adding a host to a replica set, see *Add Members to a Replica Set* (page 66).

Sharding Considerations

If the new replica set is part of a *sharded cluster*, change the shard host information in the *config database* by doing the following:

1. Connect to one of the sharded cluster's `mongos` instances and issue a command in the following form:

```
db.getSiblingDB("config").shards.save( {_id: "<name>", host: "<replica-set>/<member,><member,><..."
```

Replace `<name>` with the name of the shard. Replace `<replica-set>` with the name of the replica set. Replace `<member,><member,><>` with the list of the members of the replica set.

2. Restart all `mongos` instances. If possible, restart all components of the replica sets (i.e., all `mongos` and all shard `mongod` instances).

3.1.6 Add Members to a Replica Set

On this page

- [Overview](#) (page 66)
- [Requirements](#) (page 67)
- [Procedures](#) (page 67)

Overview

This tutorial explains how to add an additional member to an existing *replica set*. For background on replication deployment patterns, see the *Replica Set Deployment Architectures* (page 19) document.

Maximum Voting Members

A replica set can have a maximum of seven *voting members* (page 28). To add a member to a replica set that already has seven voting members, you must either add the member as a *non-voting member* (page 30) or remove a vote from an existing member.

Init Scripts

In production deployments you can configure a *init script* to manage member processes.

Existing Members

You can use these procedures to add new members to an existing set. You can also use the same procedure to “re-add” a removed member. If the removed member’s data is still relatively recent, it can recover and catch up easily.

Data Files

If you have a backup or snapshot of an existing member, you can move the data files (e.g. the `dbPath` directory) to a new system and use them to quickly initiate a new member. The files must be:

- A valid copy of the data files from a member of the same replica set. See <https://docs.mongodb.org/manual/tutorial/backup-with-filesystem-snapshots> document for more information.

Important: Always use filesystem snapshots to create a copy of a member of the existing replica set. **Do not** use `mongodump` and `mongorestore` to seed a new replica set member.

- More recent than the oldest operation in the *primary’s oplog*. The new member must be able to become current by applying operations from the primary’s oplog.

Requirements

1. An active replica set.
2. A new MongoDB system capable of supporting your data set, accessible by the active replica set through the network.

Otherwise, use the MongoDB *installation tutorial* and the *Deploy a Replica Set* (page 52) tutorials.

Procedures

Prepare the Data Directory

Before adding a new member to an existing *replica set*, prepare the new member’s *data directory* using one of the following strategies:

- Make sure the new member’s data directory *does not* contain data. The new member will copy the data from an existing member.

If the new member is in a *recovering* state, it must exit and become a *secondary* before MongoDB can copy all data as part of the replication process. This process takes time but does not require administrator intervention.

- Manually copy the data directory from an existing member. The new member becomes a secondary member and will catch up to the current state of the replica set. Copying the data over may shorten the amount of time for the new member to become current.

Ensure that you can copy the data directory to the new member and begin replication within the *window allowed by the oplog* (page 40). Otherwise, the new instance will have to perform an initial sync, which completely resynchronizes the data, as described in *Resync a Member of a Replica Set* (page 87).

Use `rs.printReplicationInfo()` to check the current state of replica set members with regards to the oplog.

For background on replication deployment patterns, see the *Replica Set Deployment Architectures* (page 19) document.

Add a Member to an Existing Replica Set

1. Start the new `mongod` instance. Specify the data directory and the replica set name. The following example specifies the `/srv/mongodb/db0` data directory and the `rs0` replica set:

```
mongod --dbpath /srv/mongodb/db0 --replSet rs0
```

Take note of the host name and port information for the new `mongod` instance.

For more information on configuration options, see the `mongod` manual page.

Optional

You can specify the data directory and replica set in the `mongod.conf` configuration file, and start the `mongod` with the following command:

```
mongod --config /etc/mongod.conf
```

2. Connect to the replica set's primary.

You can only add members while connected to the primary. If you do not know which member is the primary, log into any member of the replica set and issue the `db.isMaster()` command.

3. Use `rs.add()` to add the new member to the replica set. For example, to add a member at host `mongodb3.example.net`, issue the following command:

```
rs.add("mongodb3.example.net")
```

You can include the port number, depending on your setup:

```
rs.add("mongodb3.example.net:27017")
```

4. Verify that the member is now part of the replica set. Call the `rs.conf()` method, which displays the *replica set configuration* (page 108):

```
rs.conf()
```

To view replica set status, issue the `rs.status()` method. For a description of the status fields, see <https://docs.mongodb.org/manual/reference/command/replSetGetStatus>.

Configure and Add a Member

You can add a member to a replica set by passing to the `rs.add()` method a `members` (page 110) document. The document must be in the form of a `members` (page 110) document. These documents define a replica set member in the same form as the *replica set configuration document* (page 108).

Important: Specify a value for the `_id` field of the `members` (page 110) document. MongoDB does not automatically populate the `_id` field in this case. Finally, the `members` (page 110) document must declare the `host` value. All other fields are optional.

Example

To add a member with the following configuration:

- an `_id` of 1.
- a `hostname` and `port` number of `mongodb3.example.net:27017`.
- a `priority` value within the replica set of 0.

- a configuration as hidden,

Issue the following:

```
rs.add({_id: 1, host: "mongodb3.example.net:27017", priority: 0, hidden: true})
```

3.1.7 Remove Members from Replica Set

On this page

- [Remove a Member Using `rs.remove\(\)`](#) (page 69)
- [Remove a Member Using `rs.reconfig\(\)`](#) (page 69)

To remove a member of a *replica set* use either of the following procedures.

Remove a Member Using `rs.remove()`

1. Shut down the `mongod` instance for the member you wish to remove. To shut down the instance, connect using the mongo shell and the `db.shutdownServer()` method.
2. Connect to the replica set's current *primary*. To determine the current primary, use `db.isMaster()` while connected to any member of the replica set.
3. Use `rs.remove()` in either of the following forms to remove the member:

```
rs.remove("mongodb3.example.net:27017")
rs.remove("mongodb3.example.net")
```

MongoDB disconnects the shell briefly as the replica set elects a new primary. The shell then automatically reconnects. The shell displays a `DBClientCursor::init call() failed` error even though the command succeeds.

Remove a Member Using `rs.reconfig()`

To remove a member you can manually edit the *replica set configuration document* (page 108), as described here.

1. Shut down the `mongod` instance for the member you wish to remove. To shut down the instance, connect using the mongo shell and the `db.shutdownServer()` method.
2. Connect to the replica set's current *primary*. To determine the current primary, use `db.isMaster()` while connected to any member of the replica set.
3. Issue the `rs.conf()` method to view the current configuration document and determine the position in the `members` array of the member to remove:

Example

`mongod_C.example.net` is in position 2 of the following configuration file:

```
{
  "_id" : "rs",
  "version" : 7,
  "members" : [
    {
      "_id" : 0,
```

```
        "host" : "mongod_A.example.net:27017"
    },
    {
        "_id" : 1,
        "host" : "mongod_B.example.net:27017"
    },
    {
        "_id" : 2,
        "host" : "mongod_C.example.net:27017"
    }
]
}
```

4. Assign the current configuration document to the variable `cfg`:

```
cfg = rs.conf()
```

5. Modify the `cfg` object to remove the member.
-

Example

To remove `mongod_C.example.net:27017` use the following JavaScript operation:

```
cfg.members.splice(2,1)
```

6. Overwrite the replica set configuration document with the new configuration by issuing the following:

```
rs.reconfig(cfg)
```

As a result of `rs.reconfig()` the shell will disconnect while the replica set renegotiates which member is primary. The shell displays a `DBClientCursor::init call() failed` error even though the command succeeds, and will automatically reconnect.

7. To confirm the new configuration, issue `rs.conf()`.

For the example above the output would be:

```
{
  "_id" : "rs",
  "version" : 8,
  "members" : [
    {
      "_id" : 0,
      "host" : "mongod_A.example.net:27017"
    },
    {
      "_id" : 1,
      "host" : "mongod_B.example.net:27017"
    }
  ]
}
```

3.1.8 Replace a Replica Set Member

On this page

- [Operation](#) (page 71)
- [Example](#) (page 71)

If you need to change the hostname of a replica set member without changing the configuration of that member or the set, you can use the operation outlined in this tutorial. For example if you must re-provision systems or rename hosts, you can use this pattern to minimize the scope of that change.

Operation

To change the hostname for a replica set member modify the `members[n].host` (page 110) field. The value of `members[n]._id` (page 110) field will not change when you reconfigure the set.

See [Replica Set Configuration](#) (page 108) and `rs.reconfig()` for more information.

Note: Any replica set configuration change can trigger the current *primary* to step down, which forces an *election* (page 28). During the election, the current shell session and clients connected to this replica set disconnect, which produces an error even when the operation succeeds.

Example

To change the hostname to `mongo2.example.net` for the replica set member configured at `members[0]`, issue the following sequence of commands:

```
cfg = rs.conf()
cfg.members[0].host = "mongo2.example.net"
rs.reconfig(cfg)
```

3.2 Member Configuration Tutorials

The following tutorials provide information in configuring replica set members to support specific operations, such as to provide dedicated backups, to support reporting, or to act as a cold standby.

Adjust Priority for Replica Set Member (page 72) Change the precedence given to a replica set members in an election for primary.

Prevent Secondary from Becoming Primary (page 73) Make a secondary member ineligible for election as primary.

Configure a Hidden Replica Set Member (page 74) Configure a secondary member to be invisible to applications in order to support significantly different usage, such as a dedicated backups.

Configure a Delayed Replica Set Member (page 76) Configure a secondary member to keep a delayed copy of the data set in order to provide a rolling backup.

Configure Non-Voting Replica Set Member (page 77) Create a secondary member that keeps a copy of the data set but does not vote in an election.

Convert a Secondary to an Arbiter (page 78) Convert a secondary to an arbiter.

3.2.1 Adjust Priority for Replica Set Member

On this page

- [Overview](#) (page 72)
- [Considerations](#) (page 72)
- [Procedure](#) (page 72)

Overview

The priority settings of replica set members affect the outcomes of *elections* (page 28) for primary. Use this setting to ensure that some members are more likely to become primary and that others can never become primary.

The value of the member's `members[n].priority` (page 111) setting determines the member's priority in elections. The higher the number, the higher the priority.

Considerations

To modify priorities, you update the `members` (page 110) array in the replica configuration object. The array index begins with 0. Do **not** confuse this index value with the value of the replica set member's `members[n]._id` (page 110) field in the array.

The value of `members[n].priority` (page 111) can be any floating point (i.e. decimal) number between 0 and 1000. The default value for the `members[n].priority` (page 111) field is 1.

To block a member from seeking election as primary, assign it a priority of 0. *Hidden members* (page 15), *delayed members* (page 16), and *arbiters* (page ??) all have `members[n].priority` (page 111) set to 0.

Adjust priority during a scheduled maintenance window. Reconfiguring priority can force the current primary to step down, leading to an election. Before an election the primary closes all open *client* connections.

Procedure

Step 1: Copy the replica set configuration to a variable.

In the mongo shell, use `rs.conf()` to retrieve the replica set configuration and assign it to a variable. For example:

```
cfg = rs.conf()
```

Step 2: Change each member's priority value.

Change each member's `members[n].priority` (page 111) value, as configured in the `members` (page 110) array.

```
cfg.members[0].priority = 0.5
cfg.members[1].priority = 2
cfg.members[2].priority = 2
```

This sequence of operations modifies the value of `cfg` to set the priority for the first three members defined in the `members` (page 110) array.

Step 3: Assign the replica set the new configuration.

Use `rs.reconfig()` to apply the new configuration.

```
rs.reconfig(cfg)
```

This operation updates the configuration of the replica set using the configuration defined by the value of `cfg`.

3.2.2 Prevent Secondary from Becoming Primary**On this page**

- [Overview](#) (page 73)
- [Considerations](#) (page 73)
- [Procedure](#) (page 73)
- [Related Documents](#) (page 74)

Overview

In a replica set, by default all *secondary* members are eligible to become primary through the election process. You can use the `priority` to affect the outcome of these elections by making some members more likely to become primary and other members less likely or unable to become primary.

Secondaries that cannot become primary are also unable to trigger elections. In all other respects these secondaries are identical to other secondaries.

To prevent a *secondary* member from ever becoming a *primary* in a *failover*, assign the secondary a priority of 0, as described here. For a detailed description of secondary-only members and their purposes, see [Priority 0 Replica Set Members](#) (page 14).

Considerations

When updating the replica configuration object, access the replica set members in the `members` (page 110) array with the **array index**. The array index begins with 0. Do **not** confuse this index value with the value of the `members[n]._id` (page 110) field in each document in the `members` (page 110) array.

Note: MongoDB does not permit the current *primary* to have a priority of 0. To prevent the current primary from again becoming a primary, you must first step down the current primary using `rs.stepDown()`.

Procedure

This tutorial uses a sample replica set with 5 members.

Warning:

- The `rs.reconfig()` shell method can force the current primary to step down, which causes an *election* (page 28). When the primary steps down, the `mongod` closes all client connections. While this typically takes 10-20 seconds, try to make these changes during scheduled maintenance periods.
- To successfully reconfigure a replica set, a majority of the members must be accessible. If your replica set has an even number of members, add an *arbiter* (page 64) to ensure that members can quickly obtain a majority of votes in an election for primary.

Step 1: Retrieve the current replica set configuration.

The `rs.conf()` method returns a *replica set configuration document* (page 108) that contains the current configuration for a replica set.

In a mongo shell connected to a primary, run the `rs.conf()` method and assign the result to a variable:

```
cfg = rs.conf()
```

The returned document contains a `members` (page 110) field which contains an array of member configuration documents, one document for each member of the replica set.

Step 2: Assign priority value of 0.

To prevent a secondary member from becoming a primary, update the secondary member's `members[n].priority` (page 111) to 0.

To assign a priority value to a member of the replica set, access the member configuration document using the array index. In this tutorial, the secondary member to change corresponds to the configuration document found at position 2 of the `members` (page 110) array.

```
cfg.members[2].priority = 0
```

The configuration change does not take effect until you reconfigure the replica set.

Step 3: Reconfigure the replica set.

Use `rs.reconfig()` method to reconfigure the replica set with the updated replica set configuration document.

Pass the `cfg` variable to the `rs.reconfig()` method:

```
rs.reconfig(cfg)
```

Related Documents

- `members[n].priority` (page 111)
- *Adjust Priority for Replica Set Member* (page 72)
- *Replica Set Reconfiguration*
- *Replica Set Elections* (page 28)

3.2.3 Configure a Hidden Replica Set Member

On this page

- *Considerations* (page 75)
- *Examples* (page 75)
- *Related Documents* (page 76)

Hidden members are part of a *replica set* but cannot become *primary* and are invisible to client applications. Hidden members may vote in *elections* (page 28). For a more information on hidden members and their uses, see *Hidden Replica Set Members* (page 15).

Considerations

The most common use of hidden nodes is to support *delayed members* (page 16). If you only need to prevent a member from becoming primary, configure a *priority 0 member* (page 14).

If the `settings.chainingAllowed` (page 113) setting allows secondary members to sync from other secondaries, MongoDB by default prefers non-hidden members over hidden members when selecting a sync target. MongoDB will only choose hidden members as a last resort. If you want a secondary to sync from a hidden member, use the `replSetSyncFrom` database command to override the default sync target. See the documentation for `replSetSyncFrom` before using the command.

See also:

Manage Chained Replication (page 95)

Changed in version 2.0: For *sharded clusters* running with replica sets before 2.0, if you reconfigured a member as hidden, you *had* to restart `mongos` to prevent queries from reaching the hidden member.

Examples

Member Configuration Document

To configure a secondary member as hidden, set its `members[n].priority` (page 111) value to `0` and set its `members[n].hidden` (page 111) value to `true` in its member configuration:

```
{
  "_id" : <num>
  "host" : <hostname:port>,
  "priority" : 0,
  "hidden" : true
}
```

Configuration Procedure

The following example hides the secondary member currently at the index `0` in the `members` (page 110) array. To configure a *hidden member*, use the following sequence of operations in a `mongo` shell connected to the primary, specifying the member to configure by its array index in the `members` (page 110) array:

```
cfg = rs.conf()
cfg.members[0].priority = 0
cfg.members[0].hidden = true
rs.reconfig(cfg)
```

After re-configuring the set, this secondary member has a priority of `0` so that it cannot become primary and is hidden. The other members in the set will not advertise the hidden member in the `isMaster` or `db.isMaster()` output.

When updating the replica configuration object, access the replica set members in the `members` (page 110) array with the **array index**. The array index begins with `0`. Do **not** confuse this index value with the value of the `members[n]._id` (page 110) field in each document in the `members` (page 110) array.

Warning:

- The `rs.reconfig()` shell method can force the current primary to step down, which causes an *election* (page 28). When the primary steps down, the `mongod` closes all client connections. While this typically takes 10-20 seconds, try to make these changes during scheduled maintenance periods.
- To successfully reconfigure a replica set, a majority of the members must be accessible. If your replica set has an even number of members, add an *arbiter* (page 64) to ensure that members can quickly obtain a majority of votes in an election for primary.

Related Documents

- [Replica Set Reconfiguration](#)
- [Replica Set Elections](#) (page 28)
- [Read Preference](#) (page 34)

3.2.4 Configure a Delayed Replica Set Member

On this page

- [Example](#) (page 76)
- [Related Documents](#) (page 77)

To configure a delayed secondary member, set its `members[n].priority` (page 111) value to 0, its `members[n].hidden` (page 111) value to `true`, and its `members[n].slaveDelay` (page 112) value to the number of seconds to delay.

Important: The length of the secondary `members[n].slaveDelay` (page 112) must fit within the window of the `oplog`. If the `oplog` is shorter than the `members[n].slaveDelay` (page 112) window, the delayed member cannot successfully replicate operations.

When you configure a delayed member, the delay applies both to replication and to the member's *oplog*. For details on delayed members and their uses, see [Delayed Replica Set Members](#) (page 16).

Example

The following example sets a 1-hour delay on a secondary member currently at the index 0 in the `members` (page 110) array. To set the delay, issue the following sequence of operations in a `mongo` shell connected to the primary:

```
cfg = rs.conf()
cfg.members[0].priority = 0
cfg.members[0].hidden = true
cfg.members[0].slaveDelay = 3600
rs.reconfig(cfg)
```

After the replica set reconfigures, the delayed secondary member cannot become *primary* and is hidden from applications. The `members[n].slaveDelay` (page 112) value delays both replication and the member's *oplog* by 3600 seconds (1 hour).

When updating the replica configuration object, access the replica set members in the `members` (page 110) array with the **array index**. The array index begins with 0. Do **not** confuse this index value with the value of the `members[n]._id` (page 110) field in each document in the `members` (page 110) array.

Warning:

- The `rs.reconfig()` shell method can force the current primary to step down, which causes an *election* (page 28). When the primary steps down, the `mongod` closes all client connections. While this typically takes 10-20 seconds, try to make these changes during scheduled maintenance periods.
- To successfully reconfigure a replica set, a majority of the members must be accessible. If your replica set has an even number of members, add an *arbiter* (page 64) to ensure that members can quickly obtain a majority of votes in an election for primary.

Related Documents

- `members[n].slaveDelay` (page 112)
- *Replica Set Reconfiguration*
- *Oplog Size* (page 40)
- *Change the Size of the Oplog* (page 80) tutorial
- *Replica Set Elections* (page 28)

3.2.5 Configure Non-Voting Replica Set Member**On this page**

- [Example](#) (page 77)
- [Related Documents](#) (page 78)

Non-voting members allow you to add additional members for read distribution beyond the maximum seven voting members. To configure a member as non-voting, set its `members[n].votes` (page 112) value to 0.

Example

To disable the ability to vote in elections for the fourth, fifth, and sixth replica set members, use the following command sequence in the `mongo` shell connected to the primary. You identify each replica set member by its array index in the `members` (page 110) array:

```
cfg = rs.conf()
cfg.members[3].votes = 0
cfg.members[4].votes = 0
cfg.members[5].votes = 0
rs.reconfig(cfg)
```

This sequence gives 0 votes to the fourth, fifth, and sixth members of the set according to the order of the `members` (page 110) array in the output of `rs.conf()`. This setting allows the set to elect these members as *primary* but does not allow them to vote in elections. Place voting members so that your designated primary or primaries can reach a majority of votes in the event of a network partition.

When updating the replica configuration object, access the replica set members in the `members` (page 110) array with the **array index**. The array index begins with 0. Do **not** confuse this index value with the value of the `members[n]._id` (page 110) field in each document in the `members` (page 110) array.

Warning:

- The `rs.reconfig()` shell method can force the current primary to step down, which causes an *election* (page 28). When the primary steps down, the `mongod` closes all client connections. While this typically takes 10-20 seconds, try to make these changes during scheduled maintenance periods.
- To successfully reconfigure a replica set, a majority of the members must be accessible. If your replica set has an even number of members, add an *arbiter* (page 64) to ensure that members can quickly obtain a majority of votes in an election for primary.

In general and when possible, all members should have only 1 vote. This prevents intermittent ties, deadlocks, or the wrong members from becoming primary. Use `members[n].priority` (page 111) to control which members are more likely to become primary.

Related Documents

- `members[n].votes` (page 112)
- *Replica Set Reconfiguration*
- *Replica Set Elections* (page 28)

3.2.6 Convert a Secondary to an Arbiter

On this page

- [Convert Secondary to Arbiter and Reuse the Port Number](#) (page 78)
- [Convert Secondary to Arbiter Running on a New Port Number](#) (page 79)

If you have a *secondary* in a *replica set* that no longer needs to hold data but that needs to remain in the set to ensure that the set can *elect a primary* (page 28), you may convert the secondary to an *arbiter* (page ??) using either procedure in this tutorial. Both procedures are operationally equivalent:

- You may operate the arbiter on the same port as the former secondary. In this procedure, you must shut down the secondary and remove its data before restarting and reconfiguring it as an arbiter.

For this procedure, see *Convert Secondary to Arbiter and Reuse the Port Number* (page 78).

- Run the arbiter on a new port. In this procedure, you can reconfigure the server as an arbiter before shutting down the instance running as a secondary.

For this procedure, see *Convert Secondary to Arbiter Running on a New Port Number* (page 79).

Convert Secondary to Arbiter and Reuse the Port Number

1. If your application is connecting directly to the secondary, modify the application so that MongoDB queries don't reach the secondary.
2. Shut down the secondary.
3. Remove the *secondary* from the *replica set* by calling the `rs.remove()` method. Perform this operation while connected to the current *primary* in the `mongo` shell:

```
rs.remove("<hostname>:<port>")
```

4. Verify that the replica set no longer includes the secondary by calling the `rs.conf()` method in the mongo shell:

```
rs.conf()
```

5. Move the secondary's data directory to an archive folder. For example:

```
mv /data/db /data/db-old
```

Optional

You may remove the data instead.

6. Create a new, empty data directory to point to when restarting the `mongod` instance. You can reuse the previous name. For example:

```
mkdir /data/db
```

7. Restart the `mongod` instance for the secondary, specifying the port number, the empty data directory, and the replica set. You can use the same port number you used before. Issue a command similar to the following:

```
mongod --port 27021 --dbpath /data/db --replSet rs
```

8. In the mongo shell convert the secondary to an arbiter using the `rs.addArb()` method:

```
rs.addArb("<hostname><:port>")
```

9. Verify the arbiter belongs to the replica set by calling the `rs.conf()` method in the mongo shell.

```
rs.conf()
```

The arbiter member should include the following:

```
"arbiterOnly" : true
```

Convert Secondary to Arbiter Running on a New Port Number

1. If your application is connecting directly to the secondary or has a connection string referencing the secondary, modify the application so that MongoDB queries don't reach the secondary.
2. Create a new, empty data directory to be used with the new port number. For example:

```
mkdir /data/db-temp
```

3. Start a new `mongod` instance on the new port number, specifying the new data directory and the existing replica set. Issue a command similar to the following:

```
mongod --port 27021 --dbpath /data/db-temp --replSet rs
```

4. In the mongo shell connected to the current primary, convert the new `mongod` instance to an arbiter using the `rs.addArb()` method:

```
rs.addArb("<hostname><:port>")
```

5. Verify the arbiter has been added to the replica set by calling the `rs.conf()` method in the mongo shell.

```
rs.conf()
```

The arbiter member should include the following:

```
"arbiterOnly" : true
```

6. Shut down the secondary.
7. Remove the *secondary* from the *replica set* by calling the `rs.remove()` method in the mongo shell:

```
rs.remove("<hostname><:port>")
```

8. Verify that the replica set no longer includes the old secondary by calling the `rs.conf()` method in the mongo shell:

```
rs.conf()
```

9. Move the secondary's data directory to an archive folder. For example:

```
mv /data/db /data/db-old
```

Optional

You may remove the data instead.

3.3 Replica Set Maintenance Tutorials

The following tutorials provide information in maintaining existing replica sets.

***Change the Size of the Oplog* (page 80)** Increase the size of the *oplog* which logs operations. In most cases, the default oplog size is sufficient.

***Perform Maintenance on Replica Set Members* (page 83)** Perform maintenance on a member of a replica set while minimizing downtime.

***Force a Member to Become Primary* (page 85)** Force a replica set member to become primary.

***Resync a Member of a Replica Set* (page 87)** Sync the data on a member. Either perform initial sync on a new member or resync the data on an existing member that has fallen too far behind to catch up by way of normal replication.

***Configure Replica Set Tag Sets* (page 88)** Assign tags to replica set members for use in targeting read and write operations to specific members.

***Reconfigure a Replica Set with Unavailable Members* (page 92)** Reconfigure a replica set when a majority of replica set members are down or unreachable.

***Manage Chained Replication* (page 95)** Disable or enable chained replication. Chained replication occurs when a secondary replicates from another secondary instead of the primary.

***Change Hostnames in a Replica Set* (page 96)** Update the replica set configuration to reflect changes in members' hostnames.

***Configure a Secondary's Sync Target* (page 99)** Specify the member that a secondary member synchronizes from.

3.3.1 Change the Size of the Oplog

On this page

- [Overview](#) (page 81)
- [Procedure](#) (page 81)

The *oplog* exists internally as a *capped collection*, so you cannot modify its size in the course of normal operations. In most cases the *default oplog size* (page 40) is an acceptable size; however, in some situations you may need a larger or smaller oplog. For example, you might need to change the oplog size if your applications perform large numbers of multi-updates or deletes in short periods of time.

This tutorial describes how to resize the oplog. For a detailed explanation of oplog sizing, see *Oplog Size* (page 40). For details how oplog size affects *delayed members* and affects *replication lag*, see *Delayed Replica Set Members* (page 16).

Overview

To change the size of the oplog, you must perform maintenance on each member of the replica set in turn. The procedure requires: stopping the `mongod` instance and starting as a standalone instance, modifying the oplog size, and restarting the member.

Important: Always start rolling replica set maintenance with the secondaries, and finish with the maintenance on primary member.

Procedure

- Restart the member in standalone mode.

Tip

Always use `rs.stepDown()` to force the primary to become a secondary, before stopping the server. This facilitates a more efficient election process.

- Recreate the oplog with the new size and with an old oplog entry as a seed.
- Restart the `mongod` instance as a member of the replica set.

Restart a Secondary in Standalone Mode on a Different Port

Shut down the `mongod` instance for one of the non-primary members of your replica set. For example, to shut down, use the `db.shutdownServer()` method:

```
db.shutdownServer()
```

Restart this `mongod` as a standalone instance running on a different port and *without* the `--replSet` parameter. Use a command similar to the following:

```
mongod --port 37017 --dbpath /srv/mongoddb
```

Create a Backup of the Oplog (Optional)

Optionally, backup the existing oplog on the standalone instance, as in the following example:

```
mongodump --db local --collection 'oplog.rs' --port 37017
```

Recreate the Oplog with a New Size and a Seed Entry

Save the last entry from the oplog. For example, connect to the instance using the mongo shell, and enter the following command to switch to the `local` database:

```
use local
```

In mongo shell scripts you can use the following operation to set the `db` object:

```
db = db.getSiblingDB('local')
```

Ensure that the `temp` temporary collection is empty by dropping the collection:

```
db.temp.drop()
```

Use the `db.collection.save()` method and a sort on reverse *natural order* to find the last entry and save it to a temporary collection:

```
db.temp.save( db.oplog.rs.find( { }, { ts: 1, h: 1 } ).sort( {$natural : -1} ).limit(1).next() )
```

To see this oplog entry, use the following operation:

```
db.temp.find()
```

Remove the Existing Oplog Collection

Drop the old `oplog.rs` collection in the `local` database. Use the following command:

```
db = db.getSiblingDB('local')
db.oplog.rs.drop()
```

This returns `true` in the shell.

Create a New Oplog

Use the `create` command to create a new oplog of a different size. Specify the `size` argument in bytes. A value of `2 * 1024 * 1024 * 1024` will create a new oplog that's 2 gigabytes:

```
db.runCommand( { create: "oplog.rs", capped: true, size: (2 * 1024 * 1024 * 1024) } )
```

Upon success, this command returns the following status:

```
{ "ok" : 1 }
```

Insert the Last Entry of the Old Oplog into the New Oplog

Insert the previously saved last entry from the old oplog into the new oplog. For example:

```
db.oplog.rs.save( db.temp.findOne() )
```

To confirm the entry is in the new oplog, use the following operation:

```
db.oplog.rs.find()
```

Restart the Member

Restart the `mongod` as a member of the replica set on its usual port. For example:

```
db.shutdownServer()  
mongod --replSet rs0 --dbpath /srv/mongod
```

The replica set member will recover and “catch up” before it is eligible for election to primary.

Repeat Process for all Members that may become Primary

Repeat this procedure for all members you want to change the size of the oplog. Repeat the procedure for the primary as part of the following step.

Change the Size of the Oplog on the Primary

To finish the rolling maintenance operation, step down the primary with the `rs.stepDown()` method and repeat the oplog resizing procedure above.

3.3.2 Perform Maintenance on Replica Set Members

On this page

- [Overview](#) (page 83)
- [Procedure](#) (page 83)

Overview

Replica sets allow a MongoDB deployment to remain available during the majority of a maintenance window.

This document outlines the basic procedure for performing maintenance on each of the members of a replica set. Furthermore, this particular sequence strives to minimize the amount of time that the *primary* is unavailable and controlling the impact on the entire deployment.

Use these steps as the basis for common replica set operations, particularly for procedures such as upgrading to the latest version of MongoDB and *changing the size of the oplog* (page 80).

Procedure

For each member of a replica set, starting with a secondary member, perform the following sequence of events, ending with the primary:

- Restart the `mongod` instance as a standalone.
- Perform the task on the standalone instance.
- Restart the `mongod` instance as a member of the replica set.

Step 1: Stop a secondary.

In the `mongo` shell, shut down the `mongod` instance:

```
db.shutdownServer()
```

Step 2: Restart the secondary as a standalone on a different port.

At the operating system shell prompt, restart `mongod` as a standalone instance running on a different port and *without* the `--replSet` parameter:

```
mongod --port 37017 --dbpath /srv/mongodb
```

Always start `mongod` with the same user, even when restarting a replica set member as a standalone instance.

Step 3: Perform maintenance operations on the secondary.

While the member is a standalone, use the `mongo` shell to perform maintenance:

```
mongo --port 37017
```

Step 4: Restart `mongod` as a member of the replica set.

After performing all maintenance tasks, use the following procedure to restart the `mongod` as a member of the replica set on its usual port.

From the `mongo` shell, shut down the standalone server after completing the maintenance:

```
db.shutdownServer()
```

Restart the `mongod` instance as a member of the replica set using its normal command-line arguments or configuration file.

The secondary takes time to *catch up to the primary* (page 41). From the `mongo` shell, use the following command to verify that the member has caught up from the `RECOVERING` (page 118) state to the `SECONDARY` (page 117) state.

```
rs.status()
```

Step 5: Perform maintenance on the primary last.

To perform maintenance on the primary after completing maintenance tasks on all secondaries, use `rs.stepDown()` in the `mongo` shell to step down the primary and allow one of the secondaries to be elected the new primary. Specify a 300 second waiting period to prevent the member from being elected primary again for five minutes:

```
rs.stepDown(300)
```

After the primary steps down, the replica set will elect a new primary. See *Replica Set Elections* (page 28) for more information about replica set elections.

3.3.3 Force a Member to Become Primary

On this page

- [Overview](#) (page 85)
- [Consideration](#) (page 85)
- [Procedures](#) (page 85)

Overview

You can force a *replica set* member to become *primary* by giving it a higher `members[n].priority` (page 111) value than any other member in the set.

Optionally, you also can force a member never to become primary by setting its `members[n].priority` (page 111) value to 0, which means the member can never seek *election* (page 28) as primary. For more information, see *Priority 0 Replica Set Members* (page 14).

For more information on priorities, see `members[n].priority` (page 111).

Consideration

A majority of the configured members of a replica set *must* be available for a set to reconfigure a set or elect a primary. See *Replica Set Elections* (page 28) for more information.

Procedures

Force a Member to be Primary by Setting its Priority High

This procedure assumes your current *primary* is `m1.example.net` and that you'd like to instead make `m3.example.net` primary. The procedure also assumes you have a three-member *replica set* with the configuration below. For more information on configurations, see *Replica Set Configuration Use*.

This procedure assumes this configuration:

```
{
  "_id" : "rs",
  "version" : 7,
  "members" : [
    {
      "_id" : 0,
      "host" : "m1.example.net:27017"
    },
    {
      "_id" : 1,
      "host" : "m2.example.net:27017"
    },
    {
      "_id" : 2,
      "host" : "m3.example.net:27017"
    }
  ]
}
```

1. In a mongo shell connected to the primary, use the following sequence of operations to make `m3.example.net` the primary:

```
cfg = rs.conf()
cfg.members[0].priority = 0.5
cfg.members[1].priority = 0.5
cfg.members[2].priority = 1
rs.reconfig(cfg)
```

The last statement calls `rs.reconfig()` with the modified configuration document to configure `m3.example.net` to have a higher `members[n].priority` (page 111) value than the other mongod instances.

The following sequence of events occur:

- `m3.example.net` and `m2.example.net` sync with `m1.example.net` (typically within 10 seconds).
 - `m1.example.net` sees that it no longer has highest priority and, in most cases, steps down. `m1.example.net` *does not* step down if `m3.example.net`'s sync is far behind. In that case, `m1.example.net` waits until `m3.example.net` is within 10 seconds of its optime and then steps down. This minimizes the amount of time with no primary following failover.
 - The step down forces an election in which `m3.example.net` becomes primary based on its priority setting.
2. Optionally, if `m3.example.net` is more than 10 seconds behind `m1.example.net`'s optime, and if you don't need to have a primary designated within 10 seconds, you can force `m1.example.net` to step down by running:

```
db.adminCommand({replSetStepDown: 86400, force: 1})
```

This prevents `m1.example.net` from being primary for 86,400 seconds (24 hours), even if there is no other member that can become primary. When `m3.example.net` catches up with `m1.example.net` it will become primary.

If you later want to make `m1.example.net` primary again while it waits for `m3.example.net` to catch up, issue the following command to make `m1.example.net` seek election again:

```
rs.freeze()
```

The `rs.freeze()` provides a wrapper around the `replSetFreeze` database command.

Force a Member to be Primary Using Database Commands

Changed in version 1.8.

Consider a *replica set* with the following members:

- `mdb0.example.net` - the current *primary*.
- `mdb1.example.net` - a *secondary*.
- `mdb2.example.net` - a *secondary*.

To force a member to become primary use the following procedure:

1. In a mongo shell, run `rs.status()` to ensure your replica set is running as expected.
2. In a mongo shell connected to the mongod instance running on `mdb2.example.net`, freeze `mdb2.example.net` so that it does not attempt to become primary for 120 seconds.

```
rs.freeze(120)
```

3. In a mongo shell connected the mongod running on `mdb0.example.net`, step down this instance that the mongod is not eligible to become primary for 120 seconds:

```
rs.stepDown(120)
```

`mdb1.example.net` becomes primary.

Note: During the transition, there is a short window where the set does not have a primary.

For more information, consider the `rs.freeze()` and `rs.stepDown()` methods that wrap the `replSetFreeze` and `replSetStepDown` commands.

3.3.4 Resync a Member of a Replica Set

On this page

- [Procedures](#) (page 87)

A *replica set* member becomes “stale” when its replication process falls so far behind that the *primary* overwrites oplog entries the member has not yet replicated. The member cannot catch up and becomes “stale.” When this occurs, you must completely resynchronize the member by removing its data and performing an *initial sync* (page 42).

This tutorial addresses both resyncing a stale member and to creating a new member using seed data from another member. When syncing a member, choose a time when the system has the bandwidth to move a large amount of data. Schedule the synchronization during a time of low usage or during a maintenance window.

MongoDB provides two options for performing an initial sync:

- Restart the `mongod` with an empty data directory and let MongoDB’s normal initial syncing feature restore the data. This is the more simple option but may take longer to replace the data.

See [Procedures](#) (page 87).

- Restart the machine with a copy of a recent data directory from another member in the replica set. This procedure can replace the data more quickly but requires more manual steps.

See [Sync by Copying Data Files from Another Member](#) (page 88).

Procedures

Automatically Sync a Member

Warning: During initial sync, `mongod` will remove the content of the `dbPath`.

This procedure relies on MongoDB’s regular process for *initial sync* (page 42). This will store the current data on the member. For an overview of MongoDB initial sync process, see the [Replication Processes](#) (page 39) section.

If the instance has no data, you can simply follow the [Add Members to a Replica Set](#) (page 66) or [Replace a Replica Set Member](#) (page 70) procedure to add a new member to a replica set.

You can also force a `mongod` that is already a member of the set to to perform an initial sync by restarting the instance without the content of the `dbPath` as follows:

1. Stop the member's `mongod` instance. To ensure a clean shutdown, use the `db.shutdownServer()` method from the `mongo` shell or on Linux systems, the `mongod --shutdown` option.
2. Delete all data and sub-directories from the member's data directory. By removing the data `dbPath`, MongoDB will perform a complete resync. Consider making a backup first.

At this point, the `mongod` will perform an initial sync. The length of the initial sync process depends on the size of the database and network connection between members of the replica set.

Initial sync operations can impact the other members of the set and create additional traffic to the primary and can only occur if another member of the set is accessible and up to date.

Sync by Copying Data Files from Another Member

This approach “seeds” a new or stale member using the data files from an existing member of the replica set. The data files **must** be sufficiently recent to allow the new member to catch up with the `oplog`. Otherwise the member would need to perform an initial sync.

Copy the Data Files You can capture the data files as either a snapshot or a direct copy. However, in most cases you cannot copy data files from a running `mongod` instance to another because the data files will change during the file copy operation.

Important: If copying data files, you must copy the content of the `local` database.

You *cannot* use a `mongodump` backup for the data files, **only a snapshot backup**. For approaches to capturing a consistent snapshot of a running `mongod` instance, see the <https://docs.mongodb.org/manual/core/backups> documentation.

Sync the Member After you have copied the data files from the “seed” source, start the `mongod` instance and allow it to apply all operations from the `oplog` until it reflects the current state of the replica set.

3.3.5 Configure Replica Set Tag Sets

On this page

- [Differences Between Read Preferences and Write Concerns \(page 88\)](#)
- [Add Tag Sets to a Replica Set \(page 89\)](#)
- [Custom Multi-Datacenter Write Concerns \(page 90\)](#)
- [Configure Tag Sets for Functional Segregation of Read and Write Operations \(page 91\)](#)

Tag sets let you customize *write concern* and *read preferences* for a *replica set*. MongoDB stores tag sets in the replica set configuration object, which is the document returned by `rs.conf()`, in the `members[n].tags` (page 112) embedded document.

This section introduces the configuration of tag sets. For an overview on tag sets and their use, see `w: <tag set>` and *Tag Sets* (page 37).

Differences Between Read Preferences and Write Concerns

Custom read preferences and write concerns evaluate tag sets in different ways:

- Read preferences consider the value of a tag when selecting a member to read from.
- Write concerns do not use the value of a tag to select a member except to consider whether or not the value is unique.

For example, a tag set for a read operation may resemble the following document:

```
{ "disk": "ssd", "use": "reporting" }
```

To fulfill such a read operation, a member would need to have both of these tags. Any of the following tag sets would satisfy this requirement:

```
{ "disk": "ssd", "use": "reporting" }
{ "disk": "ssd", "use": "reporting", "rack": "a" }
{ "disk": "ssd", "use": "reporting", "rack": "d" }
{ "disk": "ssd", "use": "reporting", "mem": "r" }
```

The following tag sets would *not* be able to fulfill this query:

```
{ "disk": "ssd" }
{ "use": "reporting" }
{ "disk": "ssd", "use": "production" }
{ "disk": "ssd", "use": "production", "rack": "k" }
{ "disk": "spinning", "use": "reporting", "mem": "32" }
```

Add Tag Sets to a Replica Set

Given the following replica set configuration:

```
{
  "_id" : "rs0",
  "version" : 1,
  "members" : [
    {
      "_id" : 0,
      "host" : "mongodb0.example.net:27017"
    },
    {
      "_id" : 1,
      "host" : "mongodb1.example.net:27017"
    },
    {
      "_id" : 2,
      "host" : "mongodb2.example.net:27017"
    }
  ]
}
```

You could add tag sets to the members of this replica set with the following command sequence in the mongo shell:

```
conf = rs.conf()
conf.members[0].tags = { "dc": "east", "use": "production" }
conf.members[1].tags = { "dc": "east", "use": "reporting" }
conf.members[2].tags = { "use": "production" }
rs.reconfig(conf)
```

After this operation the output of `rs.conf()` would resemble the following:

```
{
  "_id" : "rs0",
  "version" : 2,
  "members" : [
    {
      "_id" : 0,
      "host" : "mongodb0.example.net:27017",
      "tags" : {
        "dc": "east",
        "use": "production"
      }
    },
    {
      "_id" : 1,
      "host" : "mongodb1.example.net:27017",
      "tags" : {
        "dc": "east",
        "use": "reporting"
      }
    },
    {
      "_id" : 2,
      "host" : "mongodb2.example.net:27017",
      "tags" : {
        "use": "production"
      }
    }
  ]
}
```

Important: In tag sets, all tag values must be strings.

Custom Multi-Datacenter Write Concerns

Given a five member replica set with members in two data centers:

1. a facility VA tagged `dc_va`
2. a facility GTO tagged `dc_gto`

Create a custom write concern to require confirmation from two data centers using replica set tags, using the following sequence of operations in the mongo shell:

1. Create a replica set configuration JavaScript object `conf`:

```
conf = rs.conf()
```

2. Add tags to the replica set members reflecting their locations:

```
conf.members[0].tags = { "dc_va": "rack1" }
conf.members[1].tags = { "dc_va": "rack2" }
conf.members[2].tags = { "dc_gto": "rack1" }
conf.members[3].tags = { "dc_gto": "rack2" }
conf.members[4].tags = { "dc_va": "rack1" }
rs.reconfig(conf)
```

3. Create a custom `settings.getLastErrorModes` (page 113) setting to ensure that the write operation will propagate to at least one member of each facility:

```
conf.settings = { getLastErrorModes: { MultipleDC : { "dc_va": 1, "dc_gto": 1 } } }
```

4. Reconfigure the replica set using the modified `conf` configuration object:

```
rs.reconfig(conf)
```

To ensure that a write operation propagates to at least one member of the set in both data centers, use the `MultipleDC` write concern mode as follows:

```
db.users.insert( { id: "xyz", status: "A" }, { writeConcern: { w: "MultipleDC" } } )
```

Alternatively, if you want to ensure that each write operation propagates to at least 2 racks in each facility, reconfigure the replica set as follows in the `mongo` shell:

1. Create a replica set configuration object `conf`:

```
conf = rs.conf()
```

2. Redefine the `settings.getLastErrorModes` (page 113) value to require two different values of both `dc_va` and `dc_gto`:

```
conf.settings = { getLastErrorModes: { MultipleDC : { "dc_va": 2, "dc_gto": 2} } }
```

3. Reconfigure the replica set using the modified `conf` configuration object:

```
rs.reconfig(conf)
```

Now, the following write operation will only return after the write operation propagates to at least two different racks in the each facility:

Changed in version 2.6: A new protocol for *write operations* integrates write concerns with the write operations. Previous versions used the `getLastError` command to specify the write concerns.

```
db.users.insert( { id: "xyz", status: "A" }, { writeConcern: { w: "MultipleDC" } } )
```

Configure Tag Sets for Functional Segregation of Read and Write Operations

Given a replica set with tag sets that reflect:

- data center facility,
- physical rack location of instance, and
- storage system (i.e. disk) type.

Where each member of the set has a tag set that resembles one of the following:²

```
{"dc_va": "rack1", disk:"ssd", ssd: "installed" }
{"dc_va": "rack2", disk:"raid"}
{"dc_gto": "rack1", disk:"ssd", ssd: "installed" }
{"dc_gto": "rack2", disk:"raid"}
{"dc_va": "rack1", disk:"ssd", ssd: "installed" }
```

To target a read operation to a member of the replica set with a disk type of `ssd`, you could use the following tag set:

```
{ disk: "ssd" }
```

However, to create comparable write concern modes, you would specify a different set of `settings.getLastErrorModes` (page 113) configuration. Consider the following sequence of operations in the `mongo` shell:

² Since read preferences and write concerns use the value of fields in tag sets differently, larger deployments may have some redundancy.

1. Create a replica set configuration object `conf`:

```
conf = rs.conf()
```

2. Redefine the `settings.getLastErrorModes` (page 113) value to configure two write concern modes:

```
conf.settings = {
  "getErrorModes" : {
    "ssd" : {
      "ssd" : 1
    },
    "MultipleDC" : {
      "dc_va" : 1,
      "dc_gto" : 1
    }
  }
}
```

3. Reconfigure the replica set using the modified `conf` configuration object:

```
rs.reconfig(conf)
```

Now you can specify the `MultipleDC` write concern mode, as in the following, to ensure that a write operation propagates to each data center.

Changed in version 2.6: A new protocol for *write operations* integrates write concerns with the write operations. Previous versions used the `getErrorModes` command to specify the write concerns.

```
db.users.insert( { id: "xyz", status: "A" }, { writeConcern: { w: "MultipleDC" } } )
```

Additionally, you can specify the `ssd` write concern mode to ensure that a write operation propagates to at least one instance with an SSD.

3.3.6 Reconfigure a Replica Set with Unavailable Members

On this page

- [Reconfigure by Forcing the Reconfiguration](#) (page 92)
- [Reconfigure by Replacing the Replica Set](#) (page 93)

To reconfigure a *replica set* when a **majority** of members are available, use the `rs.reconfig()` operation on the current *primary*, following the example in the *Replica Set Reconfiguration Procedure*.

This document provides the following options for re-configuring a replica set when *only* a **minority** of members are accessible:

- [Reconfigure by Forcing the Reconfiguration](#) (page 92)
- [Reconfigure by Replacing the Replica Set](#) (page 93)

You may need to use one of these procedures, for example, in a geographically distributed replica set, where *no* local group of members can reach a majority. See *Replica Set Elections* (page 28) for more information on this situation.

Reconfigure by Forcing the Reconfiguration

Changed in version 2.0.

This procedure lets you recover while a majority of *replica set* members are down or unreachable. You connect to any surviving member and use the `force` option to the `rs.reconfig()` method.

The `force` option forces a new configuration onto the member. Use this procedure only to recover from catastrophic interruptions. Do not use `force` every time you reconfigure. Also, do not use the `force` option in any automatic scripts and do not use `force` when there is still a *primary*.

To force reconfiguration:

1. Back up a surviving member.
2. Connect to a surviving member and save the current configuration. Consider the following example commands for saving the configuration:

```
cfg = rs.conf()

printjson(cfg)
```

3. On the same member, remove the down and unreachable members of the replica set from the `members` (page 110) array by setting the array equal to the surviving members alone. Consider the following example, which uses the `cfg` variable created in the previous step:

```
cfg.members = [cfg.members[0] , cfg.members[4] , cfg.members[7]]
```

4. On the same member, reconfigure the set by using the `rs.reconfig()` command with the `force` option set to `true`:

```
rs.reconfig(cfg, {force : true})
```

This operation forces the secondary to use the new configuration. The configuration is then propagated to all the surviving members listed in the `members` array. The replica set then elects a new primary.

Note: When you use `force : true`, the version number in the replica set configuration increases significantly, by tens or hundreds of thousands. This is normal and designed to prevent set version collisions if you accidentally force re-configurations on both sides of a network partition and then the network partitioning ends.

5. If the failure or partition was only temporary, shut down or decommission the removed members as soon as possible.

Reconfigure by Replacing the Replica Set

Use the following procedure **only** for versions of MongoDB prior to version 2.0. If you're running MongoDB 2.0 or later, use the above procedure, *Reconfigure by Forcing the Reconfiguration* (page 92).

These procedures are for situations where a *majority* of the *replica set* members are down or unreachable. If a majority is *running*, then skip these procedures and instead use the `rs.reconfig()` command according to the examples in *replica-set-reconfiguration-usage*.

If you run a pre-2.0 version and a majority of your replica set is down, you have the two options described here. Both involve replacing the replica set.

Reconfigure by Turning Off Replication

This option replaces the *replica set* with a *standalone* server.

1. Stop the surviving `mongod` instances. To ensure a clean shutdown, use an existing *init script* or use the `db.shutdownServer()` method.

For example, to use the `db.shutdownServer()` method, connect to the server using the `mongo` shell and issue the following sequence of commands:

```
use admin
db.shutdownServer()
```

2. Create a backup of the data directory (i.e. `dbPath`) of the surviving members of the set.

Optional

If you have a backup of the database you may instead remove this data.

3. Restart one of the `mongod` instances *without* the `--replSet` parameter.

The data is now accessible and provided by a single server that is not a replica set member. Clients can use this server for both reads and writes.

When possible, re-deploy a replica set to provide redundancy and to protect your deployment from operational interruption.

Reconfigure by “Breaking the Mirror”

This option selects a surviving *replica set* member to be the new *primary* and to “seed” a new replica set. In the following procedure, the new primary is `db0.example.net`. MongoDB copies the data from `db0.example.net` to all the other members.

1. Stop the surviving `mongod` instances. To ensure a clean shutdown, use an existing *init script* or use the `db.shutdownServer()` method.

For example, to use the `db.shutdownServer()` method, connect to the server using the `mongo` shell and issue the following sequence of commands:

```
use admin
db.shutdownServer()
```

2. Move the data directories (i.e. `dbPath`) for all the members except `db0.example.net`, so that all the members except `db0.example.net` have empty data directories. For example:

```
mv /data/db /data/db-old
```

3. Move the data files for local database (i.e. `local.*`) so that `db0.example.net` has no local database. For example

```
mkdir /data/local-old
mv /data/db/local* /data/local-old/
```

4. Start each member of the replica set normally.
5. Connect to `db0.example.net` in a `mongo` shell and run `rs.initiate()` to initiate the replica set.
6. Add the other set members using `rs.add()`. For example, to add a member running on `db1.example.net` at port 27017, issue the following command:

```
rs.add("db1.example.net:27017")
```

MongoDB performs an initial sync on the added members by copying all data from `db0.example.net` to the added members.

See also:

[Resync a Member of a Replica Set](#) (page 87)

3.3.7 Manage Chained Replication

On this page

- [Disable Chained Replication](#) (page 95)
- [Re-enable Chained Replication](#) (page 96)

Starting in version 2.0, MongoDB supports chained replication. A chained replication occurs when a *secondary* member replicates from another secondary member instead of from the *primary*. This might be the case, for example, if a secondary selects its replication target based on ping time and if the closest member is another secondary.

Chained replication can reduce load on the primary. But chained replication can also result in increased replication lag, depending on the topology of the network.

New in version 2.2.2.

You can use the `settings.chainingAllowed` (page 113) setting in *Replica Set Configuration* (page 108) to disable chained replication for situations where chained replication is causing lag.

MongoDB enables chained replication by default. This procedure describes how to disable it and how to re-enable it.

Note: If chained replication is disabled, you still can use `replSetSyncFrom` to specify that a secondary replicates from another secondary. But that configuration will last only until the secondary recalculates which member to sync from.

Disable Chained Replication

To disable chained replication, set the `settings.chainingAllowed` (page 113) field in *Replica Set Configuration* (page 108) to `false`.

You can use the following sequence of commands to set `settings.chainingAllowed` (page 113) to `false`:

1. Copy the configuration settings into the `cfg` object:

```
cfg = rs.config()
```

2. Take note of whether the current configuration settings contain the `settings` embedded document. If they do, skip this step.

Warning: To avoid data loss, skip this step if the configuration settings contain the `settings` embedded document.

If the current configuration settings **do not** contain the `settings` embedded document, create the embedded document by issuing the following command:

```
cfg.settings = { }
```

3. Issue the following sequence of commands to set `settings.chainingAllowed` (page 113) to `false`:

```
cfg.settings.chainingAllowed = false
rs.reconfig(cfg)
```

Re-enable Chained Replication

To re-enable chained replication, set `settings.chainingAllowed` (page 113) to `true`. You can use the following sequence of commands:

```
cfg = rs.config()
cfg.settings.chainingAllowed = true
rs.reconfig(cfg)
```

3.3.8 Change Hostnames in a Replica Set

On this page

- [Overview](#) (page 96)
- [Assumptions](#) (page 96)
- [Change Hostnames while Maintaining Replica Set Availability](#) (page 97)
- [Change All Hostnames at the Same Time](#) (page 98)

For most *replica sets*, the hostnames in the `members[n].host` (page 110) field never change. However, if organizational needs change, you might need to migrate some or all host names.

Note: Always use resolvable hostnames for the value of the `members[n].host` (page 110) field in the replica set configuration to avoid confusion and complexity.

Overview

This document provides two separate procedures for changing the hostnames in the `members[n].host` (page 110) field. Use either of the following approaches:

- *Change hostnames without disrupting availability* (page 97). This approach ensures your applications will always be able to read and write data to the replica set, but the approach can take a long time and may incur downtime at the application layer.

If you use the first procedure, you must configure your applications to connect to the replica set at both the old and new locations, which often requires a restart and reconfiguration at the application layer and which may affect the availability of your applications. Re-configuring applications is beyond the scope of this document.

- *Stop all members running on the old hostnames at once* (page 98). This approach has a shorter maintenance window, but the replica set will be unavailable during the operation.

See also:

Replica Set Reconfiguration Process, *Deploy a Replica Set* (page 52), and *Add Members to a Replica Set* (page 66).

Assumptions

Given a *replica set* with three members:

- `database0.example.com:27017` (the *primary*)
- `database1.example.com:27017`
- `database2.example.com:27017`

And with the following `rs.conf()` output:

```
{
  "_id" : "rs",
  "version" : 3,
  "members" : [
    {
      "_id" : 0,
      "host" : "database0.example.com:27017"
    },
    {
      "_id" : 1,
      "host" : "database1.example.com:27017"
    },
    {
      "_id" : 2,
      "host" : "database2.example.com:27017"
    }
  ]
}
```

The following procedures change the members' hostnames as follows:

- `mongodb0.example.net:27017` (the primary)
- `mongodb1.example.net:27017`
- `mongodb2.example.net:27017`

Use the most appropriate procedure for your deployment.

Change Hostnames while Maintaining Replica Set Availability

This procedure uses the above *assumptions* (page 96).

1. For each *secondary* in the replica set, perform the following sequence of operations:
 - (a) Stop the secondary.
 - (b) Restart the secondary at the new location.
 - (c) Open a `mongo` shell connected to the replica set's primary. In our example, the primary runs on port 27017 so you would issue the following command:

```
mongo --port 27017
```

- (d) Use `rs.reconfig()` to update the *replica set configuration document* (page 108) with the new hostname.

For example, the following sequence of commands updates the hostname for the secondary at the array index 1 of the `members` array (i.e. `members[1]`) in the replica set configuration document:

```
cfg = rs.conf()
cfg.members[1].host = "mongodb1.example.net:27017"
rs.reconfig(cfg)
```

For more information on updating the configuration document, see *replica-set-reconfiguration-usage*.

- (e) Make sure your client applications are able to access the set at the new location and that the secondary has a chance to catch up with the other members of the set.

Repeat the above steps for each non-primary member of the set.

2. Open a `mongo` shell connected to the primary and step down the primary using the `rs.stepDown()` method:

```
rs.stepDown()
```

The replica set elects another member to become primary.

3. When the step down succeeds, shut down the old primary.
4. Start the `mongod` instance that will become the new primary in the new location.
5. Connect to the current primary, which was just elected, and update the *replica set configuration document* (page 108) with the hostname of the node that is to become the new primary.

For example, if the old primary was at position 0 and the new primary's hostname is `mongodb0.example.net:27017`, you would run:

```
cfg = rs.conf()
cfg.members[0].host = "mongodb0.example.net:27017"
rs.reconfig(cfg)
```

6. Open a `mongo` shell connected to the new primary.
7. To confirm the new configuration, call `rs.conf()` in the `mongo` shell.

Your output should resemble:

```
{
  "_id" : "rs",
  "version" : 4,
  "members" : [
    {
      "_id" : 0,
      "host" : "mongodb0.example.net:27017"
    },
    {
      "_id" : 1,
      "host" : "mongodb1.example.net:27017"
    },
    {
      "_id" : 2,
      "host" : "mongodb2.example.net:27017"
    }
  ]
}
```

Change All Hostnames at the Same Time

This procedure uses the above *assumptions* (page 96).

1. Stop all members in the *replica set*.
2. Restart each member *on a different port* and *without* using the `--replSet` run-time option. Changing the port number during maintenance prevents clients from connecting to this host while you perform maintenance. Use the member's usual `--dbpath`, which in this example is `/data/db1`. Use a command that resembles the following:

```
mongod --dbpath /data/db1/ --port 37017
```

3. For each member of the replica set, perform the following sequence of operations:

- (a) Open a mongo shell connected to the mongod running on the new, temporary port. For example, for a member running on a temporary port of 37017, you would issue this command:

```
mongo --port 37017
```

- (b) Edit the replica set configuration manually. The replica set configuration is the only document in the `system.replset` collection in the `local` database. Edit the replica set configuration with the new hostnames and correct ports for all the members of the replica set. Consider the following sequence of commands to change the hostnames in a three-member set:

```
use local

cfg = db.system.replset.findOne( { "_id": "rs" } )

cfg.members[0].host = "mongodb0.example.net:27017"

cfg.members[1].host = "mongodb1.example.net:27017"

cfg.members[2].host = "mongodb2.example.net:27017"

db.system.replset.update( { "_id": "rs" } , cfg )
```

- (c) Stop the mongod process on the member.

4. After re-configuring all members of the set, start each mongod instance in the normal way: use the usual port number and use the `--replSet` option. For example:

```
mongod --dbpath /data/db1/ --port 27017 --replSet rs
```

5. Connect to one of the mongod instances using the mongo shell. For example:

```
mongo --port 27017
```

6. To confirm the new configuration, call `rs.conf()` in the mongo shell.

Your output should resemble:

```
{
  "_id" : "rs",
  "version" : 4,
  "members" : [
    {
      "_id" : 0,
      "host" : "mongodb0.example.net:27017"
    },
    {
      "_id" : 1,
      "host" : "mongodb1.example.net:27017"
    },
    {
      "_id" : 2,
      "host" : "mongodb2.example.net:27017"
    }
  ]
}
```

3.3.9 Configure a Secondary's Sync Target

On this page

- [Overview](#) (page 100)
- [Considerations](#) (page 100)
- [Procedure](#) (page 101)

Overview

Secondaries capture data from the primary member to maintain an up to date copy of the sets' data. However, by default secondaries may automatically change their sync targets to secondary members based on changes in the ping time between members and the state of other members' replication. See *Replica Set Data Synchronization* (page 41) and *Manage Chained Replication* (page 95) for more information.

For some deployments, implementing a custom replication sync topology may be more effective than the default sync target selection logic. MongoDB provides the ability to specify a host to use as a sync target.

To override the default sync target selection logic, you may manually configure a *secondary* member's sync target to temporarily pull *oplog* entries. The following provide access to this functionality:

- `replSetSyncFrom` command, or
- `rs.syncFrom()` helper in the mongo shell

Considerations

Sync Logic

Only modify the default sync logic as needed, and always exercise caution. `rs.syncFrom()` will not affect an in-progress initial sync operation. To affect the sync target for the initial sync, run `rs.syncFrom()` operation *before* initial sync.

If you run `rs.syncFrom()` during initial sync, MongoDB produces no error messages, but the sync target will not change until after the initial sync operation.

Persistence

`replSetSyncFrom` and `rs.syncFrom()` provide a temporary override of default behavior. `mongod` will revert to the default sync behavior in the following situations:

- The `mongod` instance restarts.
- The connection between the `mongod` and the sync target closes.

Changed in version 2.4: The sync target falls more than 30 seconds behind another member of the replica set; the `mongod` will revert to the default sync target.

Target

The member to sync from must be a valid source for data in the set. To sync from a member, the member must:

- Have data. It cannot be an arbiter, in startup or recovering mode, and must be able to answer data queries.
- Be accessible.

- Be a member of the same set in the replica set configuration.
- Build indexes with the `members [n] .buildIndexes` (page 110) setting.
- A different member of the set, to prevent syncing from itself.

If you attempt to replicate from a member that is more than 10 seconds behind the current member, `mongod` will log a warning but will still replicate from the lagging member.

If you run `replSetSyncFrom` during initial sync, MongoDB produces no error messages, but the sync target will not change until after the initial sync operation.

Procedure

To use the `replSetSyncFrom` command in the mongo shell:

```
db.adminCommand( { replSetSyncFrom: "hostname<:port>" } );
```

To use the `rs.syncFrom()` helper in the mongo shell:

```
rs.syncFrom("hostname<:port>");
```

3.4 Troubleshoot Replica Sets

On this page

- [Check Replica Set Status \(page 101\)](#)
- [Check the Replication Lag \(page 101\)](#)
- [Test Connections Between all Members \(page 103\)](#)
- [Socket Exceptions when Rebooting More than One Secondary \(page 103\)](#)
- [Check the Size of the Oplog \(page 104\)](#)
- [Oplog Entry Timestamp Error \(page 104\)](#)
- [Duplicate Key Error on local.slaves \(page 105\)](#)

This section describes common strategies for troubleshooting *replica set* deployments.

3.4.1 Check Replica Set Status

To display the current state of the replica set and current state of each member, run the `rs.status()` method in a mongo shell connected to the replica set's *primary*. For descriptions of the information displayed by `rs.status()`, see <https://docs.mongodb.org/manual/reference/command/replSetGetStatus>.

Note: The `rs.status()` method is a wrapper that runs the `replSetGetStatus` database command.

3.4.2 Check the Replication Lag

Replication lag is a delay between an operation on the *primary* and the application of that operation from the *oplog* to the *secondary*. Replication lag can be a significant issue and can seriously affect MongoDB *replica set* deployments. Excessive replication lag makes “lagged” members ineligible to quickly become primary and increases the possibility that distributed read operations will be inconsistent.

To check the current length of replication lag:

- In a mongo shell connected to the primary, call the `rs.printSlaveReplicationInfo()` method.

Returns the `syncedTo` value for each member, which shows the time when the last oplog entry was written to the secondary, as shown in the following example:

```
source: m1.example.net:27017
  syncedTo: Thu Apr 10 2014 10:27:47 GMT-0400 (EDT)
  0 secs (0 hrs) behind the primary
source: m2.example.net:27017
  syncedTo: Thu Apr 10 2014 10:27:47 GMT-0400 (EDT)
  0 secs (0 hrs) behind the primary
```

A *delayed member* (page 16) may show as 0 seconds behind the primary when the inactivity period on the primary is greater than the `members[n].slaveDelay` (page 112) value.

Note: The `rs.status()` method is a wrapper around the `replSetGetStatus` database command.

- Monitor the rate of replication by watching the oplog time in the “replica” graph in the [MongoDB Cloud Manager](#)³ and in [Ops Manager](#), an on-premise solution available in [MongoDB Enterprise Advanced](#)⁴. For more information see the [MongoDB Cloud Manager documentation](#)⁵ and [Ops Manager documentation](#)⁶.

Possible causes of replication lag include:

- **Network Latency**

Check the network routes between the members of your set to ensure that there is no packet loss or network routing issue.

Use tools including `ping` to test latency between set members and `traceroute` to expose the routing of packets network endpoints.

- **Disk Throughput**

If the file system and disk device on the secondary is unable to flush data to disk as quickly as the primary, then the secondary will have difficulty keeping state. Disk-related issues are incredibly prevalent on multi-tenant systems, including virtualized instances, and can be transient if the system accesses disk devices over an IP network (as is the case with Amazon’s EBS system.)

Use system-level tools to assess disk status, including `iostat` or `vmstat`.

- **Concurrency**

In some cases, long-running operations on the primary can block replication on secondaries. For best results, configure `write concern` to require confirmation of replication to secondaries. This prevents write operations from returning if replication cannot keep up with the write load.

Use the *database profiler* to see if there are slow queries or long-running operations that correspond to the incidences of lag.

- **Appropriate Write Concern**

If you are performing a large data ingestion or bulk load operation that requires a large number of writes to the primary, particularly with `unacknowledged write concern`, the secondaries will not be able to read the oplog fast enough to keep up with changes.

To prevent this, request `write acknowledgment write concern` after every 100, 1,000, or another interval to provide an opportunity for secondaries to catch up with the primary.

³<https://cloud.mongodb.com/?jmp=docs>

⁴<https://www.mongodb.com/products/mongodb-enterprise-advanced?jmp=docs>

⁵<https://docs.cloud.mongodb.com/>

⁶<https://docs.opsmanager.mongodb.com/current/>

For more information see:

- *Write Concern*
- *Replica Set Write Concern*
- *Olog Size* (page 40)

3.4.3 Test Connections Between all Members

All members of a *replica set* must be able to connect to every other member of the set to support replication. Always verify connections in both “directions.” Networking topologies and firewall configurations can prevent normal and required connectivity, which can block replication.

Consider the following example of a bidirectional test of networking:

Example

Given a replica set with three members running on three separate hosts:

- `m1.example.net`
- `m2.example.net`
- `m3.example.net`

1. Test the connection from `m1.example.net` to the other hosts with the following operation set from `m1.example.net`:

```
mongo --host m2.example.net --port 27017
```

```
mongo --host m3.example.net --port 27017
```

2. Test the connection from `m2.example.net` to the other two hosts with the following operation set from `m2.example.net`, as in:

```
mongo --host m1.example.net --port 27017
```

```
mongo --host m3.example.net --port 27017
```

You have now tested the connection between `m2.example.net` and `m1.example.net` in both directions.

3. Test the connection from `m3.example.net` to the other two hosts with the following operation set from the `m3.example.net` host, as in:

```
mongo --host m1.example.net --port 27017
```

```
mongo --host m2.example.net --port 27017
```

If any connection, in any direction fails, check your networking and firewall configuration and reconfigure your environment to allow these connections.

3.4.4 Socket Exceptions when Rebooting More than One Secondary

When you reboot members of a replica set, ensure that the set is able to elect a primary during the maintenance. This means ensuring that a majority of the set’s `members[n].votes` (page 112) are available.

When a set’s active members can no longer form a majority, the set’s *primary* steps down and becomes a *secondary*. The former primary closes all open connections to client applications. Clients attempting to write to the former primary receive socket exceptions and *Connection reset* errors until the set can elect a primary.

Example

Given a three-member replica set where every member has one vote, the set can elect a primary if at least two members can connect to each other. If you reboot the two secondaries at once, the primary steps down and becomes a secondary. Until at least another secondary becomes available, i.e. at least one of the rebooted secondaries also becomes available, the set has no primary and cannot elect a new primary.

For more information on votes, see *Replica Set Elections* (page 28). For related information on connection errors, see *faq-keepalive*.

3.4.5 Check the Size of the Oplog

A larger *oplog* can give a replica set a greater tolerance for lag, and make the set more resilient.

To check the size of the oplog for a given *replica set* member, connect to the member in a `mongo` shell and run the `rs.printReplicationInfo()` method.

The output displays the size of the oplog and the date ranges of the operations contained in the oplog. In the following example, the oplog is about 10 MB and is able to fit about 26 hours (94400 seconds) of operations:

```
configured oplog size: 10.10546875MB
log length start to end: 94400 (26.22hrs)
oplog first event time: Mon Mar 19 2012 13:50:38 GMT-0400 (EDT)
oplog last event time: Wed Oct 03 2012 14:59:10 GMT-0400 (EDT)
now: Wed Oct 03 2012 15:00:21 GMT-0400 (EDT)
```

The oplog should be long enough to hold all transactions for the longest downtime you expect on a secondary. At a minimum, an oplog should be able to hold minimum 24 hours of operations; however, many users prefer to have 72 hours or even a week's work of operations.

For more information on how oplog size affects operations, see:

- *Oplog Size* (page 40),
- *Delayed Replica Set Members* (page 16), and
- *Check the Replication Lag* (page 101).

Note: You normally want the oplog to be the same size on all members. If you resize the oplog, resize it on all members.

To change oplog size, see the *Change the Size of the Oplog* (page 80) tutorial.

3.4.6 Oplog Entry Timestamp Error

Consider the following error in `mongod` output and logs:

```
replSet error fatal couldn't query the local local.oplog.rs collection. Terminating mongod after 30
<timestamp> [rsStart] bad replSet oplog entry?
```

Often, an incorrectly typed value in the `ts` field in the last *oplog* entry causes this error. The correct data type is `Timestamp`.

Check the type of the `ts` value using the following two queries against the oplog collection:

```

db = db.getSiblingDB("local")
db.oplog.rs.find().sort({$natural:-1}).limit(1)
db.oplog.rs.find({ts:{$type:17}}).sort({$natural:-1}).limit(1)

```

The first query returns the last document in the oplog, while the second returns the last document in the oplog where the `ts` value is a Timestamp. The `$type` operator allows you to select *BSON type* 17, is the Timestamp data type.

If the queries don't return the same document, then the last document in the oplog has the wrong data type in the `ts` field.

Example

If the first query returns this as the last oplog entry:

```

{ "ts" : {t: 1347982456000, i: 1},
  "h" : NumberLong("8191276672478122996"),
  "op" : "n",
  "ns" : "",
  "o" : { "msg" : "Reconfig set", "version" : 4 } }

```

And the second query returns this as the last entry where `ts` has the Timestamp type:

```

{ "ts" : Timestamp(1347982454000, 1),
  "h" : NumberLong("6188469075153256465"),
  "op" : "n",
  "ns" : "",
  "o" : { "msg" : "Reconfig set", "version" : 3 } }

```

Then the value for the `ts` field in the last oplog entry is of the wrong data type.

To set the proper type for this value and resolve this issue, use an update operation that resembles the following:

```

db.oplog.rs.update( { ts: { t:1347982456000, i:1 } },
                   { $set: { ts: new Timestamp(1347982456000, 1)}})

```

Modify the timestamp values as needed based on your oplog entry. This operation may take some period to complete because the update must scan and pull the entire oplog into memory.

3.4.7 Duplicate Key Error on `local.slaves`

Changed in version 3.0.0.

MongoDB 3.0.0 removes the `local.slaves` (page 116) collection. For `local.slaves` error in earlier versions of MongoDB, refer to the appropriate version of the MongoDB Manual.

Replication Reference

On this page

- [Replication Methods in the mongo Shell \(page 107\)](#)
- [Replication Database Commands \(page 108\)](#)
- [Replica Set Reference Documentation \(page 108\)](#)

4.1 Replication Methods in the mongo Shell

Name	Description
<code>rs.add()</code>	Adds a member to a replica set.
<code>rs.addArb()</code>	Adds an <i>arbiter</i> to a replica set.
<code>rs.conf()</code>	Returns the replica set configuration document.
<code>rs.freeze()</code>	Prevents the current member from seeking election as primary for a period of time.
<code>rs.help()</code>	Returns basic help text for <i>replica set</i> functions.
<code>rs.initiate()</code>	Initializes a new replica set.
<code>rs.printReplicationInfo()</code>	Prints a report of the status of the replica set from the perspective of the primary.
<code>rs.printSlaveReplicaInfo()</code>	Prints a report of the status of the replica set from the perspective of the secondaries.
<code>rs.reconfig()</code>	Re-configures a replica set by applying a new replica set configuration object.
<code>rs.remove()</code>	Remove a member from a replica set.
<code>rs.slaveOk()</code>	Sets the <code>slaveOk</code> property for the current connection. Deprecated. Use <code>readPref()</code> and <code>Mongo.setReadPref()</code> to set <i>read preference</i> .
<code>rs.status()</code>	Returns a document with information about the state of the replica set.
<code>rs.stepDown()</code>	Causes the current <i>primary</i> to become a secondary which forces an <i>election</i> .
<code>rs.syncFrom()</code>	Sets the member that this replica set member will sync from, overriding the default sync target selection logic.

4.2 Replication Database Commands

Name	Description
<code>replSetFreeze</code>	Prevents the current member from seeking election as <i>primary</i> for a period of time.
<code>replSetGetStatus</code>	Returns a document that reports on the status of the replica set.
<code>replSetInitiate</code>	Initializes a new replica set.
<code>replSetMaintenance</code>	Enables or disables a maintenance mode, which puts a <i>secondary</i> node in a RECOVERING state.
<code>replSetReconfig</code>	Applies a new configuration to an existing replica set.
<code>replSetStepDown</code>	Forces the current <i>primary</i> to <i>step down</i> and become a <i>secondary</i> , forcing an election.
<code>replSetSyncFrom</code>	Explicitly override the default logic for selecting a member to replicate from.
<code>resync</code>	Forces a <code>mongod</code> to re-synchronize from the <i>master</i> . For master-slave replication only.
<code>applyOps</code>	Internal command that applies <i>oplog</i> entries to the current data set.
<code>isMaster</code>	Displays information about this member's role in the replica set, including whether it is the master.
<code>replSetGetConfig</code>	Returns the replica set's configuration object.

4.3 Replica Set Reference Documentation

Replica Set Configuration (page 108) Complete documentation of the *replica set* configuration object returned by `rs.conf()`.

The local Database (page 114) Complete documentation of the content of the `local` database that `mongod` instances use to support replication.

Replica Set Member States (page 116) Reference for the replica set member states.

Read Preference Reference (page 118) Complete documentation of the five read preference modes that the MongoDB drivers support.

4.3.1 Replica Set Configuration

On this page

- [Example Output \(page 108\)](#)
- [Replica Set Configuration Fields \(page 109\)](#)

You can access the configuration of a *replica set* using the `rs.conf()` method or the `replSetGetConfig` command.

To modify the configuration for a replica set, use the `rs.reconfig()` method, passing a configuration document to the method. See `rs.reconfig()` for more information.

Example Output

The following document provides a representation of a replica set configuration document. The configuration of your replica set may include only a subset of these settings:

```
{
  _id: <string>,
  version: <int>,
```

```

protocolVersion: <number>,
members: [
  {
    _id: <int>,
    host: <string>,
    arbiterOnly: <boolean>,
    buildIndexes: <boolean>,
    hidden: <boolean>,
    priority: <number>,
    tags: <document>,
    slaveDelay: <int>,
    votes: <number>
  },
  ...
],
settings: {
  chainingAllowed : <boolean>,
  heartbeatIntervalMillis : <int>,
  heartbeatTimeoutSecs: <int>,
  electionTimeoutMillis : <int>,
  getLastErrorModes : <document>,
  getLastErrorDefaults : <document>
}
}

```

Replica Set Configuration Fields

`_id`

Type: string

The name of the replica set. Once set, you cannot change the name of a replica set.

`_id` (page 109) *must* be identical to the `replication.replSetName` or the value of `-replSet` specified to `mongod` on the command line.

See

`replSetName` or `--replSet` for information on setting the replica set name.

version

Type: int

An incrementing number used to distinguish revisions of the replica set configuration object from previous iterations of the configuration.

configsvr

New in version 3.2.

Type: boolean

Default: false

Indicates whether the replica set is used for a sharded cluster's config servers. Set to `true` if the replica set is for a sharded cluster's config servers.

See also:

3.2-rel-notes-sharded-cluster

`protocolVersion`

New in version 3.2.

Type: number

Default: 1 for new replica sets

Version of the *replica set election protocol*.

Set to 1 to enable the *replication election enhancements* introduced in MongoDB 3.2.

By default, new replica sets in MongoDB 3.2 use `protocolVersion: 1`. Previous versions of MongoDB use version 0 of the protocol and cannot run as members of a replica set configuration that specifies `protocolVersion 1`.

`members`

`members`

Type: array

An array of member configuration documents, one for each member of the replica set. The `members` (page 110) array is a zero-indexed array.

Each member-specific configuration document can contain the following fields:

`members[n]._id`

Type: integer

An integer identifier of every member in the replica set. Values must be between 0 and 255 inclusive. Each replica set member must have a unique `_id`<`members[n]._id`>. Once set, you cannot change the `_id` (page 110) of a member.

Note: When updating the replica configuration object, access the replica set members in the `members` (page 110) array with the **array index**. The array index begins with 0. Do **not** confuse this index value with the value of the `members[n]._id` (page 110) field in each document in the `members` (page 110) array.

`members[n].host`

Type: string

The hostname and, if specified, the port number, of the set member.

The hostname name must be resolvable for every host in the replica set.

Warning: `members[n].host` (page 110) cannot hold a value that resolves to `localhost` or the local interface unless *all* members of the set are on hosts that resolve to `localhost`.

`members[n].arbiterOnly`

Optional.

Type: boolean

Default: false

A boolean that identifies an arbiter. A value of `true` indicates that the member is an arbiter.

When using the `rs.addArb()` method to add an arbiter, the method automatically sets `members[n].arbiterOnly` (page 110) to `true` for the added member.

`members[n].buildIndexes`*Optional.**Type:* boolean*Default:* true

A boolean that indicates whether the mongod builds *indexes* on this member. You can only set this value when adding a member to a replica set. You cannot change `members[n].buildIndexes` (page 110) field after the member has been added to the set. To add a member, see `rs.add()` and `rs.reconfig()`.

Do not set to `false` for mongod instances that receive queries from clients.

Setting `buildIndexes` to `false` may be useful if **all** the following conditions are true:

- you are only using this instance to perform backups using `mongodump`, *and*
- this member will receive no queries, *and*
- index creation and maintenance overburdens the host system.

Even if set to `false`, secondaries *will* build indexes on the `_id` field in order to facilitate operations required for replication.

Warning: If you set `members[n].buildIndexes` (page 110) to `false`, you must also set `members[n].priority` (page 111) to 0. If `members[n].priority` (page 111) is not 0, MongoDB will return an error when attempting to add a member with `members[n].buildIndexes` (page 110) equal to `false`.

To ensure the member receives no queries, you should make all instances that do not build indexes hidden.

Other secondaries cannot replicate from a member where `members[n].buildIndexes` (page 110) is `false`.

`members[n].hidden`*Optional.**Type:* boolean*Default:* false

When this value is `true`, the replica set hides this instance and does not include the member in the output of `db.isMaster()` or `isMaster`. This prevents read operations (i.e. queries) from ever reaching this host by way of secondary *read preference*.

See also:

Hidden Replica Set Members (page 15)

`members[n].priority`*Optional.**Type:* Number, between 0 and 1000.*Default:* 1.0

A number that indicates the relative eligibility of a member to become a *primary*.

Specify higher values to make a member *more* eligible to become *primary*, and lower values to make the member *less* eligible. A member with a `members[n].priority` (page 111) of 0 is ineligible to become primary.

Changing the balance of priority in a replica set will trigger one or more elections. If a lower priority secondary is elected over a higher priority secondary, replica set members will continue to call elections until the highest priority available member becomes primary.

See also:

[Replica Set Elections](#) (page 28).

`members[n].tags`

Optional.

Type: document

Default: none

A *tag set* document containing mappings of arbitrary keys and values. These documents describe replica set members in order to customize `write concern` and *read preference* (page 118) and thereby allow configurable data center awareness.

This field is only present if there are tags assigned to the member. See [Configure Replica Set Tag Sets](#) (page 88) for more information.

Use `replicaset.members[n].tags` to configure write concerns in conjunction with `settings.getLastErrorModes` (page 113) and `settings.getLastErrorDefaults` (page 113).

Important: In tag sets, all tag values must be strings.

For more information on configuring tag sets for read preference and write concern, see [Configure Replica Set Tag Sets](#) (page 88).

`members[n].slaveDelay`

Optional.

Type: integer

Default: 0

The number of seconds “behind” the primary that this replica set member should “lag”.

Use this option to create *delayed members* (page 16). Delayed members maintain a copy of the data that reflects the state of the data at some time in the past.

See also:

[Delayed Replica Set Members](#) (page 16)

`members[n].votes`

Optional.

Type: integer

Default: 1

The number of votes a server will cast in a *replica set election* (page 28). The number of votes each member has is either 1 or 0, and *arbiters* (page ??) always have exactly 1 vote.

A replica set can have up to 50 members but only 7 voting members. If you need more than 7 members in one replica set, set `members[n].votes` (page 112) to 0 for the additional non-voting members.

Changed in version 3.0.0: Members cannot have `members[n].votes` (page 112) greater than 1. For details, see [3.0-compatibility-repl-set-config](#).

settings**settings***Optional.**Type:* document

A document that contains configuration options that apply to the whole replica set.

The `settings` (page 113) document contain the following fields:

settings.chainingAllowed

New in version 2.2.4.

*Optional.**Type:* boolean*Default:* true

When `settings.chainingAllowed` (page 113) is `true`, the replica set allows *secondary* members to replicate from other secondary members. When `settings.chainingAllowed` (page 113) is `false`, secondaries can replicate only from the *primary*.

See also:

Manage Chained Replication (page 95)

settings.getLastErrorDefaults*Optional.**Type:* document

A document that specifies the *write concern* (page 32) for the replica set. The replica set will use this write concern only when *write operations* or `getLastError` specify no other write concern.

If `settings.getLastErrorDefaults` (page 113) is not set, the default write concern for the replica set only requires confirmation from the primary.

settings.getErrorModes*Optional.**Type:* document

A document used to define an extended *write concern* through the use of `members[n].tags` (page 112). The extended *write concern* can provide *data-center awareness*.

For example, the following document defines an extended write concern named `eastCoast` and associates with a write to a member that has the `east` tag.

```
{ getLastErrorModes: { eastCoast: { "east": 1 } } }
```

Write operations to the replica set can use the extended write concern, e.g. { w: "eastCoast" }.

See *Configure Replica Set Tag Sets* (page 88) for more information and example.

settings.heartbeatTimeoutSecs*Optional.**Type:* int*Default:* 10

Number of seconds that the replica set members wait for a successful heartbeat from each other. If a member does not respond in time, other members mark the delinquent member as inaccessible.

`settings.electionTimeoutMillis`

New in version 3.2.

Optional.

Type: int

Default: 10000 (10 seconds)

The time limit in milliseconds for detecting when a replica set's primary is unreachable:

- Higher values result in slower failovers but decreased sensitivity to primary node or network slowness or spottiness.
- Lower values result in faster failover, but increased sensitivity to primary node or network slowness or spottiness.

The setting only applies when using `protocolVersion: 1`.

`settings.heartbeatIntervalMillis`

New in version 3.2.

Internal use only.

The frequency in milliseconds of the heartbeats.

4.3.2 The `local` Database

On this page

- [Overview](#) (page 114)
- [Collection on all `mongod` Instances](#) (page 114)
- [Collections on Replica Set Members](#) (page 115)
- [Collections used in Master/Slave Replication](#) (page 116)

Overview

Every `mongod` instance has its own `local` database, which stores data used in the replication process, and other instance-specific data. The `local` database is invisible to replication: collections in the `local` database are not replicated.

In replication, the `local` database store stores internal replication data for each member of a *replica set*. The `local` stores the following collections:

Changed in version 2.4: When running with authentication (i.e. `authorization`), authenticating to the `local` database is **not** equivalent to authenticating to the `admin` database. In previous versions, authenticating to the `local` database provided access to all databases.

Collection on all `mongod` Instances

`local.startup_log`

On startup, each `mongod` instance inserts a document into `startup_log` (page 114) with diagnostic information about the `mongod` instance itself and host information. `startup_log` (page 114) is a capped collection. This information is primarily useful for diagnostic purposes.

Example

Consider the following prototype of a document from the `startup_log` (page 114) collection:

```
{
  "_id" : "<string>",
  "hostname" : "<string>",
  "startTime" : ISODate("<date>"),
  "startTimeLocal" : "<string>",
  "cmdLine" : {
    "dbpath" : "<path>",
    "<option>" : <value>
  },
  "pid" : <number>,
  "buildinfo" : {
    "version" : "<string>",
    "gitVersion" : "<string>",
    "sysInfo" : "<string>",
    "loaderFlags" : "<string>",
    "compilerFlags" : "<string>",
    "allocator" : "<string>",
    "versionArray" : [ <num>, <num>, <...> ],
    "javascriptEngine" : "<string>",
    "bits" : <number>,
    "debug" : <boolean>,
    "maxBsonObjectSize" : <number>
  }
}
```

Documents in the `startup_log` (page 114) collection contain the following fields:

`local.startup_log._id`

Includes the system hostname and a millisecond epoch value.

`local.startup_log.hostname`

The system's hostname.

`local.startup_log.startTime`

A UTC *ISODate* value that reflects when the server started.

`local.startup_log.startTimeLocal`

A string that reports the `startTime` (page 115) in the system's local time zone.

`local.startup_log.cmdLine`

An embedded document that reports the mongod runtime options and their values.

`local.startup_log.pid`

The process identifier for this process.

`local.startup_log.buildinfo`

An embedded document that reports information about the build environment and settings used to compile this mongod. This is the same output as `buildInfo`. See `buildInfo`.

Collections on Replica Set Members

`local.system.replset`

`local.system.replset` (page 115) holds the replica set's configuration object as its single document. To view the object's configuration information, issue `rs.conf()` from the `mongo` shell. You can also query this collection directly.

`local.oplog.rs`

`local.oplog.rs` (page 115) is the capped collection that holds the *oplog*. You set its size at creation using the `oplogSizeMB` setting. To resize the oplog after replica set initiation, use the *Change the Size of the Oplog* (page 80) procedure. For additional information, see the *Oplog Size* (page 40) section.

`local.replset.minvalid`

This contains an object used internally by replica sets to track replication status.

`local.slaves`

Removed in version 3.0: Replica set members no longer mirror replication status of the set to the `local.slaves` (page 116) collection. Use `rs.status()` instead.

Collections used in Master/Slave Replication

In *master/slave* replication, the `local` database contains the following collections:

- On the master:

`local.oplog.$main`

This is the oplog for the master-slave configuration.

`local.slaves`

Removed in version 3.0: MongoDB no longer stores information about each slave in the `local.slaves` (page 116) collection. Use `db.serverStatus({ repl: 1 })` instead.

- On each slave:

`local.sources`

This contains information about the slave's master server.

4.3.3 Replica Set Member States

On this page

- [States](#) (page 117)

Each member of a replica set has a state that reflects its disposition within the set.

Number	Name	State Description
0	STARTUP (page 117)	Not yet an active member of any set. All members start up in this state. The <code>mongod</code> parses the <i>replica set configuration document</i> (page 71) while in <code>STARTUP</code> (page 117).
1	PRIMARY (page 117)	The member in state <i>primary</i> (page 10) is the only member that can accept write operations.
2	SECONDARY (page 117)	A member in state <i>secondary</i> (page 10) is replicating the data store. Data is available for reads, although they may be stale.
3	RECOVERING (page 118)	Can vote. Members either perform startup self-checks, or transition from completing a <i>rollback</i> (page 31) or <i>resync</i> (page 87).
5	STARTUP2 (page 117)	The member has joined the set and is running an initial sync.
6	UNKNOWN (page 118)	The member's state, as seen from another member of the set, is not yet known.
7	ARBITER (page 117)	<i>Arbiters</i> (page ??) do not replicate data and exist solely to participate in elections.
8	DOWN (page 118)	The member, as seen from another member of the set, is unreachable.
9	ROLLBACK (page 118)	This member is actively performing a <i>rollback</i> (page 31). Data is not available for reads.
10	REMOVED (page 118)	This member was once in a replica set but was subsequently removed.

States

Core States

PRIMARY

Members in `PRIMARY` (page 117) state accept write operations. A replica set has at most one primary at a time. A `SECONDARY` (page 117) member becomes primary after an *election* (page 28). Members in the `PRIMARY` (page 117) state are eligible to vote.

SECONDARY

Members in `SECONDARY` (page 117) state replicate the primary's data set and can be configured to accept read operations. Secondaries are eligible to vote in elections, and may be elected to the `PRIMARY` (page 117) state if the primary becomes unavailable.

ARBITER

Members in `ARBITER` (page 117) state do not replicate data or accept write operations. They are eligible to vote, and exist solely to break a tie during elections. Replica sets should only have a member in the `ARBITER` (page 117) state if the set would otherwise have an even number of members, and could suffer from tied elections. There should only be at most one arbiter configured in any replica set.

See *Replica Set Members* (page 9) for more information on core states.

Other States

STARTUP

Each member of a replica set starts up in `STARTUP` (page 117) state. `mongod` then loads that member's replica set configuration, and transitions the member's state to `STARTUP2` (page 117). Members in `STARTUP` (page 117) are not eligible to vote, as they are not yet a recognized member of any replica set.

STARTUP2

Each member of a replica set enters the `STARTUP2` (page 117) state as soon as `mongod` finishes loading

that member's configuration, at which time it becomes an active member of the replica set. The member then decides whether or not to undertake an initial sync. If a member begins an initial sync, the member remains in [STARTUP2](#) (page 117) until all data is copied and all indexes are built. Afterwards, the member transitions to [RECOVERING](#) (page 118).

RECOVERING

A member of a replica set enters [RECOVERING](#) (page 118) state when it is not ready to accept reads. The [RECOVERING](#) (page 118) state can occur during normal operation, and doesn't necessarily reflect an error condition. Members in the [RECOVERING](#) (page 118) state are eligible to vote in elections, but are not eligible to enter the [PRIMARY](#) (page 117) state.

A member transitions from [RECOVERING](#) (page 118) to [SECONDARY](#) (page 117) after replicating enough data to guarantee a consistent view of the data for client reads. The only difference between [RECOVERING](#) (page 118) and [SECONDARY](#) (page 117) states is that [RECOVERING](#) (page 118) prohibits client reads and [SECONDARY](#) (page 117) permits them. [SECONDARY](#) (page 117) state does not guarantee anything about the staleness of the data with respect to the primary.

Due to overload, a *secondary* may fall far enough behind the other members of the replica set such that it may need to *resync* (page 87) with the rest of the set. When this happens, the member enters the [RECOVERING](#) (page 118) state and requires manual intervention.

Error States

Members in any error state can't vote.

UNKNOWN

Members that have never communicated status information to the replica set are in the [UNKNOWN](#) (page 118) state.

DOWN

Members that lose their connection to the replica set are seen as [DOWN](#) (page 118) by the remaining members of the set.

REMOVED

Members that are removed from the replica set enter the [REMOVED](#) (page 118) state. When members enter the [REMOVED](#) (page 118) state, the logs will mark this event with a `replSet REMOVED` message entry.

ROLLBACK

Whenever the replica set replaces a *primary* in an election, the old primary may contain documents that did not replicate to the *secondary* members. In this case, the old primary member reverts those writes. During *rollback* (page 31), the member will have [ROLLBACK](#) (page 118) state.

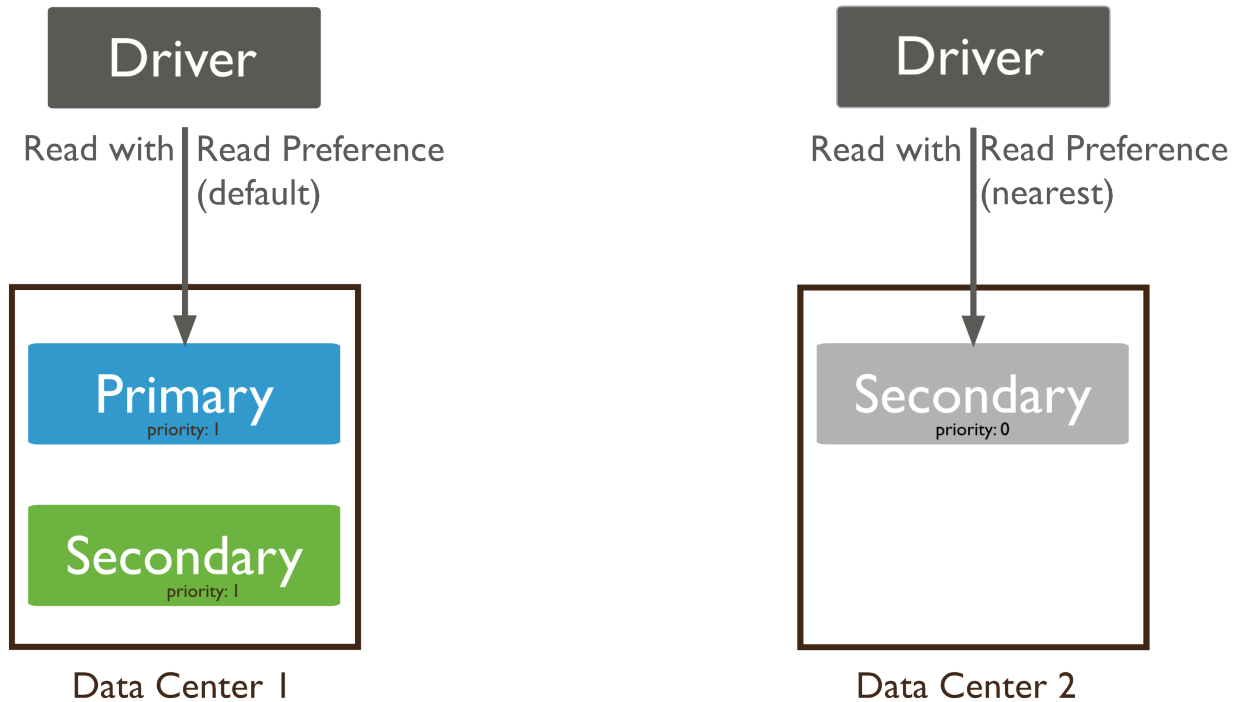
FATAL

A member in [FATAL](#) (page 118) encountered an unrecoverable error. The member must be shut down and restarted; a resync may be required as well.

4.3.4 Read Preference Reference

On this page

- [Read Preference Modes](#) (page 120)
- [Use Cases](#) (page 121)
- [Read Preferences for Database Commands](#) (page 122)



Read preference describes how MongoDB clients route read operations to the members of a *replica set*.

By default, an application directs its read operations to the *primary* member in a *replica set*.

In MongoDB, in a replica set with one primary member ¹,

- With "local" `readConcern`, reads from the primary reflect the latest writes in absence of a failover;
- With "majority" `readConcern`, read operations from the primary or the secondaries have *eventual consistency*.

Read Preference Mode	Description
<code>primary</code> (page 120)	Default mode. All operations read from the current replica set <i>primary</i> .
<code>primaryPreferred</code> (page 120)	In most situations, operations read from the <i>primary</i> but if it is unavailable, operations read from <i>secondary</i> members.
<code>secondary</code> (page 120)	All operations read from the <i>secondary</i> members of the replica set.
<code>secondaryPreferred</code> (page 120)	In most situations, operations read from <i>secondary</i> members but if no <i>secondary</i> members are available, operations read from the <i>primary</i> .
<code>nearest</code> (page 120)	Operations read from member of the <i>replica set</i> with the least network latency, irrespective of the member's type.

Note: The read preference does not affect the visibility of data; i.e. clients can see the results of writes before they are made *durable*:

- Regardless of `write concern`, other clients using "local" (i.e. the default) `readConcern` can see the result of a write operation before the write operation is acknowledged to the issuing client.

¹ In *some circumstances* (page 121), two nodes in a replica set may *transiently* believe that they are the primary, but at most, one of them will be able to complete writes with { `w: "majority"` } write concern. The node that can complete { `w: "majority"` } writes is the current primary, and the other node is a former primary that has not yet recognized its demotion, typically due to a *network partition*. When this occurs, clients that connect to the former primary may observe stale data despite having requested read preference `primary` (page 120), and new writes to the former primary will eventually roll back.

- Clients using "local" (i.e. the default) readConcern can read data which may be subsequently *rolled back* (page 31).

For more information on read isolation level in MongoDB, see <https://docs.mongodb.org/manual/core/read-isolation>

Read Preference Modes

primary

All read operations use only the current replica set *primary*.⁵ This is the default read mode. If the primary is unavailable, read operations produce an error or throw an exception.

The `primary` (page 120) read preference mode is not compatible with read preference modes that use *tag sets* (page 37). If you specify a tag set with `primary` (page 120), the driver will produce an error.

primaryPreferred

In most situations, operations read from the *primary* member of the set. However, if the primary is unavailable, as is the case during *failover* situations, operations read from secondary members.

When the read preference includes a *tag set* (page 37), the client reads first from the primary, if available, and then from *secondaries* that match the specified tags. If no secondaries have matching tags, the read operation produces an error.

Since the application may receive data from a secondary, read operations using the `primaryPreferred` (page 120) mode may return stale data in some situations.

Warning: Changed in version 2.2: mongos added full support for read preferences. When connecting to a mongos instance older than 2.2, using a client that supports read preference modes, `primaryPreferred` (page 120) will send queries to secondaries.

secondary

Operations read *only* from the *secondary* members of the set. If no secondaries are available, then this read operation produces an error or exception.

Most sets have at least one secondary, but there are situations where there may be no available secondary. For example, a set with a primary, a secondary, and an *arbiter* may not have any secondaries if a member is in recovering state or unavailable.

When the read preference includes a *tag set* (page 37), the client attempts to find secondary members that match the specified tag set and directs reads to a random secondary from among the *nearest group* (page 38). If no secondaries have matching tags, the read operation produces an error.²

Read operations using the `secondary` (page 120) mode may return stale data.

secondaryPreferred

In most situations, operations read from *secondary* members, but in situations where the set consists of a single *primary* (and no other members), the read operation will use the set's primary.

When the read preference includes a *tag set* (page 37), the client attempts to find a secondary member that matches the specified tag set and directs reads to a random secondary from among the *nearest group* (page 38). If no secondaries have matching tags, the client ignores tags and reads from the primary.

Read operations using the `secondaryPreferred` (page 120) mode may return stale data.

nearest

The driver reads from the *nearest* member of the *set* according to the *member selection* (page 38) process. Reads

² If your set has more than one secondary, and you use the `secondary` (page 120) read preference mode, consider the following effect. If you have a *three member replica set* (page 21) with a primary and two secondaries, and one secondary becomes unavailable, all `secondary` (page 120) queries must target the remaining secondary. This will double the load on this secondary. Plan and provide capacity to support this as needed.

in the `nearest` (page 120) mode do not consider the member's *type*. Reads in `nearest` (page 120) mode may read from both primaries and secondaries.

Set this mode to minimize the effect of network latency on read operations without preference for current or stale data.

If you specify a *tag set* (page 37), the client attempts to find a replica set member that matches the specified tag set and directs reads to an arbitrary member from among the *nearest group* (page 38).

Read operations using the `nearest` (page 120) mode may return stale data.

Note: All operations read from a member of the nearest group of the replica set that matches the specified read preference mode. The `nearest` (page 120) mode prefers low latency reads over a member's *primary* or *secondary* status.

For `nearest` (page 120), the client assembles a list of acceptable hosts based on tag set and then narrows that list to the host with the shortest ping time and all other members of the set that are within the "local threshold," or acceptable latency. See *Member Selection* (page 38) for more information.

Use Cases

Depending on the requirements of an application, you can configure different applications to use different read preferences, or use different read preferences for different queries in the same application. Consider the following applications for different read preference strategies.

Maximize Consistency

To avoid *stale* reads, use `primary` (page 120) read preference and "majority" `readConcern`. If the primary is unavailable, e.g. during elections or when a majority of the replica set is not accessible, read operations using `primary` (page 120) read preference produce an error or throw an exception. In some circumstances, it may be possible for a replica set to temporarily have two primaries; however, only one primary will be capable of confirming writes with the "majority" write concern.

- A partial *network partition* may segregate a primary (P_{old}) into a partition with a minority of the nodes, while the other side of the partition contains a majority of nodes. The partition with the majority will elect a new primary (P_{new}), but for a brief period, the old primary (P_{old}) may still continue to serve reads and writes, as it has not yet detected that it can only see a minority of nodes in the replica set. During this period, if the old primary (P_{old}) is still visible to clients as a primary, reads from this primary may reflect stale data.
- A primary (P_{old}) may become unresponsive, which will trigger an election and a new primary (P_{new}) can be elected, serving reads and writes. If the unresponsive primary (P_{old}) starts responding again, two primaries will be visible for a brief period. The brief period will end when P_{old} steps down. However, during the brief period, clients might read from the old primary P_{old} , which can provide stale data.

To increase consistency, you can disable automatic *failover*; however, disabling automatic failover sacrifices availability.

Maximize Availability

To permit read operations when possible, use `primaryPreferred` (page 120). When there's a primary you will get consistent reads⁵, but if there is no primary you can still query *secondaries*. However, when using this read mode, consider the situation described in *Reduce load on the primary* (page 122).

Minimize Latency

To always read from a low-latency node, use `nearest` (page 120). The driver or `mongos` will read from the nearest member and those no more than 15 milliseconds³ further away than the nearest member.

`nearest` (page 120) does *not* guarantee consistency. If the nearest member to your application server is a secondary with some replication lag, queries could return stale data. `nearest` (page 120) only reflects network distance and does not reflect I/O or CPU load.

Query From Geographically Distributed Members

If the members of a replica set are geographically distributed, you can create replica tags based that reflect the location of the instance and then configure your application to query the members nearby.

For example, if members in “east” and “west” data centers are *tagged* (page 88) `{ 'dc' : 'east' }` and `{ 'dc' : 'west' }`, your application servers in the east data center can read from nearby members with the following read preference:

```
db.collection.find().readPref( { mode: 'nearest',
                                tags: [ { 'dc': 'east' } ] } )
```

Although `nearest` (page 120) already favors members with low network latency, including the tag makes the choice more predictable.

Reduce load on the primary

To shift read load from the primary, use mode `secondary` (page 120). Although `secondaryPreferred` (page 120) is tempting for this use case, it carries some risk: if all secondaries are unavailable and your set has enough *arbiters* to prevent the primary from stepping down, then the primary will receive all traffic from clients. If the primary is unable to handle this load, queries will compete with writes. For this reason, use `secondary` (page 120) to distribute read load to replica sets, not `secondaryPreferred` (page 120).

Read Preferences for Database Commands

Because some *database commands* read and return data from the database, all of the official drivers support full *read preference mode semantics* (page 120) for the following commands:

- `group`
- `mapReduce`⁴
- `aggregate`⁵
- `collStats`
- `dbStats`
- `count`
- `distinct`
- `geoNear`

³ This threshold is configurable. See `localPingThresholdMs` for `mongos` or your driver documentation for the appropriate setting.

⁴ Only “inline” `mapReduce` operations that do not write data support read preference, otherwise these operations must run on the *primary* members.

⁵ Using the `$out` pipeline operator forces the aggregation pipeline to run on the primary.

- geoSearch
- parallelCollectionScan

New in version 2.4: mongos adds support for routing commands to shards using read preferences. Previously mongos sent all commands to shards' primaries.