

TencentDB for MongoDB

Upgrade and Expansion

Product Introduction



Copyright Notice

©2013-2018 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Upgrade and Expansion

 Console Expansion

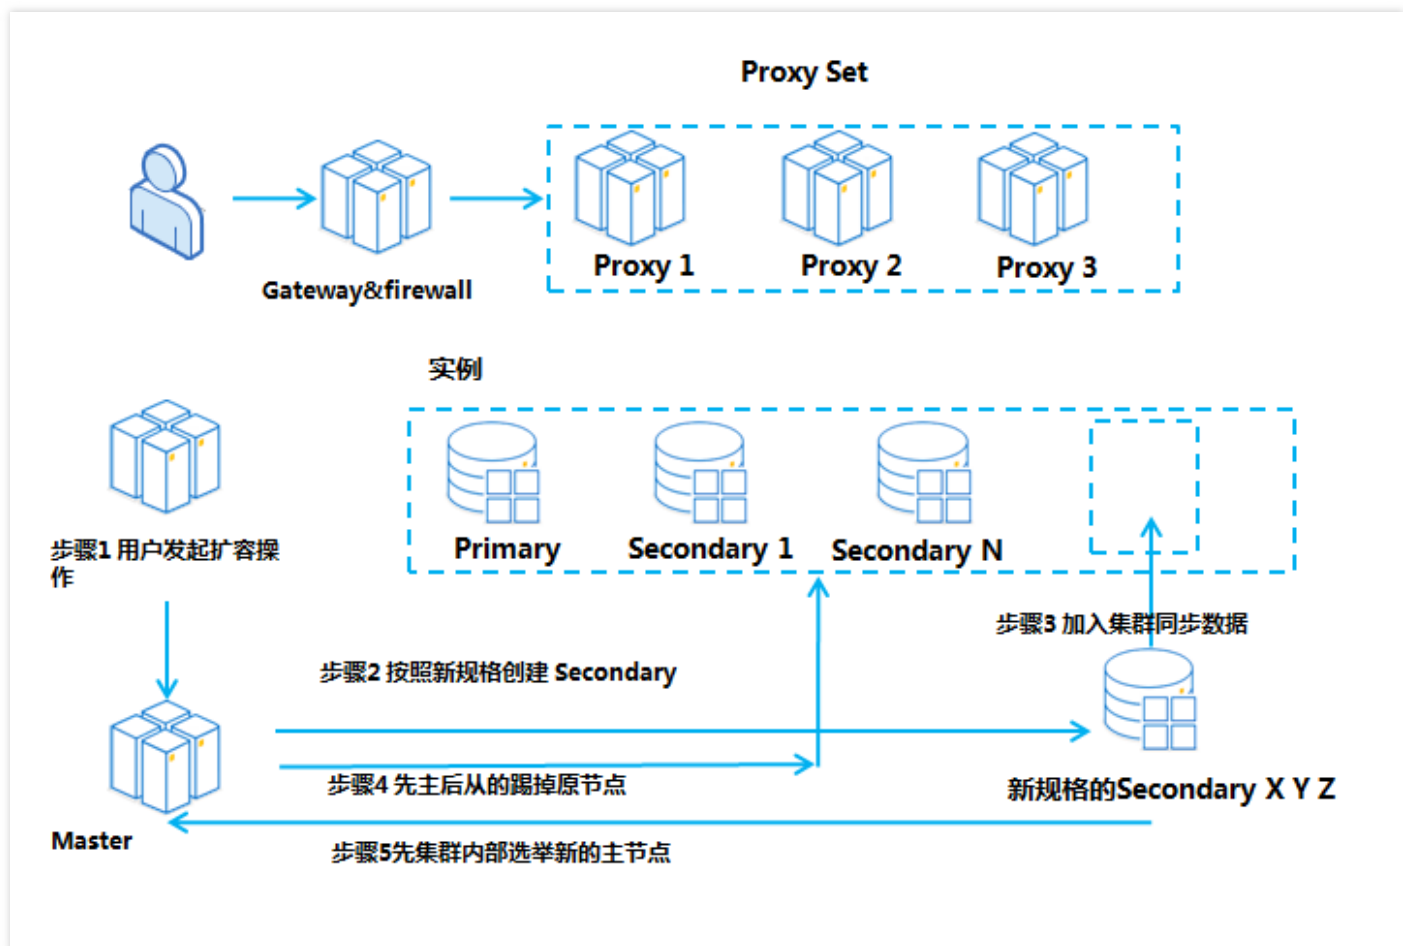
 Online Expansion

Upgrade and Expansion Console Expansion

Last updated : 2018-09-14 15:46:33

Procedure for Online Expansion

1. Initiate expansion operation on the WEB console or via API
2. The system creates an appropriate number of secondary nodes according to the new specifications.
3. Add the new secondary nodes into the cluster instance for data synchronization.
4. After the data synchronization is completed for the last secondary node, kick off the original nodes one by one. Secondary nodes are first kicked off before primary nodes.
5. When there is no primary node in the cluster, a new primary node will be selected.



Online Expansion

Last updated : 2018-09-14 15:47:00

Procedure for Online Expansion

1. Initiate expansion operation on the WEB console or via API
2. The system creates an appropriate number of secondary nodes according to the new specifications.
3. Add the new secondary nodes into the cluster instance for data synchronization.
4. After the data synchronization is completed for the last secondary node, kick off the original nodes one by one. Secondary nodes are first kicked off before primary nodes.
5. When there is no primary node in the cluster, a new primary node will be selected.

