WHAT IS CLAIMED IS:

1. A system, comprising:

a plurality of storage nodes, wherein each of the plurality of storage nodes maintains respective replicas of a plurality of data volumes in block-based storage devices implemented at the storage node, wherein a write request received from a client for a data volume of the plurality of data volumes with replication enabled is completed at each respective replica of the data volume prior to acknowledging the write request to the client; and

each of the plurality of storage nodes, configured to:

determine that a peer storage node of the plurality of storage nodes maintaining a same replica of a particular data volume of the plurality of data volumes as is maintained at the storage node is unavailable for replication;

in response to determining that the peer storage node maintaining the same replica of the particular data volume as is maintained at the storage node is unavailable, disable replication for the particular data volume, comprising:

complete subsequently received write requests from the client for the particular data volume at the storage node; and

acknowledge the completed subsequently received write requests from the client for the particular data volume at the storage node as complete without completing the subsequently received write requests from the client for the particular data volume at any other storage node of the plurality of storage nodes.

2. The system of claim 1, further comprising:

a control plane that authorizes enablement and disablement of replication for the plurality of data volumes;

wherein each storage node is configured to perform said disabling replication for the particular data volume in response to receiving authorization to disable replication from the control plane; and

the control plane, configured to:

in response to receiving a request from a storage node for one or more candidate storage nodes to maintain another replica of the particular data volume:

determine that replication for the particular data volume is to be enabled; and

in response to determining that replication for the particular data volume is to be enabled, send respective identifiers for the one or more candidate storage nodes of the plurality of storage nodes to the requesting storage node in order for the requesting storage node to replicate the particular data volume to at least one of the one or more candidate storage nodes.

3. The system of claim 2, further comprising:

the requesting storage node, configured to:

receive the respective identifiers for the one or more candidate storage nodes;

based, at least in part, on the candidate identifiers, identify a particular one of the one or more candidate storage nodes to maintain another replica of the data volume;

perform one or more replication operations to update the other replica of the data volume in order to duplicate the replica of the particular data volume maintained at the requesting storage node; and

enable replication for the particular data volume such that write requests received at the requesting storage node are sent to the other storage node for completion prior to acknowledging the write requests as complete.

4. The system of claim 1, wherein the plurality of storage nodes together implement a network-based block-based storage service, wherein write requests are received from one or more virtual compute instances implemented by a network-based virtual compute service, wherein the network-based block-based storage service and the network-based virtual compute service are implemented together as part of a same network-based services provider network.

5. A method, comprising:

performing, by one or more computing devices:

determining that a storage node of a plurality of storage nodes maintaining respective replicas of a data volume is unavailable for replication, wherein write requests received for the data volume are completed at each of the plurality of storage nodes in order to perform replication for the data volume, and wherein write requests are acknowledged as complete when complete at each storage node of the plurality of storage nodes; and

in response to said determining that the storage node maintaining a replica of the data volume is unavailable, processing subsequently received write requests for the data volume according to a reduced durability state for the data volume such that replication for the data volume is disabled for the storage node, comprising:

for each of the subsequently received write requests for the data volume:

completing the subsequently received write request at two or more storage nodes of the plurality of storage nodes available for replication, wherein the write request is completed at a fewer number of storage nodes than the plurality of storage nodes; and

in response to completing the subsequently received write request at the two or more storage nodes that are fewer in number than the plurality of storage nodes previously used to acknowledge the write requests, acknowledging the subsequently received write request as complete.

6. The method of claim 5, wherein the plurality of storage nodes includes a master storage node and a slave storage node, wherein the storage node that is determined to be unavailable for replication is a master storage node, wherein the write requests for the data volume are received at the master storage node, wherein the master storage node sends the write requests received at the master storage node to the slave storage node to be completed at the slave storage node, and wherein the method further comprises:

in response to said determining that the storage node is unavailable, promoting the slave storage node to be a new master storage node for the data volume, wherein the subsequently received write requests for the data volume are received at the new master storage node;

wherein said completing the subsequently received write request at the one or more storage nodes of the plurality of storage nodes available for replication is performed such that the subsequently received write request is completed at the new master storage node without completing the write request at any other storage node of the plurality of storage nodes.

7. The method of claim 5, further comprising:

determining that replication for the data volume is to be enabled for another storage node in order to process write requests for the data volume according to an increased durability state; and

in response to determining that replication for the data volume is to be enabled:

identifying the other storage node to maintain another replica of the data volume;

performing one or more replication operations to update the other replica of the data volume to replicate the data volume as maintained at the one or more storage nodes; and

enabling replication for the other storage node such that additional write requests are completed at the other storage node prior to acknowledging the additional write requests as complete.

8. The method of claim 7, wherein the other storage node is different from the storage node that is determined to be unavailable, and wherein the one or more replication operations performed to update the other replica send a copy of all of the data volume maintained at the one or more storage nodes to be stored as the other replica of the data volume.

9. The method of claim 7, wherein the plurality of storage nodes is implemented as part of a larger collection of storage nodes storing a plurality of different data volumes, wherein

said determining that replication for the data volume is to be enabled is based, at least in part, on one or more of:

a storage capacity among the larger collection of storage nodes to maintain the other replica of the data volume;

network utilization among the larger collection of storage nodes; or

a state of durability for the data volume.

10. The method of claim 9, wherein at least some of the one or more computing devices implement a control plane for the larger collection of storage nodes, and wherein said determining that replication for the data volume is to be enabled is performed by the control plane.

11. The method of claim 7, wherein the plurality of storage nodes is implemented as part of a larger collection of storage nodes storing a plurality of different data volumes, wherein write requests for at least some of the plurality of different data volumes are processed according to a respective reduced durability state, and wherein said determining that replication for the data volume is to be enabled is based, at least in part, on a prioritization scheme for increasing durability for the at least some data volumes with the respective reduced durability state.

12. The method of claim 7, wherein the plurality of storage nodes is implemented as part of a larger collection of storage nodes storing a plurality of different data volumes, wherein said determining that the storage node is unavailable, said processing the subsequently received write requests for the data volume according to the reduced durability state, said determining that replication is to be enabled, said identifying the other storage node, said performing the one or more replication operations, and said enabling replication for the other storage are performed for at least some storage nodes storing different ones of the plurality of different data volumes.

13. The method of claim 5, wherein the plurality of storage nodes is part of a larger collection of storage nodes implementing a network-based block-based storage service, wherein write requests are received from one or more virtual compute instances implemented by a network-based virtual compute service, wherein the network-based block-based storage service

and the network-based virtual compute service are implemented together as part of a same provider network.

14. A non-transitory, computer-readable storage medium, storing program instructions that when executed by one or more computing devices cause the one or more computing devices to implement:

determining that a storage node of a plurality of storage nodes maintaining respective replicas of a data volume is unavailable for replication, wherein write requests received for the data volume are completed at each of the plurality of storage nodes in order to perform replication for the data volume, wherein write requests are acknowledged as complete when complete at each storage node of the plurality of storage nodes; and

in response to said determining that the storage node is unavailable, processing subsequently received write requests for the data volume according to a reduced durability state for the data volume such that replication for the data volume is disabled for the storage node, comprising:

for each of the subsequently received write requests for the data volume:

completing the subsequently received write request at two or more storage nodes of the plurality of storage nodes available for replication, wherein the write request is completed at a fewer number of storage nodes than the plurality of storage nodes; and

in response to completing the subsequently received write request at the two or more storage nodes that are fewer in number than the plurality of storage nodes previously used to acknowledge the write requests, acknowledging the subsequently received write request as complete.

15. The non-transitory, computer-readable storage medium of claim 14, wherein the plurality of storage nodes includes a master storage node and a slave storage node, wherein the storage node determined to be unavailable is the slave storage node, wherein the master storage

node sends the write requests received at the master storage node from a client to the slave storage node to be completed at the slave storage node, and wherein said completing the subsequently received write request at the one or more storage nodes of the plurality of storage nodes available for replication is performed such that the subsequently received write request is completed at the master storage node without completing the write request at any other storage node of the plurality of storage nodes.

16. The non-transitory, computer-readable storage medium of claim 14, wherein the program instructions further cause the one or more computing devices to implement:

determining that replication for the data volume is to be enabled for another storage node in order to process write requests for the data volume according to an increased durability state; and

in response to determining that replication for the data volume is to be enabled:

identifying the other storage node to maintain another replica of the data volume;

performing one or more replication operations to update the other replica of the data volume to replicate the data volume as maintained at the one or more storage nodes; and

enabling replication for the other storage node such that additional write requests are completed at the other storage node prior to acknowledging the additional write requests as complete.

17. The non-transitory, computer-readable storage medium of claim 16, wherein the other storage node is the storage node, wherein the other replica of the data volume is the respective replica of the data volume, and wherein, in said performing the one or more replication operations to update the other replica of the data volume, the program instructions cause the one or more computing devices to implement:

sending data chunks of the data volume that were modified subsequent to disabling replication for the data volume for the storage node in order to update the respective replica of the data volume.

18. The non-transitory, computer-readable storage medium of claim 16, wherein the plurality of storage nodes is implemented as part of a larger collection of storage nodes storing a plurality of different data volumes, and wherein said determining that replication for the data volume is to be enabled is based, at least in part, on one or more of:

a storage capacity among the larger collection of storage nodes to maintain the other replica of the data volume;

network utilization among the larger collection of storage nodes; or

a state of durability for the data volume.

19. The non-transitory, computer-readable storage medium of claim 16, wherein the plurality of storage nodes is implemented as part of a larger collection of storage nodes storing a plurality of different data volumes, wherein write requests for at least some of the plurality of different data volumes are processed according to a respective reduced durability state, and wherein said determining that replication for the data volume is to be enabled is based, at least in part, on a prioritization scheme for increasing durability for the at least some data volumes with the respective reduced durability state.

20. The non-transitory, computer-readable storage medium of claim 14, wherein the plurality of storage nodes is part of a larger collection of storage nodes implementing a network-based block-based storage service, wherein write requests are received from one or more virtual compute instances implemented by a network-based virtual compute service, wherein the network-based block-based storage service and the network-based virtual compute service are implemented together as part of a same provider network.