WHAT IS CLAIMED IS:

1. A system implementing a storage node, comprising:

a plurality of block-based storage devices maintaining a plurality of data volumes;

a system memory, comprising a page cache, wherein the page cache is periodically or aperiodically flushed to the plurality of block-based storage devices in order to persistently update the plurality of data volumes;

a persistent storage device maintaining a page cache write log maintaining a plurality of log records describing updates to the page cache, wherein the plurality of log records describe a state of the page cache to be restored to the page cache in the event of a system failure resulting in a loss of data in the page cache;

a log management module, configured to:

monitor a rate of incoming write requests to the storage node that are to be logged to the page cache write log;

based, at least in part, on the rate of incoming write requests, dynamically modify a durability property for at least one data volume of the plurality of data volumes, comprising:

enable page cache write logging for the at least one data volume;

or

disable page cache write logging for the at least one data volume.

2. The system of claim 1, wherein, to dynamically modify the durability property for the at least one data volume, the log management module is further configured to:

select the at least one data volume to modify based, at least in part, on:

a respective size for each of the plurality of data volumes; or

a respective amount of write requests directed toward each of the plurality of data volumes.

3. The system of claim 1, wherein the log management module is further configured to:

receive a write request for a particular data volume;

determine that page cache write logging is disabled for the particular data volume; and

in response to updating the page cache according to the write request, acknowledge the write request as complete without storing a log record for the write request in the page cache write log.

4. The system of claim 1, wherein the storage node is implemented as one of a plurality of storage nodes implementing a network-based block-based storage service, wherein the network-based block-based storage service maintains data volumes for a plurality of different clients.

5. A method, comprising:

performing, by one or more computing devices:

maintaining, at a storage node, a page cache write log in a persistent storage device storing a plurality of log records each describing updates to a page cache at the storage node according to write requests for a plurality of data volumes maintained at the storage node;

monitoring a rate of incoming write requests to the storage node that are to be logged to the page cache write log;

based, at least in part on the rate of incoming write requests, dynamically modifying a durability property for at least one data volume of the plurality of data volumes, comprising:

enabling page cache write logging for the at least one data volume; or

disabling page cache write logging for the at least one data volume.

6. The method of claim 5, wherein said dynamically modifying the durability property for at least one data volume, further comprising:

selecting the at least one data volume to modify based, at least in part, on:

a respective size for each of the plurality of data volumes; or

a respective amount of write requests directed toward each of the plurality of data volumes.

7. The method of claim 6, wherein the storage node is one of a plurality of storage nodes implementing a block-based storage system, wherein different ones of the plurality of storage nodes maintain respective replicas of the plurality of data volumes maintained at the storage node, wherein said selecting the at least one data volume to modify is further based, at least in part, on respective durability for the respective replicas of the plurality of data volumes maintained at the different storage nodes.

8. The method of claim 5, wherein the durability property for all of the plurality of data volumes at the storage node is modified to be the same.

9. The method of claim 5, further comprising:

receiving a write request for a particular data volume of the plurality of data volumes;

determining that page cache write logging is disabled for the particular data volume; and

in response to updating the page cache according to the write request, acknowledging the write request as complete without storing a log record for the write request in the page cache write log.

10. The method of claim 9, further comprising:

receiving another write request for another data volume of the plurality of data volumes;

determining that page cache write logging is enabled for the other data volume;

in response to determining that page cache write logging is enabled for the other data volume:

updating the page cache according to the other write request;

storing a log record describing the update to the page cache for the other write request in the page cache write log; and

in response to updating the page cache according to the other write request and storing the log record describing the update to the page cache for the other write request in the page cache write log, acknowledging the other write request as complete.

11. The method of claim 5,

wherein said enabling page cache write logging for the at least one data volume is performed in response to determining that the rate of incoming write requests is below a logging rate floor; and

wherein said disabling page cache write logging for the at least one data volume is performed in response to determining that the rate of incoming write requests exceeds a logging rate ceiling.

12. The method of claim 5, further comprising:

upon recovery from a system failure:

identifying data volumes of the plurality data volumes with page cache write logging enabled as current; and

making the identified data volumes available for processing input/output (I/O) requests.

13. The method of claim 5, wherein the storage node is one of a plurality of storage nodes implementing a network-based block-based storage service, wherein write requests are received from virtual compute instances implemented by a network-based virtual compute service, wherein the network-based blocks storage service and the network-based virtual compute service are implemented together as part of the same network-based services 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:

maintaining, at a storage node, a page cache write log in a persistent storage device storing a plurality of log records each describing updates to a page cache at the storage node according to write requests for a plurality of data volumes maintained at the storage node;

monitoring a rate of incoming write requests to the storage node that are to be logged to the page cache write log;

based, at least in part on the rate of incoming write requests, dynamically modifying a durability property for at least one data volume of the plurality of data volumes, comprising:

enabling page cache write logging for the at least one data volume; or

disabling page cache write logging for the at least one data volume.

15. The non-transitory, computer-readable storage medium of claim 14, wherein, in said dynamically modifying the durability property for at least one data volume, the program instructions cause the one or more computing devices to implement:

selecting the at least one data volume to modify based, at least in part, on:

a respective size for each of the plurality of data volumes; or

a respective amount of write requests directed toward each of the plurality of data volumes.

16. The non-transitory, computer-readable storage medium of claim 15, wherein the storage node is one of a plurality of storage nodes implementing a block-based storage system, wherein different ones of the plurality of storage nodes maintain respective replicas of the plurality of data volumes maintained at the storage node, wherein said selecting the at least one data volume to modify is further based, at least in part, on respective durability for the respective replicas of the plurality of data volumes maintained at the different storage nodes.

17. The non-transitory, computer-readable storage medium of claim 14,

wherein said enabling page cache write logging for the at least one data volume is performed in response to determining that the rate of incoming write requests is below a logging rate floor;

wherein said disabling page cache write logging for the at least one data volume is performed in response to determining that the rate of incoming write requests exceeds a logging rate ceiling.

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

receiving a write request for a particular data volume of the plurality of data volumes;

determining that page cache write logging is disabled for the particular data volume; and

in response to updating the page cache according to the write request, acknowledging the write request as complete without storing a log record for the write request in the page cache write log.

19. The non-transitory, computer-readable storage medium of claim 18, wherein the program instructions further cause the one or more computing devices to implement:

receiving another write request for another data volume of the plurality of data volumes;

determining that page cache write logging is enabled for the other data volume;

in response to determining that page cache write logging is enabled for the other data volume:

updating the page cache according to the other write request;

storing a log record describing the update to the page cache for the other write request in the page cache write log; and

in response to updating the page cache according to the other write request and storing the log record describing the update to the page cache for the other write request in the page cache write log, acknowledging the other write request as complete.

20. The non-transitory, computer-readable storage medium of claim 14, wherein the storage node is implemented as one of a plurality of storage nodes implementing a network-based block-based storage service, wherein the network-based block-based storage service maintains data volumes for a plurality of different clients.