

Managing File Space and Database Capacity

Xpand monitors the amount of space available within your cluster and proactively warns of potential capacity issues. The thresholds for determining a cluster's capacity are configurable and described below.

- [Types of Storage](#)
- [Checking Storage Utilization](#)
- [Global Variables](#)



See Also: [Allocating Disk Space for Fault Tolerance and Availability](#) or [In-Memory Tables](#) for guidance in allocating memory and configuring alerts.

Types of Storage

To understand how to manage the Device and Database Utilization, one must first understand how Xpand allocates disk space. Xpand creates and allocates space in two different files:

device1 (main storage)

The main device1 storage is used for all database data, undo logs, temporary tables, binlogs, Xpand system tables, as well as temporary storage used for query execution. The initial size of the device1 file is auto-detected by the Xpand installer, but can also be configured manually. Post-installation, the device1 file's size can be extended using [ALTER CLUSTER RESIZE DEVICES](#). To decrease the size of the device1 file, see [Decreasing device1 Size](#).

Xpand expects the device1 file to be the same size on every node. By default, on database startup, Xpand will automatically attempt to resize the device1 file on each node to match the largest device1 file in the cluster. To disable this feature, set `device_auto_resize_to_largest=false`.

Temporary storage is used for sorting and grouping of large query results and is stored in device1. There are two global variables to control temp space usage:

- `device_temporary_space_limit_bytes` limits the amount of space usable for temporary storage.
- `device_temporary_space_preallocate_bytes` specifies the amount of space that will be pre-allocated for temp space (guaranteed for use by temp).

Setting `device_temporary_space_limit_bytes` allows additional temp space to be used, but does not guarantee additional space will be available for temp. Increasing these values takes effect immediately, while decreasing it takes effect after database restart.



Prior to Xpand 9.2, temp space was stored in a separate file called device1-temp, but with v9.2, temp space is now managed within the device1 file.

device1-redo (write-ahead log)

The write-ahead log (WAL) is stored in the device1-redo file. This size of this file is 4GB and is not configurable.

Checking Storage Utilization

See how much space is in use by using [The CLX Command-Line Administration Tool](#)

```

shell> /opt/clustrix/bin/clx space
nid | Hostname | Status
| Undo | Perm | WAL | Temp | Used | DB Total | FS Free
+-----+-----+-----+-----+-----+-----+-----+-----+
16 | eukanuba003 | OK | 321.8M (0.04%) | 674.7G (79.4%) | 1024.0M (0.12%) | 0 (0.00%) | 760.1G
(89.4%) | 850.0G | 113.9G
17 | karma183 | OK | 313.5M (0.04%) | 664.6G (78.2%) | 1024.0M (0.12%) | 0 (0.00%) | 750.1G
(88.2%) | 850.0G | 113.9G
18 | eukanuba002 | OK | 324.3M (0.04%) | 669.5G (78.8%) | 1024.0M (0.12%) | 0 (0.00%) | 755.0G
(88.8%) | 850.0G | 113.9G
19 | eukanuba001 | OK | 339.7M (0.04%) | 671.0G (78.9%) | 1024.0M (0.12%) | 0 (0.00%) | 756.4G
(89.0%) | 850.0G | 113.9G
20 | eukanuba005 | OK | 277.3M (0.03%) | 668.7G (78.7%) | 1024.0M (0.12%) | 0 (0.00%) | 754.1G
(88.7%) | 850.0G | 113.9G
21 | eukanuba004 | OK | 420.3M (0.05%) | 678.6G (79.8%) | 1024.0M (0.12%) | 0 (0.00%) | 764.1G
(89.9%) | 850.0G | 113.9G
22 | eukanuba006 | OK | 397.0M (0.05%) | 670.4G (78.9%) | 1024.0M (0.12%) | 0 (0.00%) | 755.9G
(88.9%) | 850.0G | 113.9G
23 | karma184 | OK | 479.9M (0.06%) | 674.8G (79.4%) | 1024.0M (0.12%) | 0 (0.00%) | 760.3G
(89.5%) | 850.0G | 113.9G
+-----+-----+-----+-----+-----+-----+-----+-----+
(89.1%) | 6.6T | 910.9G | 2.8G (0.04%) | 5.2T (79.0%) | 8.0G (0.12%) | 0 (0.00%) | 5.9T

```

Global Variables

The default values for these global variables are optimal for most workloads.

Variable	Description	Default
device_auto_resize_to_largest	Automatically resize all (online) devices in the cluster to match the largest device	true
device_temporary_space_limit_bytes	Maximum number of bytes allowed to be used for temporary containers.	5368709120
device_temporary_space_preallocate_bytes	The amount of space that will be pre-allocated for temporary storage	5368709120

Database Storage Thresholds

Global variables establish the database storage thresholds for a cluster. When the first level of thresholds are exceeded, alerts are sent. If storage utilization continues to increase, user queries will begin to fail once the next set of thresholds are exceeded. Finally, if storage utilization continues to grow, system queries (including for critical internal processes) will be killed. Once the database is completely full, the database may become inoperable. See [Issue Resolution](#) below for suggestions on freeing space.

Variables

The following variables are used to set thresholds for device1 utilization.

Variable	Description	Default Value	Allowed Values
databasefull_message_interval_s	Database almost full message interval in seconds.	120	Minimum: 10 Maximum: 600
databasefull_user_warn_percentage	Warn about user queries when space usage surpasses this percentage.	80	Minimum: 50 Maximum: databasefull_user_error_percentage - 1
databasefull_user_error_percentage	Fail user queries when space usage surpasses this percentage.	90	Minimum: databasefull_user_warn_percentage + 1 Maximum: databasefull_system_warn_percentage - 1
databasefull_system_warn_percentage	Warn about system queries when space usage surpasses this percentage.	95	Minimum: databasefull_user_error_percentage + 1 Maximum: databasefull_system_error_percentage - 1
databasefull_system_error_percentage	Fail user queries when space usage surpasses this percentage.	97	Minimum: databasefull_system_warn_percentage + 1 Maximum: >99

User queries are transactions which originate with an end user whereas system queries are internal Xpand processes such the [Rebalancer](#), [binlog deletes](#), [statd](#), etc.

Alert Messages

The following alerts are triggered when the corresponding global variable is exceeded. This is evaluated each time Xpand allocates space and any alerts necessary are sent every `databasefull_message_interval_s` seconds. If multiple alerts are detected, only the most critical will appear. To learn more about Xpand's Alerter, see [Database Alerts](#).

Global Variable Evaluated	Alert Triggered	Level	Description	Message Shown
<code>databasefull_user_warn_percentage</code>	<code>DATABASE_SPACE_LOW</code>	warning	Database space low	Database space is nn% used. Soon user queries will fail.
<code>databasefull_user_error_percentage</code>	<code>DATABASE_SPACE_EXTRME</code>	warning	Database space extreme	Database space is nn% used. User queries will now fail.
<code>databasefull_system_warn_percentage</code>	<code>DATABASE_SPACE_CRITICAL</code>	critical	Database space critical	Database space is nn% used. User queries will fail, and soon system queries will fail.
<code>databasefull_system_error_percentage</code>	<code>DATABASE_SPACE_EXHAUSTED</code>	critical	Database space exhausted	Database space is nn% used. User queries and system queries will now fail.

Resolving Low Space Issues

When you receive any of the alerts above, some action will be necessary to prevent the capacity of device1 from reaching the next threshold.

Some resolutions to consider:

- Add nodes to the cluster by [Expanding Your Cluster's Capacity - Flex Up](#).
- Increase available space on the cluster by:
 - [Trimming Binlogs](#)
 - Deleting data
- Enlarge the size of the device1 file on all nodes by using [ALTER CLUSTER RESIZE DEVICES](#).
- Terminate and reschedule long running transactions such as ALTERs, Backups, and long-running transactions. These halt garbage collection and cause the undo log to temporarily grow in size.

If you need assistance, please contact [Xpand Support](#).