

Logical Unit Status

NOTE: Please try to keep as much as possible the following conventions:

- Oracle LUNs: numbers between 10 and 30.
- Data LUNs (various NFS mount points, data writing volumes, backups): numbers between 40 and 100. (TODO: clarify this even more).
- Test LUNs: numbers over 100. This can be deleted at any time.
- Tiers 3,4,5,6,7 are for even file writing; tiers 8,9,10,11 for databases. Tier 2 user related partitions. For testing parts of tier1 can be used. For long term data storage (low IO and big space requirements) the database tiers can be used.

Things to keep in mind:

- Use 4kB block sizes whenever possible
- Make sure when creating LUNs the space does not get fragmented. E. g. do not create two LUNs and delete then the first one created. The space in the tier will be split in two by the last created LUN.
- Do not use LUN groups.

The following LUNs are available. Below are some additional comments:

LUN	Label	Owner	Status	Capacity (MB)	BS	Tiers	Tier list
11	rac_voting02	1	Ready	100	512	1	9
12	rac_voting03	1	Ready	100	512	1	10
13	rac_voting04	1	Ready	100	512	1	8
14	rac_voting05	1	Ready	100	512	1	9
15	rac_voting06	1	Ready	100	512	1	10
20	rac_ocr01	1	Ready	250	512	1	8
21	rac_ocr02	1	Ready	250	512	1	9
23	rac_spfile	1	Ready	50	512	1	10
30	rac_4Tiers	1	Ready	500000	512	4	8 9 10 11
31	rac_asm01	1	Ready	20000	512	1	8
32	rac_asm02	1	Ready	20000	512	1	9
33	rac_asm03	1	Ready	20000	512	1	10
34	rac_asm04	1	Ready	20000	512	1	11
61	nfs_home01	1	Ready	1048576	4096	1	2
62	recovery1	1	Ready	11000	512	1	1
63	logbook01	1	Ready	100000	4096	1	2
64	backuplun01	1	Ready	200000	4096	1	2
65	backuplun02	1	Ready	200000	4096	1	2
66	backuplun03	1	Ready	200000	4096	1	2
67	backuplun04	1	Ready	200000	4096	1	2
70	histo01	1	Ready	2097120	512	1	1
75	tmpdaqlun01	1	Ready	1048576	4096	2	3 5
81	data_lun01	1	Ready	10485760	4096	5	3 4 5 6 7
91	data_mdln01	1	Ready	2000	4096	1	6
92	data_mdln02	1	Ready	2000	4096	1	7
93	data_journal	1	Ready	2000	4096	1	4
94	data_journal	1	Ready	2000	4096	1	4
100	test01	1	Ready	100	4096	1	1

- 11, 12 - LUNs create for the initial voting disk of the oracle cluster. When the file system LUN 10 (now deleted) got corrupted LUNs 13-15 were created as a replacement for the voting disks and the data was restored from backups. Might be useful in the future to see if the voting configurations can be recreated from another voting disk.
- 70 - The name is misleading: only 500 GB are made available for the histograms. The partition is based on LVM. Details:

StorageConfigurationHowTos < LHCb < TWiki

physical volume = /dev/mpath/360001ff0200211010af381f7fff4600 (multipath on top of the storage luns)

volume group = tier1

logical volume = /dev/tier1/histogram_lv

The filesystem is XFS and the LUN is exported to store04.

NOTE: the LUN BS is 512, the XFS BS is 4096 while sector size is 512. This is bad as the journaling information will be flushed in 512 pieces while the minimum disk IO for the storage is 512 x (numbers of disk in a tier - 1). This is not optimal and the LUN + file systems should be recreated at some point. Keep this in mind when using this LUN.

- 100 - For testing the NFS redundancy.
- 81 - This should be recreated at some point and used include fully the tiers 3,4,5,6,7

-- RaduStoica - 10 Mar 2008

This topic: LHCb > StorageConfigurationHowTos

Topic revision: r1 - 2008-03-10 - RaduStoica



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback