FreeNAS 11.1-U4

Intel(R) Xeon(R) CPU L5520 @ 2.27GHz / 72GB RAM

forums.freenas.org
sfcredfox

3G Pool (24 disk / 146GB 3G SAS / 10K) arranged in twelve mirrors
6G Pool (12 disk / 300GB 6G SAS / 10K) arranged in 6 mirrors

Test platform is Vmware's IO Analyzer with 20GB data drive, in a VMFS5 datastore, dedicated for testing, no other load on the datastore/diskpool. ESX servers connected to FreeNAS via round robin (IOPS = 5) dual fabric 4gb/s FC. Each test was run 15 minutes, then immediately again for 60 minutes. The recorded results are from the 60 minute run. LOG device is 80GB Intel S3500 6G. Cache device is Mushkin Chronos 120GB 6G. Shared LOG/CACHE consists of two 55GB partitions on cache SSD, LOG is two 20GB partitions on Intel S3500.

SQL (16k 66%Read 100%Random) / Vmware IO Analyzer
Exchange2007 (8k_55%Read_80%Random) / Vmware IO Analyzer

Dedicated LOG/CACHE Single Test

	Order	IOPS	Read IOPS	Write IOPS	MB/s	Read MB/s	Write MB/s	RLAT/ms	WLAT/ms
SQL	3G	18179	11998	6181	284	187	97	0.37	1.62
	Order	IOPS	Read IOPS	Write IOPS	MB/s	Read MB/s	Write MB/s	RLAT/ms	WLAT/ms
Exchange	3G	20082	11044	9039	157	86	71	0.25	0.86

Dedicated LOG/CACHE Dual Test (Both pools running simultaneous test)

	Order	IOPS	Read IOPS	Write IOPS	MB/s	Read MB/s	Write MB/s	RLAT/ms	WLAT/ms
SQL	3G	15672.2	10343	5329	245	162	83	0.65	1.52
	6G	8854	5842	3011	138	91	47	1.63	2.11
	Order	IOPS	Read IOPS	Write IOPS	MB/s	Read MB/s	Write MB/s	RLAT/ms	WLAT/ms
Exchange	3G	17492.1	9620	7872	137	75	62	0.3	1
	6G	17069.8	9388	7682	133	73	60	0.3	1.28

Shared LOG/CACHE Dual Test (Both pools running simultaneous test)

SQL

Order	IOPS	Read IOPS	Write IOPS	MB/s	Read MB/s	Write MB/s	RLAT/ms	WLAT/ms
3G	16562	10932	5630	259	171	88	0.62	1.4
6G	16231	10713	5518	254	167	86	0.61	1.49

Exchange

Order	IOPS	Read IOPS	Write IOPS	MB/s	Read MB/s	Write MB/s	RLAT/ms	WLAT/ms
3G	17405	9574	7832	136	75	61	0.29	1.01
6G	16882	9285	7598	132	73	59	0.29	1.09

Presumptive Concolusions: The slight increase in performance running the tests on each pool independently is likely due to having more system ARC available for the test as it was not being split between two needy workloads. However, I do suggest that the system is supporting a total of roughly 32,000-34,000 IOPS of 8K/16K workloads which may align with some SQL and Exchange environments.

The intended focus of testing was trying to determine if there was any quantifiable benefits from using a shared SLOG on this system or if the introduction of multiple pools writing to the devices would adversely increase latency due to contention of multiple simultaneous writes.

Under the conditions that the workloads were heavily writing to the pools (small/synchronous writes), there was a noticeable, albeit only half a millisecond, performance gain of striping the SLOG, while not introducing a higher write latency. The dedicated single SLOG tests showed higher latency, likely due to only having one device to write to at a time, or reaching the throughput limit of the device, although, the testing was conducted with smaller IO sizes so as not to allow device throughput to be a significant factor, and maxing out the SSDs was avoided.

Overall conclusion: In the case of limited funds, where the workloads are conducive, it can be a benefit to share a pool/collection of devices for SLOG rather than dedicating a single device to each pool. In the production world, funding should likely not be a limiting factor and the appropriate number of devices would be purchased for each pool to meet its performance requirements, thus making sharing them unnecessary. Additionally, it would be more likely/advisable to use mirrored pairs for a multi-device SLOG for greater redundancy/reliability. This solution, if ever done, is really only suitable for the lab environment where the greatest performance solution is sought for the lowest possible cost.