Performance Analysis

Synthetic benchmarks: IOR, bonnie++, mdtest

Report contains performance studies based on IOR, mdtest and bonnie++ benchmarks applied on the NOTUR systems: hexagon, titan, njord, stallo.

2010
Alexander Oltu UniBCCS

Content

1 Introduction

These performance studies can be used to select NOTUR system for the type of application based on the file system access pattern. It contains two metadata benchmarks and one read-write test.

The metadata benchmarks measures how fast file system can create, update and delete files.

The read-write test shows how fast clients can write and read. N-to-1 file access method has been used. In this test all clients will write into a single file. IOR, the benchmark used in this test, supports MPIIO, HDF5 and POSIX interfaces. The POSIX interface has been selected to eliminate the performance impact of additional layers like MPIIO or HDF5.

All benchmarks have been executed on live systems. This means that the system load during the test can have affected the results. Therefore, these numbers reflect the possible performance during regular usage of the machines and not the peak performance.

2 Performance studies

The tests performed included two metadata benchmarks, bonnie++ and mdtest, as well as one file system read-write test using the IOR application. The metadata performance is measured in IOPS (number of operations per second).

2.1 Bonnie++

This test was not executed on njord, due to lack of a way to start multiple non-MPI applications at the same time.

The bonnie++ source code was slightly modified to support the creation of directories with randomly generated names.

The test has been performed using 16000 zero-sized files, they were created, 'stat'ed and deleted by each node. The default striping pattern have been used on the Lustre filesystems (hexagon and stallo).

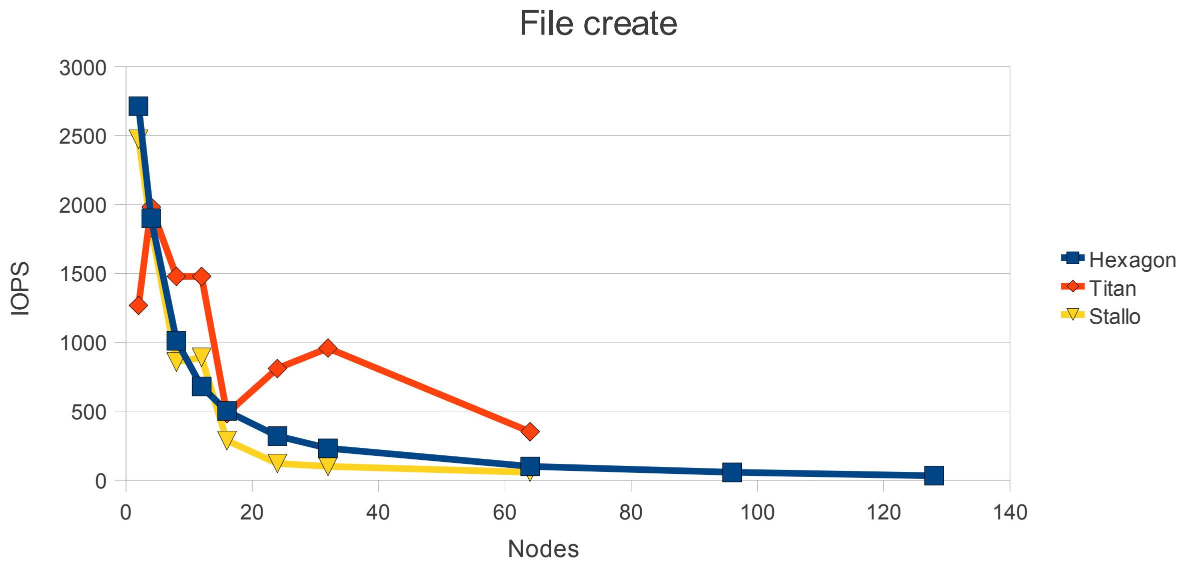

For a low amount of nodes the Lustre based machines had the best file create performance. In addition, results show that the performance on the Lustre based machines decreases proportionally with increasing number of clients. Due to the dedicated MDS (meta data server) the Lustre file system for each file create and file delete operation has to perform two operations: the first on the metadata server and the second on the object storage server.

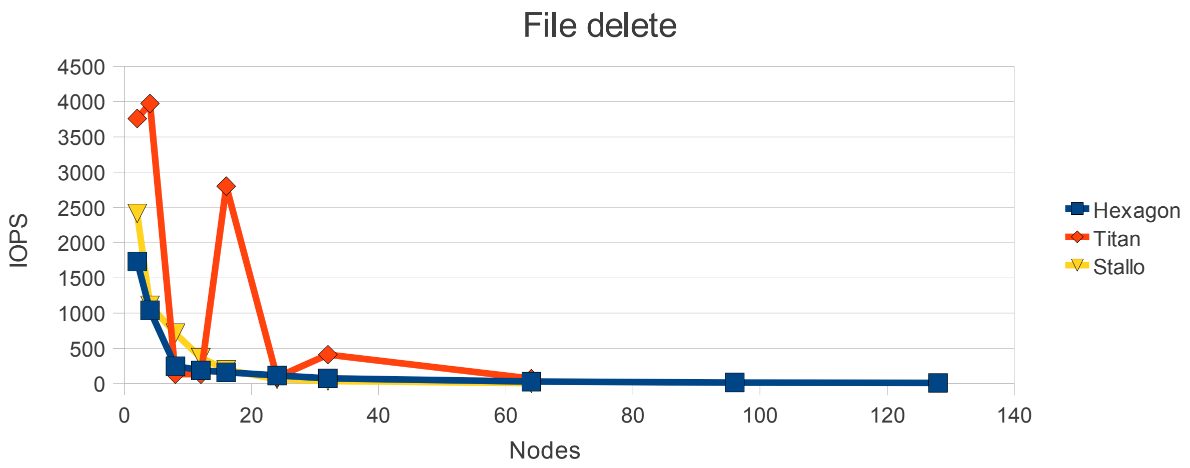

The titan file system has a performance drop in create and delete operations when using 16 nodes, this result can have been affected by ATLAS jobs running on the titan system or possible by the specifics of the GPFS file system.

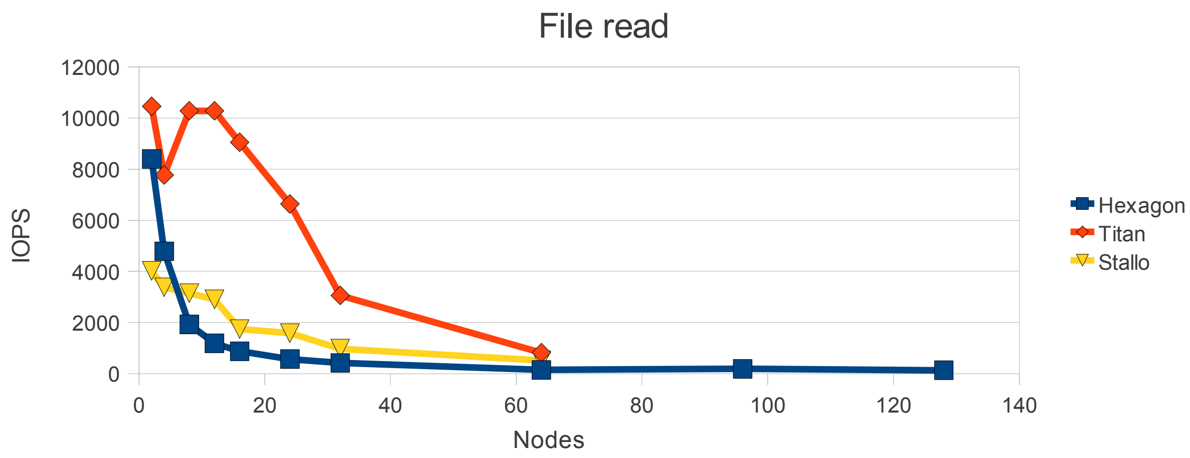

The best read performance in IOPS was on the titans GPFS file system.

A different striping pattern has been tested on the hexagons Lustre file system. Using a single stripe showed slightly higher performance for up to 96 nodes. For the 36 stripe pattern (the maximum for hexagon) the performance was very low. This can be explained by how the Lustre file system creates files (two object operations). Increasing the number of stripes therefore increases the overall number of operations required.

2.2 Mdtest

The mdtest had lower and more typical load pattern. Each client created, 'stat'ed and deleted 10 files. The test has been rerun 200 times for increased accuracy. All results are given in IOPS.

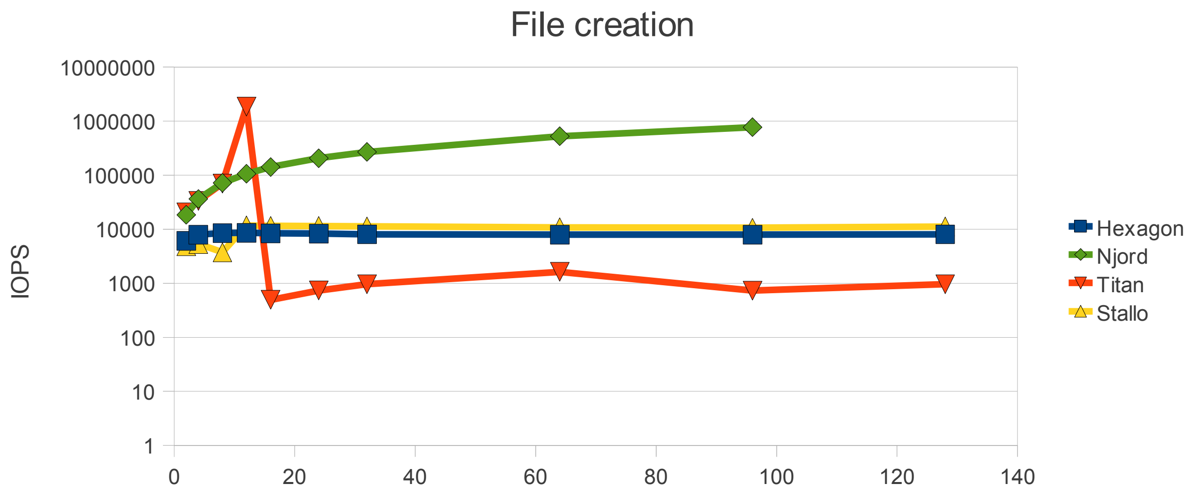

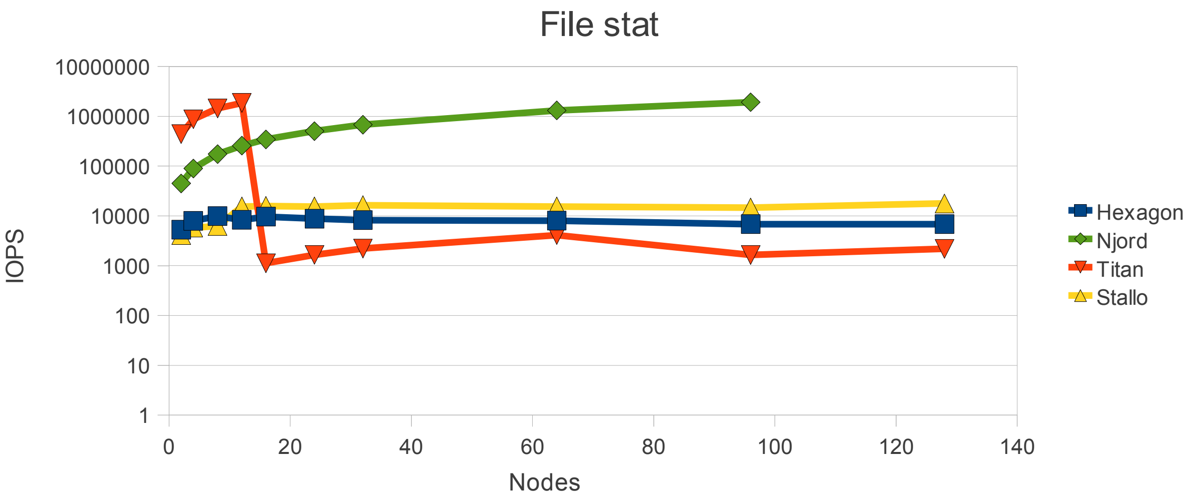

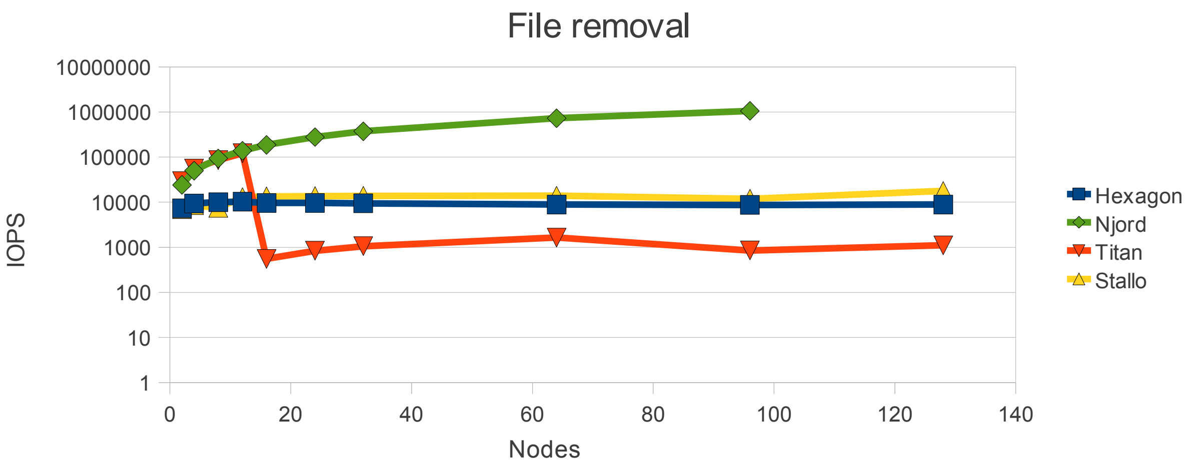

uni BCCS

It is assumed that the GPFS local caching had a higher influence during this lower load (operation on 10 files per client). This can be clearly seen in the results from the njord system. For the titan system GPFS cache has less influence on the performance when using more than 16 clients. This surprisingly different behavior for the two GPFS file systems, can be related to the Linux and AIX implementation, version number or hardware type.

2.3 IOR

In this test each client wrote 1GB into the same file. This operation has then been repeated three times and the average used. All results are represented in MIB/s (mebibytes per second), to get MB/s the number must be multiplied by 1.048.

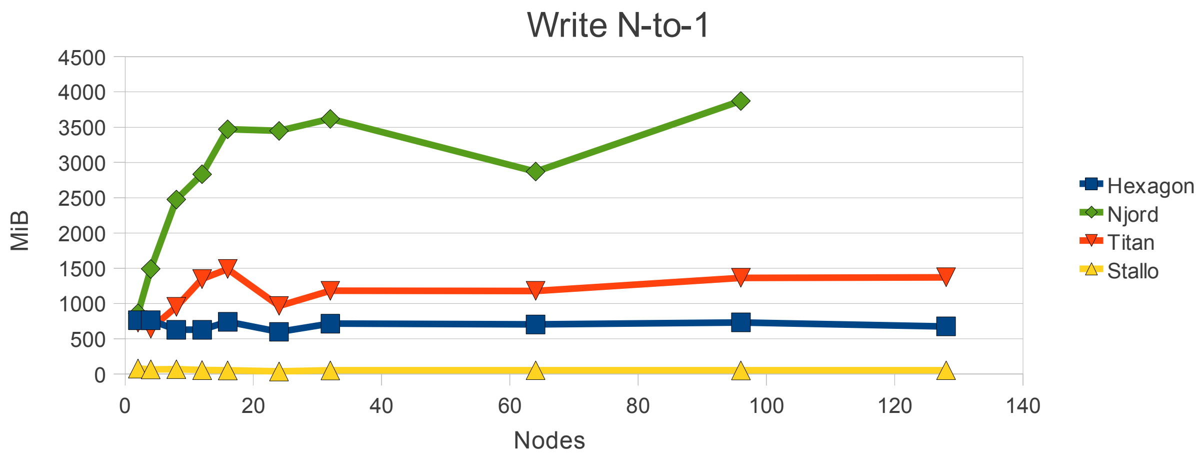

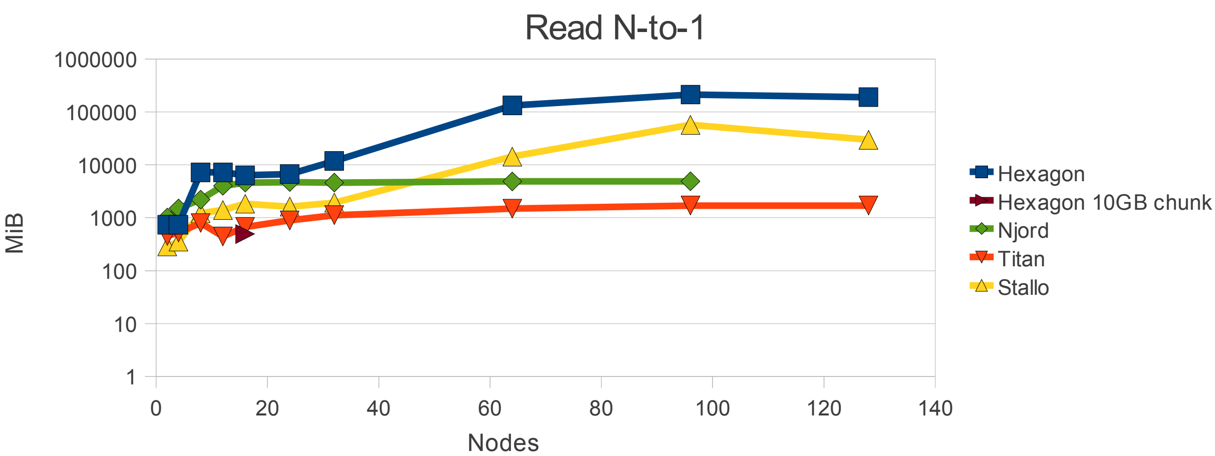

Due to file system space-constrains, the size for each client was selected to be 1GB, this will affect read performance due to caching effects, especially on hexagon. To eliminate this effect the file size must be bigger than the available memory on each node.

In the N-to-1 test njord shows very good write performance. To be sure that this is not caused by the caching effect, two reruns have been performed: one with two nodes and 25GB chunk size for each client and another with 16 nodes and 25GB chunk size for each client. In these reruns the results were nearly identical, showing that caching had minimal effect.

On hexagon, one rerun using 16 nodes and a 10GB chunk size were executed and showed a lower performance due to the reduced cache effect.

The poor write performance on the stallo system can be related to the disk upgrade, but a retest using two nodes on October 18th gave the same result.

3 Appendix A – applied commands and results

3.1 Bonnie++

The version used has patch to create random directory names.

$submit_pref ./$EXE -q -s 0 -n 16:0:0:1 >>run.csv 2>>$stderr

Results:

		Hexagon					**Titan**						**Stallo**					
		Sequential create					Sequential create						Sequential create					
Version	N of files K	Create IOPS	%CP	Read IOPS	%CP	Delete IOPS	Create IOPS	%CP	Read IOPS	%CP	Delete IOPS	%CP	Create IOPS	%CP	Read IOPS	%CP	Delete IOPS	%CP
2	**Nodes**																	
Avg		2713	11.5	8397	39	1733.5	1267.5	3	10456	14	3759.5	9.5	2472.5	17.5	4017	26.5	2414.5	19.5
4	**Nodes**																	
Avg		1899	8.25	4786	21	1040.5	1982.3	4.75	7770	8.75	3972.5	12	1831.8	13.25	3381	26	1110	9
8	**Nodes**																	
Avg		1011.38	4.125	1930	8.5	248.13	1477.4	3.5	10279	11.75	134.75	0	855.88	5.875	3148	24.25	715.25	5.625
12	**Nodes**																	
Avg		679.583	2.75	1184	5.3333	188.58	1477.4	3.5	10279	11.75	134.75	0	890.92	6.083	2900	21.17	367.5	2.917
16	**Nodes**																	
Avg		501	2	880.1	3.75	161.19	482	1	9052	10	2801	8	286.81	2	1745	13.13	196.81	1
24	**Nodes**																	
Avg		319.583	1	568.3	2.0833	116.33	809.17	1.708	6638	6.875	89.542	0	121.08	0.042	1585	11.5	63.818	3
32	**Nodes**																	
Avg		230.938	0.969	420.2	1.2813	76.031	958.28	2.031	3062	2.9688	412.19	0.969	99.5	0	976.47	6.969	37.656	0
64	**Nodes**																	
Avg		99.4688	0	147	0	29.547	350.88	0.516	821.1	0.4375	71.375	0.063	57.938	0	519.72	3.063	5.6563	0
96	**Nodes**																	
Avg		56.6146	0	187	0.1042	17.115												
128	**Nodes**																	
Avg		31.9141	0	131.1	0	11.398												

3.2 Mdtest

```
$submit_pref ./$EXE -n 10 -i 200 -u -t >>$stderr
```

Results:

2	**2 tasks, 20 files/directories**															
Machine:	Hexagon				Njord				Titan				Stallo			
Operation	Max	Min	Mean	Std Dev	Max	Min	Mean	Std Dev	Max	Min	Mean	Std Dev	Max	Min	Mean	Std Dev
Directory creation:	7391.495	2373.619	6512.691	528.08	2966.059	576.699	2492.151	340.715	2257.065	491.05	1904.812	201.206	7312.245	600.636	5616.898	1524.804
Directory stat:	13890.724	3273.348	12756.093	1036.008	44739.243	38321.645	42576.235	1226.178	487709.767	138884.238	394973.148	61599.049	16206.739	2473.421	12407.099	3005.184
Directory removal:	7138.026	1607.598	6299.773	573.745	10081.25	2040.18	8448.471	1788.08	8280.956	952.926	7044.434	876.843	6466.205	1412.224	4683.184	1421.602
File creation:	6961.5	3036.71	6078.82	366.152	19117.156	4101.808	18515.53	1427.394	23366.596	2797.601	20566.258	2333.182	6720.564	2.301	4830.661	1806.152
File stat:	5866.98	1931.079	5258.208	507.409	47393.266	40741.175	45023.55	1124.312	476625.455	60436.657	437427.941	44185.465	6101.693	1198.681	4151.785	1150.475
File removal:	8196.002	74.551	7252.064	917.666	26116.463	1556.675	24068.52	2300.183	32313.59	19742.546	29460.312	1553.567	8972.733	12.15	7035.097	1910.474
Tree creation:	2864.962	786.776	2554.252	309.873	137.383	15.265	91.714	15.427	89.157	25.871	74.976	9.379	2941.307	0.124	1926.154	653.105
Tree removal:	2462.891	320.005	2087.012	319.445	116.999	33.11	96.356	17.584	120.208	43.697	97.91	13.07	2105.574	234.358	1428.582	462.478
4	**4 tasks, 40 files/directories**															
Directory creation:	9496.896	5843.684	8318.959	501.547	5527.913	879.355	4726.254	722.282	4357.832	210.665	3666.9	576.648	11880.198	1436.627	7294.062	3088.702
Directory stat:	20769.022	12783.615	17516.18	1366.677	89098.332	58314.967	85211.896	3485.093	975419.535	338933.657	641290.548	140726.456	27196.006	2907.613	16362.285	7255.106
Directory removal:	13408.9	6333.176	7907.465	548.299	19427.068	1401.687	14213.556	4137.382	15747.34	791.781	13148.707	2017.13	10230.634	5.638	6145.887	2752.005
File creation:	9041.884	6771.833	7958.51	433.86	37600.215	3805.911	36460.605	2543.81	44834.89	261.426	33594.42	10541.59	11132.857	5.745	5328.05	3947.123
File stat:	9128.471	666.233	7963.56	1210.061	94307.004	76748.472	89541.795	2744.81	1111073.907	154344.213	869731.471	126123.287	10654.907	9.529	5798.561	2512.04
File removal:	11229.73	324.491	9356.672	1499.376	52841.625	9172.389	49706.122	3142.354	62906.697	44691.572	55557.238	3728	14624.491	1760.168	8561.942	3441.211
Tree creation:	1550.574	750.726	1388.586	140.886	50.938	15.212	37.814	7.536	39.777	0.457	33.558	5.09	1526.867	121.935	901.141	411.871
Tree removal:	1254.653	719.435	1117.999	98.156	51.643	12.193	34.323	7.66	43.395	1.169	36.744	4.711	1040.512	0.262	594.284	278.748
8	**8 tasks, 80 files/directories**															
Directory creation:	9415.088	6105.913	8706.306	403.46	12678.316	4338.169	9394.572	1228.746	9073.425	3636.391	7535.364	967.542	13209.886	4.89	5142.861	3008.001
Directory stat:	22105.825	16116.442	18987.026	1254.711	175127.516	49056.187	165575.981	9407.837	1775366.772	602413.501	1248642.942	175661.963	31020.091	1542.321	13166.784	6094.61
Directory removal:	9022.191	7057.702	8260.276	333.859	37667.75	4487.086	22969.733	7969.277	30096.36	5718.011	21405.89	5191.719	11111.107	13.402	5180.886	2144.878
File creation:	9165.874	7588.067	8497.728	293.969	74284.773	28069.627	72293.004	3538.47	86951.107	58436.837	70308.307	2709.305	11654.891	5.229	3746.313	2968.771
File stat:	11490.85	5890.153	9875.501	715.741	184365.011	119581.012	174380.881	6491.385	2058554.11	305317.853	1437867.672	265108.582	15171.331	56.653	6261.444	2924.507
File removal:	10921.6	2190.766	10006.856	1020.016	102175.493	11487.703	92869.076	6901.956	123817.092	65421.002	85882.168	5400.682	15402.539	4.028	7452.418	2696.45
Tree creation:	799.982	478.256	725.013	61.211	24.963	12.711	18.436	2.504	20.114	1.013	16.052	2.649	686.803	5.992	298.569	156.077
Tree removal:	1054.905	294.647	559.599	53.377	23.359	9.764	16.986	2.678	19.242	6.4	16.481	1.84	485.452	37.824	176.755	95.662
12	**12 tasks, 120 files/directories**															
Directory creation:	9316.707	7711.653	8802.619	234.715	16987.292	3063.809	14209.07	1826.852	11988.007	611.805	6727.737	2288.06	12367.41	4129.099	11494.95	937.889
Directory stat:	20753.607	12819.104	19058.371	946.132	256401.671	207211.396	243816.335	7075.773	2351946.168	335097.523	1699438.88	280154.588	30127.887	6936.46	27000.139	2832.567
Directory removal:	8900.065	1432.375	8302.44	814.586	54471.481	9141.569	35207.417	10617.722	36992.245	1039.743	17882.088	8756.178	12144.496	7080.089	11131.791	717.065
File creation:	9220.277	3497.026	8581.946	450.806	110183.117	20046.859	106467.623	10811.514	109202.968	6536.238	98781.447	18134.512	13013.664	2050.161	11769.624	941.39
File stat:	11004.099	47.611	8325.211	3519.064	270309.603	236632.102	258078.798	6348.396	2796202.667	585251.721	1855879.71	330959.524	17413.987	355.534	15434.447	1600.041
File removal:	11041.516	1797.187	10227.104	693.392	147816.881	27204.826	138048.503	10984.68	138578.326	11619.912	121517.595	17543.191	14117.482	6929.966	12865.676	800.202
Tree creation:	531.059	139.782	482.872	34.65	15.738	7.227	13.192	1.435	8.89	0.633	5.868	1.554	532.205	46.356	477.161	42.775
Tree removal:	406.701	289.862	368.452	20.371	15.055	6.912	12.381	1.331	8.142	0.887	5.787	1.533	334.554	109.793	301.379	23.803
16	**16 tasks, 160 files/directories**															
Directory creation:	9069.868	4752.55	8496.499	467.677	20642.53	2580.277	17371.028	3167.474	954.983	428.753	764.262	60.106	15061.69	402.193	11910.796	1334.039
Directory stat:	20863.292	9662.486	18512.893	1236.308	343971.625	215368.626	323636.681	11274.451	1584.27	548.3	1233.712	95.795	33898.502	1584.045	28225.585	2802.107
Directory removal:	8775.842	1934.586	7959.384	1062.132	70670.666	9673.629	41570.52	15824.916	619.276	344.953	532.814	28.52	14116.888	9118.424	11800.407	827.02
File creation:	8991.849	3284.948	8397.27	510.067	146782.292	13685.911	141809.687	14623.654	544.642	343.668	494.512	25.54	14002.016	102.107	11594.122	2396.789
File stat:	10737.59	848.576	9635.334	952.884	360219.345	276395.651	343812.443	9293.598	1309.181	559.61	1122.039	72.202	18890.07	834.447	15791.139	2494.785
File removal:	11809.119	1865.737	9563.635	1609.698	203545.235	16842.058	185574.874	15455.254	707.379	336.022	550.107	33.236	15934.67	6283.131	13397.869	1184.402
Tree creation:	380.677	16.54	344.149	39.36	12.661	5.838	9.328	1.599	8.101	3.363	6.949	0.657	419.472	9.914	355.739	41.14
Tree removal:	304.796	55.441	265.271	28.4	11.666	5.424	8.84	1.544	7.89	4.235	6.748	0.67	245.813	149.232	221.517	14.554
24	**24 tasks, 240 files/directories**															
Directory creation:	8807.939	5211.853	8197.128	405.62	29335.926	2994.987	21658.814	4594.531	1344.696	614.03	1019.974	93.327	14068.944	3236.075	11883.691	1147.182
Directory stat:	20088.465	7547.389	18020.842	1146.207	500066.051	14122.632	458774.825	79277.524	2623.414	1098.655	1816.937	159.107	33103.126	9763.277	27980.036	2429.903
Directory removal:	8623.527	888.682	7820.434	959.882	94528.403	20356.17	48890.456	17306.982	1015.843	506.408	802.621	45.485	13763.098	3585.948	11891.247	1210.889
File creation:	8859.724	6050.24	8286.421	343.317	216433.662	12054.185	206548.082	14691.958	837.053	473.617	736.583	42.689	13941.319	106.979	11502.498	2530.201
File stat:	10438.461	88.158	8808.68	2274.821	531204.728	297468.369	504725.491	21120.062	1881.705	739.893	1648.422	123.886	19190.776	302.939	15438.322	3169.084
File removal:	10512.479	4207.417	9616.666	853.841	295546.964	239731.593	277451.261	7990.662	980.188	640.978	837.227	33.358	15644.307	9194.003	13556.428	975.313
Tree creation:	250.511	74.504	220.214	18.996	6.973	3.32	5.52	0.824	5.495	3.167	4.77	0.389	274.586	76.324	237.562	19.223
Tree removal:	204.75	63.557	175.631	14.984	6.634	0.13	5.217	0.81	5.299	1.871	4.528	0.444	161.133	98.395	146.525	9.788
32	**32 tasks, 320 files/directories**															
Directory creation:	8463.83	3174.662	7831.651	464.842	32405.652	657.437	21712.822	6029.258	1618.45	748.2	1213.25	105.464	13729.169	4495.397	11933.936	1064.549
Directory stat:	19293.581	11138.955	17420.656	1024.6	672433.507	430.037	599681.8	142621.998	3283.775	2061.935	2389.052	184.5	32612.739	18503.602	28364.336	2139.629
Directory removal:	8407.367	2178.502	7556.187	842.498	113599.431	2457.813	43685.097	17379.358	1181.784	137.101	1005.309	82.898	13776.518	6216.979	12351.775	877.017
File creation:	8585.099	2425.452	8049.129	573.005	287773.859	76959.706	269634.582	20551.071	1091.663	547.374	961.231	59.477	13895.469	148.208	11343.388	3144.227
File stat:	10254.395	147.963	8200.625	2737.566	715827.883	637311.149	676816.819	13571.793	2622.372	925.578	2203.256	174.985	19389.181	1396.118	16394.747	2451.572
File removal:	10209.622	2525.111	9389.472	1002.506	390281.268	25232.691	372529.221	31319.507	1206.277	932.317	1053.929	45.203	15727.964	7706.753	13838.639	1046.747
Tree creation:	180.47	122.519	160.863	9.649	4.558	0.104	2.827	0.727	3.446	0.766	2.618	0.355	197.985	49.885	176.924	14.023
Tree removal:	145.818	54.051	129.685	11.559	4.555	0.109	2.741	0.699	3.104	0.398	2.48	0.368	119.133	60.296	111.365	5.857
64	**64 tasks, 640 files/directories**															
Directory creation:	7698.973	2812.076	6612.61	426.122	44555.821	6875.274	30591.716	6444.968	3016.105	127.605	1873.868	744.648	13307	5927.193	11554.39	1083.855
Directory stat:	17768.475	10784.785	15612.935	808.175	1344867.014	39949.914	1199332.23	167695.046	6152.853	784.064	4431.093	989.302	30962.842	18494.423	27324.827	2313.383
Directory removal:	7843.325	2585.408	6953.962	660.406	123361.882	26899.766	64726.21	21280.814	2278.96	103.602	1556.282	623.341	13627.134	7116.415	12359.577	1030.202
File creation:	8563.572	6431.517	7953.035	278.786	549847.308	123271.242	527660.758	49401.85	2242.497	160.835	1633.038	495.859	13342.253	485.329	10824.746	3552.134
File stat:	9458.647	193.841	8046.208	1384.639	1385115.872	1185145.501	1309649.391	34491.312	5765.564	529.263	4099.624	906.54	18704.218	500.128	15477.208	4571.39
File removal:	9622.446	423.863	8842.509	812.583	791144.875	610080.582	727932.269	21442.355	2363.404	113.837	1640.111	648.337	15356.983	5879.548	13928.856	1146.981
Tree creation:	81.334	16.713	74.056	7.709	2.069	0.126	1.277	0.301	1.476	0.027	0.91	0.427	96.145	62.613	86.735	4.999
Tree removal:	69.72	10.192	62.831	6.925	2.071	0.109	1.238	0.28	1.41	0.059	0.793	0.41	59.513	27.467	54.642	2.914
96	**96 tasks, 960 files/directories**															
Directory creation:	7218.97	3060.933	6179.034	401.184	52901.33	10317.483	40200.818	6635.138	1222.056	703.187	1034.078	61.358	10470.055	6364.237	9367.236	748.573
Directory stat:	16906.764	11398.727	14516.051	879.785	1909213.77	27688.792	1701344.36	388612.55	2160.206	792.765	1783.401	173.494	24878.939	14875.177	22763.326	1389.845
Directory removal:	7490.567	2257.171	6635.661	805.126	245280.936	38301.96	105862.223	32315.716	912.484	356.084	807.35	77.44	12415.152	6359.935	11638.69	801.034
File creation:	8391.755	5225.082	7930.672	374.925	810167.372	194687.74	774145.961	66788.072	810.178	372.163	733.256	56.942	11998.045	395.552	10694.447	2471.664
File stat:	8991.286	116.761	6824.442	2108.548	2038750.299	1784019.424	1920685.91	41850.687	1893.408	735.803	1643.589	130.029	22002.185	332.856	14670.669	3513.127
File removal:	9499.114	485.728	8651.486	888.294	1181147.504	115132.583	1056012.246	130239.447	952.75	524.195	846.686	64.569	12583.423	8697.477	11827.33	755.663
Tree creation:	51.778	12.443	47.52	3.819	1.454	0.412	1.093	0.171	1.086	0.399	0.969	0.1	57.571	22.841	52.52	3.665
Tree removal:	44.577	9.472	40.923	3.98	1.417	0.517	1.054	0.143	1.001	0.294	0.881	0.1	36.304	20.849	33.787	1.796
128	**128 tasks, 1280 files/directories**															
Directory creation:	7079.881	2456.432	5943.45	639.803					1488.735	381.254	1320.155	166.328	18048.265	10935.415	16051.462	1430.648
Directory stat:	16764.067	11435.316	14435.663	907.864					2713.066	1046.463	2356.173	212.171	53637.209	23646.847	35524.175	3909.469
Directory removal:	7566.71	3019.095	6584.889	842.596					1197.745	507.732	1071.979	100.747	15436.773	2907.947	13185.282	1715.929
File creation:	8659.949	5652.885	8043.094	389.929					1064.505	411.211	962.598	98.125	14861.808	2005.411	11206.772	3906.405
File stat:	8708.736	155.876	6813.306	1919.071					2489.493	866.202	2180.225	215.131	22353.043	2134.002	17722.825	5411.82
File removal:	9558.799	5775.053	8863.525	442.062					1253.912	251.382	1101.047	153.338	19598.478	9653.747	17662.638	1499.477
Tree creation:	37.168	5.733	35.051	3.098					0.718	0.255	0.611	0.069	51.782	2.033	43.236	6.884
Tree removal:	34.16	4.481	30.196	3.837					0.705	0.249	0.604	0.075	31.289	2.244	27.825	4.626

3.3 IOR

```
$submit_pref ./$EXE -a POSIX -i 3 -d 32 -E -o "\$IORFILE" -s 1 -b 1G -t 1m
>>$stderr
```

Results:

Operation	Hexagon Max (MiB)	Min (MiB)	Mean (MiB)	Std Dev	Max (OPs)	Min (OPs)	Mean (OPs)	Std Dev	Mean (s)	Hexagon
2 nodes										
write	763.17	759.98	761.11	1.46	763.17	759.98	761.11	1.46	2.6908	
read	862.66	640.64	741.87	91.69	862.66	640.64	741.87	91.69	2.8021	
4										
write	763.17	759.98	761.11	1.46	763.17	759.98	761.11	1.46	2.6908	
read	862.66	640.64	741.87	91.69	862.66	640.64	741.87	91.69	2.8021	
8										
write	760.01	368.74	627.68	183.11	760.01	368.74	627.68	183.11	14.619	
read	15585	2417.52	7211.25	5942	15585	2417.52	7211.25	5942	2.0568	
12										
write	760.01	368.74	627.68	183.11	760.01	368.74	627.68	183.11	14.619	
read	15585	2417.52	7211.25	5942	15585	2417.52	7211.25	5942	2.0568	
16										
write	767.42	701.8	741.12	28.32	767.42	701.8	741.12	28.32	22.14	583.9
read	7080.59	5426.38	6342.61	686.99	7080.59	5426.38	6342.61	686.99	2.6153	493.22
24										
write	732.07	387.03	597.85	150.91	732.07	387.03	597.85	150.91	44.503	
read	9478.96	3207.79	6711.27	2612.53	9478.96	3207.79	6711.27	2612.53	4.518	
32										
write	739.72	698.68	714.49	18.03	739.72	698.68	714.49	18.03	45.891	
read	14837	5860.36	11825.8	4218.26	14837	5860.36	11825.8	4218.26	3.339	
64										
write	759.31	612.91	703.18	64.46	759.31	612.91	703.18	64.46	94.041	
read	148668	101417	132898	22261	148668	101417	132898	22260.5	0.5093	
96										
write	752.31	717.46	731.48	15.02	752.31	717.46	731.48	15.02	134.45	
read	221130	202537	213188	7828.41	221130	202537	213188	7828.41	0.4618	
128										
write	688.3	656.74	674.51	13.19	688.3	656.74	674.51	13.19	194.4	
read	298150	4408.03	190281	131996	298150	4408.03	190281	131996	10.221	

Operation	Titan Max (MiB)	Min (MiB)	Mean (MiB)	Std Dev	Max (OPs)	Min (OPs)	Mean (OPs)	Std Dev	Mean (s)
2 nodes									
write	859.92	556.9	726.44	126.31	859.92	556.9	726.44	126.31	2.915
read	518.89	345.63	454.57	77.45	518.89	345.63	454.57	77.45	4.658
4									
write	747.73	488.99	643.39	111.41	747.73	488.99	643.39	111.41	6.587
read	638.2	406.5	524.5	94.64	638.2	406.5	524.5	94.64	8.08
8									
write	1005.25	911.69	952.09	39.24	1005.3	911.69	952.09	39.24	8.619
read	882.98	701.03	797.52	74.69	882.98	701.03	797.52	74.69	10.36
12									
write	1403.95	1286.5	1346.84	48.03	1404	1286.5	1347	48.03	9.135
read	456.83	426.06	444.03	13.08	456.83	426.06	444.03	13.08	27.7
16									
write	1527.32	1420.8	1489.19	48.48	1527.3	1420.8	1489	48.48	11.01
read	676	667.66	671.22	3.51	676	667.66	671.22	3.51	24.41
24									
write	1000	927.14	964.99	29.81	1000	927.14	964.99	29.81	25.49
read	944.02	819.76	895.28	54.15	944.02	819.76	895.28	54.15	27.56
32									
write	1228.22	1116.9	1179.2	46.4	1228.2	1116.9	1179.2	46.4	27.83
read	1162.24	1066.1	1122.29	40.88	1162.2	1066.1	1122	40.88	29.24
64									
write	1196.85	1168.9	1179.14	12.58	1196.9	1168.9	1179	12.58	55.59
read	1693.52	1367.4	1485.44	147.58	1693.5	1367.4	1485	147.58	44.53
96									
write	1398.1	1335.7	1364.19	25.76	1398.1	1335.7	1364	25.76	72.09
read	1741.65	1672.2	1702.14	29.16	1741.7	1672.2	1702	29.16	57.77
128									
write	1394.69	1342.6	1371.8	21.75	1394.7	1342.6	1371.8	21.75	95.57
read	1747.62	1674.4	1705.92	30.75	1747.6	1674.4	1706	30.75	76.86

Operation	Stallo Max (MiB)	Min (MiB)	Mean (MiB)	Std Dev	Max (OPs)	Min (OPs)	Mean (OPs)	Std Dev	Mean (s)
2 nodes									
write	83.53	75	78.57	3.62	83.53	75	78.57	3.62	26.12
read	356.69	156.7	287	92.21	356.69	156.69	287	92.21	8.2346
4									
write	71.57	63.31	66.23	3.78	71.57	63.31	66.23	3.78	62.044
read	364.66	334.3	354.49	14.26	364.66	334.32	354.49	14.26	11.574
8									
write	69.56	65.34	68.11	1.96	69.56	65.34	68.11	1.96	120.37
read	1271.83	1123	1220.95	69.11	1271.83	1123.24	1221	69.11	6.732
12									
write	57.67	52.92	56	2.18	57.67	52.92	56	2.18	219.77
read	1848.22	1113	1402.67	319.78	1848.22	1112.82	1402.7	319.78	9.1817
16									
write	55.1	48.94	51.69	2.56	55.1	48.94	51.69	2.56	317.76
read	2209.17	1431	1840.24	318.82	2209.17	1431.31	1840.2	318.82	9.1923
24									
write	51.9	10.12	37.03	19.06	51.9	10.12	37.03	19.06	1134.2
read	3441.91	151.4	1623.92	1365.4	3441.91	151.39	1623.9	1365.4	62.901
32									
write	54.18	50.12	51.64	1.81	54.18	50.12	51.64	1.81	635.31
read	2284.09	1707	1932.83	251.71	2284.09	1707.19	1932.8	251.71	17.224
64									
write	55.11	49.61	53.12	2.49	55.11	49.61	53.12	2.49123	EXCEL
read	19375	11353	14368	3565	19375	11353.4	14368	3565	4.8169
96									
write	54.85	50.25	52.2	1.94	54.85	50.25	52.2	1.94188	EXCEL
read	127859	21418	57329	49875	127859	21418	57329	49875	3.229
128									
write	55.5	50.91	52.62	2.05	55.5	50.91	52.62	2.05249	EXCEL
read	50241.7	7408	29763	17537	50241.7	7408.29	29763	17537	8.1481

Operation	Njord Max (MiB)	Min (MiB)	Mean (MiB)	Std Dev	Max (OPs)	Min (OPs)	Mean (OPs)	Std Dev	Mean (s)
2 nodes									
write	978.6	799.02	859.5	84.25	978.6	799.02	859.5	84.25	2.4045
read	1377	791.37	1009	261.3	1377	791.37	1009	261.3	2.1521
4									
write	1578	1352	1491	99.57	1578	1352	1491	99.57	2.7602
read	2627	915.12	1507	792.3	2627	915.12	1507	792.3	3.4065
8									
write	2873	2052	2475	335.8	2873	2052	2475	335.8	3.3733
read	3256	1363.4	2215	784	3256	1363.4	2215	784	4.1892
12									
write	3022	2676	2833	143.1	3022	2676	2833	143.1	4.3478
read	4446	3705	4029	309.7	4446	3705	4029	309.7	3.0673
16									
write	3493	3435	3470	25.53	3493	3435	3470	25.53	4.7212
read	4694	4542	4621	62.27	4694	4542	4621	62.27	3.5464
24									
write	3762	3213	3447	231.4	3762	3213	3447	231.4	7.1604
read	4832	4653	4713	83.99	4832	4653	4713	83.99	5.2162
32									
write	3737	3458	3618	117.7	3737	3458	3618	117.7	9.0667
read	4635	4600	4615	14.46	4635	4600	4615	14.46	7.0999
64									
write	3880	1186	2871	1199	3880	1186	2871	1199	30.209
read	4888	4796	4852	39.93	4888	4796	4852	39.93	13.509
96									
write	3929	3810	3872	48.85	3929	3810	3872	48.85	25.395
read	4989	4816	4889	73.16	4989	4816	4889	73.16	20.114
128									
write									
read									