StorPool System Requirements

2019-05-17

1. Hardware components

This is a list of hardware StorPool is expected to work well on. In most cases any controller model which uses the same driver (in parenthesis) will work well. Components marked in blue are the ones we recommend based on experience. StorPool aims for wide hardware support. If your specific model of controller or drive is not on this list, please ask.

CPU:
- Nehalem generation (Xeon 5500) or newer Intel Xeon processor
- AMD Epyc 7000 series
- AMD family 10h (Opteron 6100) or newer AMD Opteron processor (*only for clients/initiators*)
- ARM
- IBM POWER8 or newer

Note: older supported CPUs and server platforms often have severe architectural I/O bottlenecks. We recommend Sandy Bridge (E3-12xx / E5-26xx) or newer CPUs.
Note: Usually one or more CPU cores are dedicated for the storage system. This significantly improves overall performance and performance consistency and avoids negative effects of CPU saturation. Please refer to the StorPool User Guide for more information on the recommended CPU setup.

RAM: 32GB (or more) ECC memory per server for normal operation of the server of which 8-20GB is typically used by StorPool. Exact memory requirement depends on the size of the drives and the amount of RAM used for caching in StorPool. The rule of thumb is to reserve 1GB RAM per 1TB of raw storage space. Contact StorPool Support for a detailed memory requirement assessment. Non-ECC memory is not supported.

HBAs and RAID controllers:
- Intel C200, C600, ICH10(82801JI) (ahci)
- LSI 2308 and 2116-based HBAs and integrated controllers (mpt2sas)
- LSI 3008, 3216, 3224-based HBAs and integrated controllers (mpt3sas)
- LSI MegaRAID controllers with 2108 and 2208 chipsets (megaraid_sas)
- LSI MegaRAID controllers with 3108 chipset (megaraid_sas)
- Dell PERC H730, H730 Mini, H730P Mini (megaraid_sas)
- HPE H240, P840ar (hpsa)
- Intel C600 SCU (isci)

- Storage controllers connecting HDDs must have power-loss protected write cache: BBU, CacheVault or Intel Optane P4801X NVMe drive
- Storage controllers connecting SSDs are used in JBOD mode.

Datacenter SATA SSDs:
- Samsung SV843, SM863, PM863, PM863a

- Intel DC S3500, S3510, S3520, S3610, S3700, S3710, S4500, S4600, S4510, S4610
- Micron M500DC, M510DC, X5100 ECO/PRO/MAX*
- HPE 1.92TB SATA 6G RI SFF SC DS SSD (868826-B21)
- Toshiba HK4 (Hawk-4) Series
- SanDisk (ex. Smart Storage Systems) CloudSpeed 1000E, Eco, Ascend
- HGST Ultrastar DC SS200
- minimum size is restricted to 960GB
- minimum number of 8 drives per cluster in same size are recommended

*Note: An SSD on this list that's **not** in blue means that we have performed validation tests on the drive and have not noticed any performance degradation or general issues which would prevent it from working with StorPool. The drives in blue also have been in production use for prolonged period of time under different workloads and have not shown performance degradation or other issues.*

* - Ensure that the drive is with the latest firmware recommended by the vendor, there are known good versions in the table below

NVMe SSDs:

- Samsung PM963
- Intel P3520, P3600, P3608, P3700, P4500, P4600, P4510, P4610
- Micron 9200 ECO, PRO
- Seagate Nytro 1551
- Toshiba PX04PMC
- Huawei ES300 V3

*Note: See the note on **Datacenter SATA SSDs** above.*

HDDs:

- Any enterprise-grade SAS or SATA hard drive
- We've had good experience with HGST Ultrastar, WD Re, Seagate and Toshiba HDDs.
- storage controllers connecting HDDs must have power-loss protected write cache: BBU, CacheVault or Intel Optane drive
- 512n, 512e and 4kn format drives are supported.
- maximum size is restricted to 4TB, to ensure sufficient performance. Using bigger size HDDs, up to 12TB, is only acceptable in large backup systems >100TB usable capacity and really cold data.

10 Gigabit Ethernet controllers:

- Mellanox ConnectX-3, ConnectX-3 Pro (mlx4_en)
- Mellanox ConnectX-4 Lx (mlx5_core)
- Intel 82599, X520, X540, X550 (ixgbe)
- Intel X710 and XL710 (i40e)
- Qlogic QLE3440-CU, QLE3442-CU (bnx2x)
- Broadcom BCM57810, BCM57840S (bnx2x)

25/40/50/56 Gigabit Ethernet controllers:

- Mellanox ConnectX-3, ConnectX-3 Pro (mlx4_en)
- Mellanox ConnectX-4 Lx, ConnectX-4, ConnectX-5 (mlx5_en)

- Intel XL710, XXV710 (i40e)

Ethernet Switches:
- StorPool works with any standards-compliant 10/25/40/50/56/100 GbE switch with jumbo frames and flow control
- Mellanox SX1012, SX1016, SN2010, SN2100B, SN2100, SN2410, SN2700
- Dell S4810, S6000

2. Network/fabric architectures

- Flat Ethernet networks
- Routed leaf-spine datacenter ethernet networks
- Ethernet with RDMA - RoCE, RoCEv2

Gigabit Ethernet networks are not supported.
For deployments utilizing the StorPool native protocol, 2 network ports per storage node are required.
For deployment utilizing iSCSI, 4 network ports per storage node are required.

3. Firmware versions

All network controllers, storage controllers, and SSDs must run the latest known-stable version of firmware available from the vendor. This is to prevent occurrence of known firmware bugs. Note that the latest stable might not always be the latest available version.

Known good versions for some devices:

Broadcom/Avago/LSI 2008,2108,2208,2308 and 2116-based controllers	19.00.00.00
Broadcom/Avago/LSI 3008	12.00.02.00
Broadcom/Avago/LSI 3108	4.650.00-6223
Broadcom/Avago/LSI 3ware 9690SA-4I	FH9X 4.10.00.027
HP HBA H240	Firmware Version: 4.52
HP Smart Array P840ar	Firmware Version: 4.52
Micron M500DC	0144
Micron M510DC	0013
Micron X5100 Eco 1.92T	D0MU037
Micron 5100 Pro 960GB	D0MU042
Micron 5100 Pro 1920GB	D0MU051

Intel S3500	D2012370

4. Example storage cluster configurations

The following hardware configurations cover the requirements outlined above and are common among StorPool customers.

4.1 Example hardware configuration for dedicated storage - All-SSD pool

- 3 copies on SSDs
- Distribute drives evenly across 3 servers
- 720 000 IOPS (random read) and <0.2 ms latency!
- Quoted usable space is before gains from snapshots/clones, thin provisioning and zeroes detection.
- 24x 3.84 TB SSDs - 27.9 TB usable (3x 10-bay nodes)
- other combinations and sizes are possible

Chassis	CSE-116TQ-R700CB - 10-bay 2.5"	3
Motherboard	X11SSL-CF	3
CPU	E3-1220v6	3
RAM	16 GB DDR4-2400 ECC UDIMM	6
NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	3
Boot drive	SATA SSD 240GB	3
Pool drive	3.84 TB datacenter SATA SSD, e.g. Intel S4510	24

4.2 Example hardware configuration for dedicated storage - All-NVMe pool

- 3 copies on NVMe SSDs
- Distribute drives evenly across 3 servers
- 2 400 000 IOPS (random read) and <0.15 ms latency!
- Quoted usable space is before gains from snapshots/clones, thin provisioning and zeroes detection.
- 30x 8 TB NVMe SSDs - 72.7 TB usable (3x 10-bay nodes)
- other combinations and sizes are possible

Barebone	SUPERMICRO AS-1113S-WN10RT	3
CPU	AMD EPYC Naples SP3 16C/32T 7351P 2.4G 64M	3
RAM	16 GB DDR4 ECC RDIMM	24
Boot drive	M.2 NVMe SSD, e.g. Intel P4101	3
NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	3
Pool drive	8 TB NVMe U.2 SSD e.g. Intel P4510	30

4.3. Example hardware configuration for dedicated storage - NVMe SSD-Hybrid pool

- 1 copy on NVMe SSDs and 2 spare copies on HDDs
- Total space on HDDs is two times the space on SSDs
- Distribute drives evenly across 3+ servers
- 360 000 IOPS (random read) and <0.15 ms latency! The performance of an all-flash array at a fraction of the cost.
- Quoted usable space is before gains from snapshots/clones, thin provisioning and zeroes detection.
- 6x 4TB NVMe SSDs, 24x 2TB HDDs - 21.8 TB usable (3x 12-bay nodes)
- other combinations and sizes are possible

Chassis	CSE-826BAC4-R920LPB - 12-bay 4x U.2	3
Motherboard	X11SPH-nCTPF (LSI3008, 2x10G SFP+, 2x OCuLink, 1x m.2)	3
CPU	Intel® Xeon® Silver 4110 - 8 cores @ 2.4 GHz	3
RAM	16 GB DDR4-2400 ECC RDIMM	6
Boot drive	M.2 NVMe SSD, e.g. Intel P4101	3
NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	3
Pool drive NVMe	4 TB NVMe U.2 SSD e.g. Intel P4510	6
Pool drive HDD	2TB Enterprise SATA HDD	24

4.4. Example hardware configuration for hyper-converged system - NVMe SSD-Hybrid pool

- 222 delivered dedicated vCPUs, excluding resources for Host OS and hyperconverged storage system (3 active nodes + 1 standby)
- 1152 delivered GB RAM (3 active nodes + 1 standby)
- 1 copy on NVMe SSDs and 2 spare copies on HDDs
- Total space on HDDs is two times the space on SSDs
- Distribute drives evenly across 3+ servers
- 480 000 IOPS (random read) and <0.15 ms latency! The performance of an all-flash array at a fraction of the cost.
- Quoted usable space is before gains from snapshots/clones, thin provisioning and zeroes detection.
- 8x 4TB NVMe SSDs, 32x 2TB HDDs - 29.1 TB usable (4x 12-bay nodes)
- other combinations and sizes are possible

Chassis	CSE-826BAC4-R920LPB - 12-bay 4x U.2	4
Motherboard	X11DPH-I	4
CPU	Xeon Gold 6138 - 20 cores @ 2.7 GHz	8

RAM	32 GB DDR4-2666 ECC RDIMM	48
NVMe contr.	AOC-SLG3-4E4T	4
NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	8
Boot drive	M.2 NVMe SSD, e.g. Intel P4101	4
Pool drive NVMe	4 TB NVMe U.2 SSD e.g. Intel P4510	8
Pool drive HDD	2TB Enterprise SATA HDD	32

4.5. Other configurations

Several other configurations have been omitted from this guide for brevity.

- SATA SSD-Hybrid
- HDD-Only
- Reverse Hybrid (2 copies on SSD and 1 copy on HDD)

Please contact StorPool support to obtain a full solution design for your use-case.

5. Operating systems

- CentOS 6, 7
- Debian 9
- Ubuntu 16.04 LTS, 18.04 LTS
- VMware ESXi 6.5 (host only, through iSCSI)
- If you are using another Linux distribution, e.g. RHEL, OEL, SuSE, just let us know. We can support StorPool on all Linux distributions with good build and packaging systems.
- x86_64 (amd64), ARM and Power architectures are supported

6. Cloud management systems and hypervisors

- OpenStack/KVM, OpenNebula/KVM, OnApp/KVM, CloudStack/KVM
- OpenStack/Hyper-V (through iSCSI)
- VMware Hypervisor with VMFS (through iSCSI)
- Proxmox, XenServer, XCP-NG (through iSCSI)
- Custom cloud management systems through StorPool API

7. Burn-in test

Hardware must pass a burn-in stress test before installing StorPool.

8. Stable server booting

When a server is restarted for whatever reason (e.g. power outage, administrator typed "reboot" at the console) it must get back up without human intervention.

9. Remote console and reboot capability

The server must have an IPMI module, including remote reboot and remote KVM capability.

10. Working kernel crash dump mechanism and debug symbols

If the Linux kernel crashes for whatever reason, we want to be able to investigate. We investigate by doing a post-mortem analysis of a kernel crash dump file from /var/crash/. StorPool requires working kernel crash dump on all servers participating in the StorPool cluster, as well as available kernel debug symbols for the running kernel.

Contacts:

For more information reach out to us.
support@storpool.com
+1 415 670 9320