

StorPool System Requirements

2018-11-19

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 family 10h (Opteron 6100) or newer AMD Opteron processor (*only for clients/initiators*)

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 (iscsi)

- Storage controllers connecting HDDs must have power-loss protected write cache: BBU or CacheVault.

- 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, S4510, S4610, [S4500](#), [S4600](#)
- 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 have 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
- 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 or CacheVault, unless HDDs are used only for third copy (after 2 safe copies on SSDs).
- 512n, 512e and 4kn format drives are supported.
- maximum size is restricted to 4TB, to gain sufficient performance.

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)

Infiniband controllers:

- Mellanox MCX353A-QCBT, MCX354A-QCBT (mlx4_ib)
- Intel (ex. Qlogic) QLE7340, QLE7342 (qib)

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
- Infiniband RDMA

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 node 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

Barebone	Supermicro SC113MFAC2-R606 - 1 RU, 8x 3.5" bays	1
Motherboard	Supermicro X11SSL-CF	1
CPU	Intel Xeon E3-1230V6	1
RAM	16GB DDR4 ECC UDIMM	2
Boot drive bracket	MCP-290-00036-0B	4
Boot drive	Intel S3520 150GB	1
NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	1
NIC alternative	Supermicro AOC-STG-I4S - Intel X710 - 4x 10GE SFP+	0
Pool drives	3.84 TB datacenter SATA SSD	8

9.3 TB usable per node

4.2 Example hardware configuration for dedicated storage - NVMe pools

Chassis	Supermicro SC813MFTQC-R407CB - 1 RU 4x 3.5" bays	1
Motherboard	Supermicro X11SRM-VF	1
CPU	Intel Xeon W-2133 - 6 cores	1
RAM	8GB DDR4-2666 ECC RDIMM	4
Backplane	1RU 4-slot NVMe backplane for SC813	1
Riser card	CSE-RR1U-E16	1
NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	1
Boot drive	Intel P3100 128 GB - internal M.2	1
Pool drives	7.68 TB NVMe U.2 SSD	4

9.3 TB usable per node.

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

Chassis	Supermicro CSE-826TQ-R500LPB - 2RU, 12x 3.5" bays	1
Motherboard	Supermicro X11SSH-F	1
CPU	Intel Xeon E3-1230 v6	1
RAM	16GB DDR4 ECC UDIMM	2
RAID controller	AOC-S3108L-H8IR-16DD - Broadcom/LSI 3108 controller	1
Cache protection	BTR-TFM8G-LSICVM02 - CacheVault kit	1

NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	1
NIC alternative	Supermicro AOC-STG-I4S - Intel X710 - 4x 10GE SFP+	0
Boot drive	Intel P3100 128 GB - internal M.2	1
Pool drives	1.92 TB datacenter SATA SSD	4
Pool drives	2TB server SAS/SATA HDD	8
	Build notes: install NIC and RAID controller in SLOT6 and SLOT7. Connect 8 3.5" slots of LSI 3108 controller. Connect 4 remaining slots to onboard Intel AHCI controller.	

7.0 TB usable per node.

4.4. Example of a hyper-converged hardware configuration - SSD-Hybrid pool

Server	Supermicro 6029BT-DNC0R - 2 servers in 2 RU, 12 drive bays	1
CPU	Intel Xeon Gold 6148 - 20 cores, 3.1 GHz all-cores	4
RAM	32GB DDR4-2600 ECC RDIMM	48
Boot drive	Intel P3100 128GB - internal M.2	2
GiE NIC	AOC-MGP-I2M-O	2
10/25G NIC	Mellanox MCX4121A-ACAT - 2x 10/25GbE SFP+	4
3.5" to 2.5" tray	MCP-220-00140-0B	4
Pool drives	7.68 TB NVMe U.2 SSD	4
Pool drives	2 TB server SATA HDD	8

7.0 TB usable per node. 2 copies on NVMe drives, 1 copy on HDDs.

4.5. Other configurations

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

- Hyper-converged All-SSD
- HDD-only for primary and DR
- HDD-only for backup and archive

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

5. Example drives pool configurations

5.1. StorPool All-SSD storage pools

- 3 copies on SSDs
- Distribute drives evenly across 3+ servers
- 100 000+ IOPS and 2000+ MB/s!
- Quoted usable space is before gains from snapshots/clones, thin provisioning and zeroes detection.
- [9x 3.84TB SSDs - 10.5 TB usable](#)

- 24x 3.84 TB SSDs - 27.9 TB usable (3x 8-bay nodes)
- other combinations and sizes are possible

5.2. StorPool SSD-Hybrid storage pools

- 2 copies on SSDs and 1 spare copy on HDDs
- Total space on SSDs is two times the space on HDDs
- Distribute drives evenly across 3+ servers
- 100 000+ IOPS and 2000+ MB/s! 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.
- 3x 7.68TB SSDs, 6x 2TB HDDs - 10.5 TB usable
- 12x 7.68TB SSDs, 24x 2TB HDDs - 41.9 TB usable
- other combinations and sizes are possible

5.3. StorPool SSD-Hybrid storage pools

- 1 copy on SSDs and 2 spare copies on HDDs
- Total space on HDDs is two times the space on SSDs
- Distribute drives evenly across 3+ servers
- 100 000+ IOPS and 2000+ MB/s! 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 1.92TB SSDs, 12x 2TB HDDs - 10.5 TB usable
- 12x 1.92TB SSDs, 24x 2TB HDDs - 20.9 TB usable (3x 12-bay nodes)
- other combinations and sizes are possible

5.4. StorPool triple-replicated HDD storage pool

- 3 copies of data on HDDs in different servers
- Distribute drives evenly across 3+ servers
- Exceptional performance for a HDD-based storage system
- Quoted usable space is before gains from snapshots/clones, thin provisioning and zeroes detection.
- 36x 4TB HDDs - 43.6 TB usable
- 108x 4TB HDDs - 130.9 TB usable
- other combinations and sizes are possible

6. Operating systems

- CentOS 6, 7
- Debian 8, 9
- Ubuntu 14.04 LTS, 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.

- Only the x86_64 (amd64) architecture is supported

7. Cloud management systems and hypervisors

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

8. Burn-in test

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

10. 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.

11. Remote console and reboot capability

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

12. 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