



Selecting Server Hardware for On-premises Deployment of FlashGrid Storage Fabric

Best Practices

rev. 2019-01-11

Overview

FlashGrid Storage Fabric software can run on a wide range of x86 servers from many server vendors – from 1-socket systems to 16-socket systems, rack-mounted, blades, or modular. This document provides recommendations for selecting server hardware for maximum performance and reliability while minimizing costs.

Identifying Optimal Server Hardware and Network Configuration

Step 1. Determine database compute requirements

In many cases DBAs have clear requirements for the compute resources they need in a particular database or cluster deployment. The following questions will help clarify the compute requirements.

- **Is Oracle ULA (Universal License Agreement) available or Oracle licenses are per CPU?**
With ULA, CPUs with the maximum number of cores recommended for best performance.
- **Is Oracle database licensed as Enterprise Edition or Standard Edition 2?**
With Standard Edition, CPUs with higher frequency recommended for maximum single-threaded performance.
- **If Enterprise Edition and no ULA, is the number of Oracle licenses fixed for certain amount of CPU cores?**
- **Are there specific preferences for the total number of CPU cores or sockets?**
- **Are there specific preferences for the number of database nodes?**
- **What type of database(s) will be used – RAC, RAC One Node, or single-instance?**
RAC databases may benefit from performance scaling with larger number of database nodes.
RAC One Node or single-instance databases typically do not require more than two or three database nodes.

Step 2. Determine database storage requirements

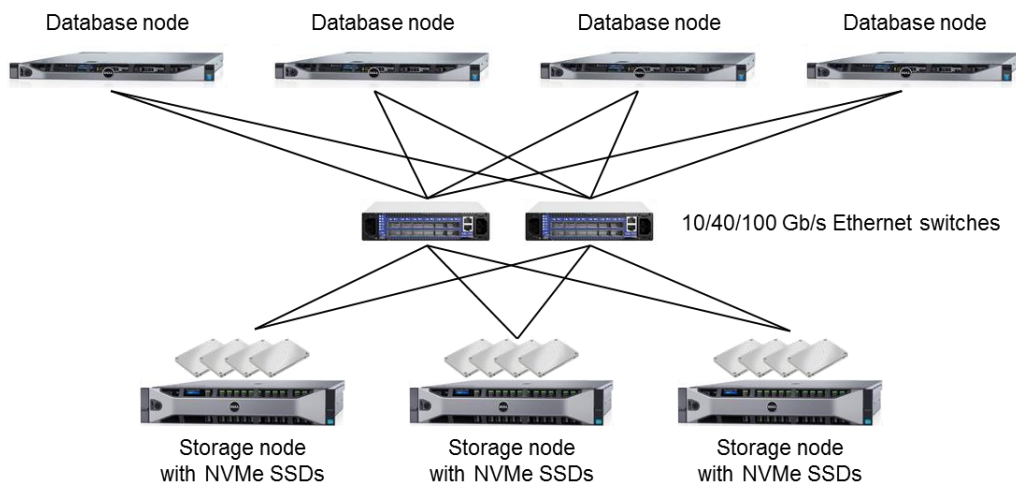
- **What is the required total usable storage capacity in the cluster – in the beginning and in 5 years?**
The total capacity should include DATA + FRA + non-database files on ACFS.
- **Is there a preference for 2-way or 3-way mirroring of the data?**
The total usable capacity must be multiplied by the factor of 2 or 3 respectively to get the total raw capacity.

Step 3. Determine the optimal way to place SSDs in servers

Does the customer have a preferred server vendor or server model? Use the table below to determine maximum capacities configurable in some popular servers.

- Place SSDs inside database nodes if total capacity is sufficient
- Add storage-only nodes for extra capacity (40 Gb/s or faster Ethernet is recommended in this case)
- Use separate database and storage nodes if large capacity is required with small number of database nodes
- In database nodes, hot-plug 2.5" SSDs are strongly recommended to avoid the need to shut down a database node when adding or replacing an SSD.

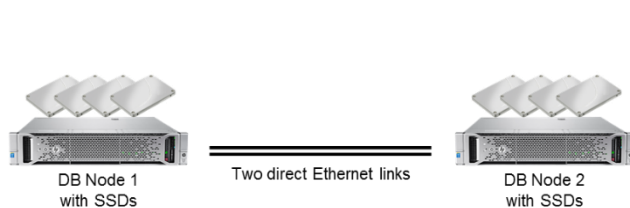
- Enterprise-class NVMe SSDs are recommended for best performance and reliability. SAS SSDs can be used if NVMe SSDs are unavailable or provide insufficient capacity. SATA SSDs should be avoided due to higher risk of reliability problems.
- It is recommended to install one extra SSD per disk group in each node. In case one SSD fails, this will allow restoring full redundancy of data even before the failed SSD is replaced.
- It is better to use larger capacity SSDs and keep some slots available for adding more SSDs in future. More SSDs (of smaller capacity) can provide higher performance, but in many cases it is more important to have the flexibility of adding more capacity when needed.
- Within one disk group, SSDs of the same type and capacity must be used. It is possible to add larger capacity SSDs to a disk group, but their capacity will be reduced to the smallest SSD capacity in the disk group.
- It is okay to have SSDs of different capacities or types in different disk groups.



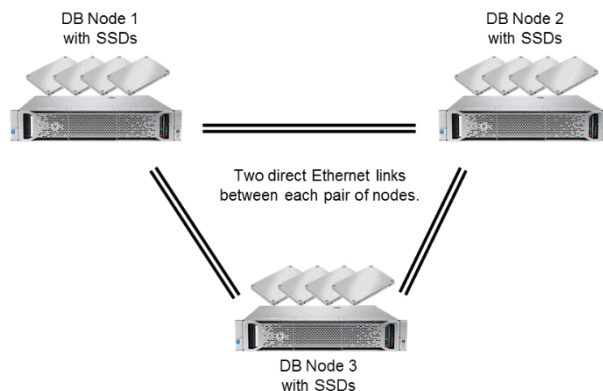
Example of a cluster configuration with separate database and storage nodes

Step 4. Determine the optimal network configuration

- Ethernet 10/40/100 Gb/s is recommended in most cases.
- Higher speed (40+ Gb/s) network is recommended when separate storage servers are used.
- In 2-node or 3-node hyper-converged configurations the storage network can be implemented as back-to-back links between the servers without network switches.
- At least two network links per server must be used for network redundancy.
- If using network switches then at least two switches must be used for the switch redundancy.
- If using RAC databases then it is recommended to use separate network links for FlashGrid Storage Fabric and for Oracle private network.
- If using RAC One Node or single-instance databases then the same network links can be shared between FlashGrid and the Oracle private network.
- Back-to-back links must not be used for Oracle private network.



Example of a 2-node cluster with switchless storage network



Example of a 3-node cluster with switchless storage network

Popular Server Models with Hot-plug 2.5" NVMe SSDs

Server model	2.5" hot-plug NVMe SSDs			Add-in PCIe card NVMe SSDs		Max total NVMe flash capacity per server
	# slots	Max capacity per SSD	Max capacity per server with 2.5" NVMe SSDs	# PCIe slots available for NVMe SSDs	Max flash capacity per server with 6.4TB add-in card SSDs	
Oracle Server X6-2L	9	3.2 TB	28.8 TB	5	32 TB	60.8 TB
Oracle Server X6-2	4	3.2 TB	12.8 TB	3	19.2 TB	32 TB
Dell PowerEdge R730xd	4	3.2 TB	12.8 TB	5	32 TB	44.8 TB
Dell PowerEdge R930	8	3.2 TB	25.6 TB	9	57.6 TB	83.2 TB
Dell PowerEdge R630	4	3.2 TB	12.8 TB	2	12.8 TB	25.6 TB
HPE ProLiant DL380 Gen9	6	2 TB	12 TB	5	32 TB	44 TB
HPE ProLiant DL560 Gen9	6	2 TB	12 TB	6	38.4 TB	50.4 TB
HPE ProLiant DL580 Gen9	5	2 TB	10 TB	8	51.2 TB	61.2 TB
Supermicro 1028U-TN10RT+	10	3.2 TB	32 TB	2	12.8 TB	44.8 TB
Supermicro 2028U-TN24R4T+	24	3.2 TB	76.8 TB	2	12.8 TB	89.6 TB
Supermicro 2028R-NR48N	48	3.2 TB	153.6 TB	2	12.8 TB	166.4 TB

Other server vendors and models can also be supported. Please contact FlashGrid to verify your preferred server configuration.

Other Server Hardware Configuration Recommendations

The following recommendations are not specific to FlashGrid Storage Fabric. However, they are important to maximize the server reliability.

- Boot drive:
 - Option 1:
 - Two SAS SSDs in RAID-1 mirror
 - Hardware RAID controller with battery or capacitor protection of write-back cache
 - Option 2:
 - A single SAS or SATA SSD
 - Disk-on-module SSD possible in some servers
 - Use of software RAID or onboard RAID without write-back cache protection is NOT recommended
 - Use of HDDs not recommended because of higher failure rates
- Redundant power supplies recommended.
- Redundant fans recommended.
- Remote access to system KVM recommended (optional feature in some systems)
- Support plan from the server vendor that includes:
 - advanced replacement of parts
 - onsite repair
 - retention of failed SSDs (if used for sensitive data)

Contact FlashGrid

For help with selecting or verifying server configuration for FlashGrid Storage Fabric deployment please contact FlashGrid by email info@flashgrid.io or phone +1-650-641-2421

Copyright © 2016-2019 FlashGrid Inc. All rights reserved.

This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document.

FlashGrid is a trademark of FlashGrid Inc. Red Hat is a registered trademark of Red Hat Inc. Oracle is a registered trademark of Oracle Corporation. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. Other names may be trademarks of their respective owners.