新え戻科学计算中心 Smarter technology for all Yang Yuanqing Scientific Computing Center

Lenovo Update SSUG 2022 London

Simon Thompson | 30th June 2022

2022 Lenovo. All rights reserved.

Lenovo Solutions for IBM Spectrum Scale™

Lenovo DE Series

- Spectrum Scale DAE/DME
- midrange SAN block storage
- Focus: midrange solutions



Lenovo DSS-G 2##

- Spectrum Scale DAE/DME
- direct-attached block storage
- Focus: high end solutions

Lenovo DSS-G 100

- Spectrum Scale ECE
- server-based block storage
- Focus: NVMe solutions





Lenovo DSS-G 4.1 Configurations

There is only one model number for ECE: DSS-G100.

But a Lenovo ECE solution will consist of multiple ECE servers where each one will be configured with at least 6 but not more than 32 servers. Storage solutions can be a hybrid mix of DSS-G and ECE Here are 3 examples.



Footnotes: 1. Fits in 48U rack.

2. Requires 2 racks.

Licensing IBM Spectrum Scale for DSS

IBM Spectrum Scale[™] for DSS-G## is licensed based on the storage.* Specifically you require one license per hard disk drive / flash module. IBM Spectrum Scale[™] for DSS-G100 (ECE) is licensed per usable TiB

	Data Access Edition	Data Management Edition	Erasure Code Edition
Multi protocol scalable file service with simultaneous access to a common set of data	٠	•	•
Facilitate data access with a global namespace, massively scalable file system, quotas and snapshots, data integrity & availability and filesets	•	•	•
Simplify management with GUI	•	•	●
Improved efficiency with QoS and Compression	٠	٠	•
Create optimized tiered storage pools based on performance, locality, or cost	•	•	•
Simplify data management with Information Lifecycle Management (ILM) tools that include policy based data placement and migration	•	•	•
Enable worldwide data access using AFM asynchronous replication	•	•	•
Asynchronous multi site Disaster Recovery		٠	•
Hybrid cloud (TCT)		•	●
Protect data with native encryption and secure erase, NIST compliant and FIPS certified.		•	•
File audit logging		٠	•
Watch folder		•	•
Erasure coding	DSS Only	DSS Only	•
Network dispersed erasure coding			•

Are any additional licenses (client, server, ...) needed for Spectrum Scale for Distributed Storage Server licensed per drive / flash module / capacity?

No, with this new model, solely the storage itself is licensed.

What license option do I have for non-DSS storage in the same Cluster, e.g. separated Metadata on traditional controller based Storage?

 For non-DSS storage there is the choice between the old socket-based licenses (Standard Edition only) and new capacity based licensed (Data Mgmt Edition only) – a per TB license.

Can I mix traditional GPFS/Spectrum Scale storage licensed per socket and new Spectrum Scale storage licensed per drive?

 Yes, however the drive-based license is only available with DSS-G. As long as a server/client accesses (cross-cluster/remote or locally) storage that is licensed per socket, it will also require a socket based client/server license.

Can I mix Standard Edition and Data Management Edition within a cluster?

 No, due to the additional features that is not possible. If a part of a cluster is Data Mgmt Edition, the whole Cluster needs to be Data Mgmt Edition.

Are the drive-based Spectrum Scale for DSS licenses transferrable?

 Existing Lenovo licenses can only be transferred from the existing Lenovo storage solution that is being decommissioned and re-used on its equivalent future or replacement Lenovo storage solution.

Where can I get more information?

 For more FAQ please have a look at <u>https://www.ibm.com/support/knowledgecenter/en/STXKQY/gpfsclustersf</u> aq.html

* Capacity (per usable TiB) licenses for Data Management Edition available as Special Bid.

Lenovo 2022 Lenovo. All rights reserved.

Mixing Licensing

Data Access Edition + Data Management Edition = NOT Permitted in same cluster Permitted when using multi-cluster

Data Access Edition + Erasure Code Edition

= Data Access Edition Features

Data Management Edition + Erasure Code Edition

=

Data Management Edition Features

Lenovo DSS-G 2XX 4.1 configuration capacities

Configuration	SR630	SR650	D3284 drive	D1224 drive	Number of drives	Number of drives
	servers	servers	enclosures	enclosures	(min total capacity)	(max total capacity)
DSS G201	0	2	0	1	24x 2.5" (19TB)*	24x 2.5" (184TB)*
DSS G202	0	2	0	2	48x 2.5" (38TB)*	48x 2.5" (369TB)*
DSS G203	0	2	0	3	72x 2.5" (58TB)*	72x 2.5" (553TB)*
DSS G204	0	2	0	4	96x 2.5" (77TB)*	96x 2.5" (737TB)*
DSS G211	0	2	1	1	24x 2.5" + 82x 3.5" (19TB + 328TB)†	24x 2.5" + 82x 3.5" (184TB + 1476TB)†
DSS G212	0	2	1	2	48x 2.5" + 82x 3.5" (38TB + 328TB)†	48x 2.5" + 82x 3.5" (369TB + 1476TB)†
DSS G221	0	2	2	1	24x 2.5" + 166x 3.5" (19TB + 664TB)†	24x 2.5" + 166x 3.5" (184TB + 2988TB)†
DSS G222	0	2	2	2	48x 2.5" + 166x 3.5" (38TB + 664TB)†	48x 2.5" + 166x 3.5" (369TB + 2988TB)†
DSS G231	0	2	3	1	24x 2.5" + 250x 3.5" (19TB + 1000TB)†	24x 2.5" + 250x 3.5" (184TB + 4500TB)†
DSS G232	0	2	3	2	48x 2.5" + 250x 3.5" (38TB + 1000TB)†	48x 2.5" + 250x 3.5" (369TB + 4500TB)†
DSS G241	0	2	4	1	24x 2.5" + 334x 3.5" (19TB + 1336TB)†	24x 2.5" + 334x 3.5" (184TB + 6012TB)†
DSS G242	0	2	4	2	48x 2.5" + 334x 3.5" (38TB + 1336TB)†	48x 2.5" + 334x 3.5" (369TB + 6012TB)†
DSS G251	0	2	5	1	24x 2.5" + 418x 3.5" (19TB + 1672TB)†	24x 2.5" + 418x 3.5" (184TB + 7524TB)†
DSS G252	0	2	5	2	48x 2.5" + 418x 3.5" (38TB + 1672TB)†	48x 2.5" + 418x 3.5" (369TB + 7524TB)†
DSS G261	0	2	6	1	24x 2.5" + 502x 3.5" (19TB + 2008TB)†	24x 2.5" + 502x 3.5" (184TB + 9036TB)†
DSS G262	0	2	6	2	48x 2.5" + 502x 3.5" (38TB + 2008TB)†	48x 2.5" + 502x 3.5" (369TB + 9036TB)†
DSS G271	0	2	7	1	24x 2.5" + 586x 3.5" (19TB + 2344TB)†	24x 2.5" + 586x 3.5" (184TB + 10548TB)†
DSS G272	0	2	7	2	48x 2.5" + 586x 3.5" (38TB + 2344TB)†	48x 2.5" + 586x 3.5" (369TB + 10548TB)†
DSS G281	0	2	8	1	24x 2.5" + 670x 3.5" (19TB + 2680TB)†	24x 2.5" + 670x 3.5" (184TB + 12060TB)†
DSS G282	0	2	8	2	48x 2.5" + 670x 3.5" (38TB + 2680TB)†	48x 2.5" + 670x 3.5" (369TB + 12060TB)†
DSS G291	0	2	9	1	24x 2.5" + 754x 3.5" (19TB + 3016TB)†	24x 2.5" + 754x 3.5" (184TB + 13572TB)†
DSS G210	0	2	1	0	82x 3.5" (328TB)**	82x 3.5" (1476TB)**
DSS G220	0	2	2	0	166x 3.5" (664TB)**	166x 3.5" (2988TB)**
DSS G230	0	2	3	0	250x 3.5" (1000TB)**	250x 3.5" (4500TB)**
DSS G240	0	2	4	0	334x 3.5" (1336TB)**	334x 3.5" (6012TB)**
DSS G250	0	2	5	0	418x 3.5" (1672TB)**	418x 3.5" (7524TB)**
DSS G260	0	2	6	0	502x 3.5" (2008TB)**	502x 3.5" (9036TB)**
DSS G270	0	2	7	0	586x 3.5" (2344TB)**	586x 3.5" (10548TB)**
DSS G280	0	2	8	0	670x 3.5" (2680TB)**	670x 3.5" (12060TB)**
DSS G290	0	2	9	0	754x 3.5" (3016TB)**	754x 3.5" (13572TB)**
DSS G2A0	0	2	10	0	838x 3.5" (3352TB)**	838x 3.5" (15084TB)**

* Capacity is based on using 800GB (min) or 7.68TB (max) 2.5-inch SSDs.

** Capacity is based on using 4TB (min) or 18TB (max) 3.5-inch HDDs in all but 2 of the drive bays in the first drive enclosure; the remaining 2 bays must have 2x SSDs for Spectrum Scale internal use.

† These models are a hybrid configuration that combines HDDs and SSDs in one building block. The number of drives and capacities are given in terms of HDD and SSD count.

What is interesting here ...?



Lenovo 2022 Lenovo. All rights reserved.

Single node performance scaling



Lenovo 2022 Lenovo. All rights reserved.

Anyone see anything wrong here?



... verify all your clients are good!



What is the best solution?

Increased performance & capacity, good for streaming bandwidth

Hybrid offerings for mixed stream and small transactional workloads

ECE with spinning disk doesn't make sense vs DSS-G 2##



12x CPU sockets 12x IB ports 96 drives 4+2P (66%)



SR650 SR650

D3284

D3284

D3284

D3284

D3284

DSS-G250

4x CPU sockets 4x IB ports 82 drives 8+2P (80%)



ECE is network heavy

- GNR RAID track written across the network
- We recommend a full non blocking fabric between ECE nodes
- READ Amplification:
 - 2x (i.e. 100Gbps of client I/O requires 200Gbps of network read between ECE servers)
- WRITE Amplification:
 - 2.4x for 8+3p
 - 2.25x for 8+2P
 - 2.5x for 4+2P
 - 2.75 for 4+3P
- We have seen a customer site where performance was limited by their network "optimisation"



An artifact of our network topology ...



Lenovo 2022 Lenovo CONFIDENTIAL. All rights reserved.

Thanks to Florian Zillner!

