

Lenovo HPC PORTFOLIO



Lenovo™

Lenovo™ ThinkSystem



SERVER




STORAGE

NETWORKING



Lenovo™

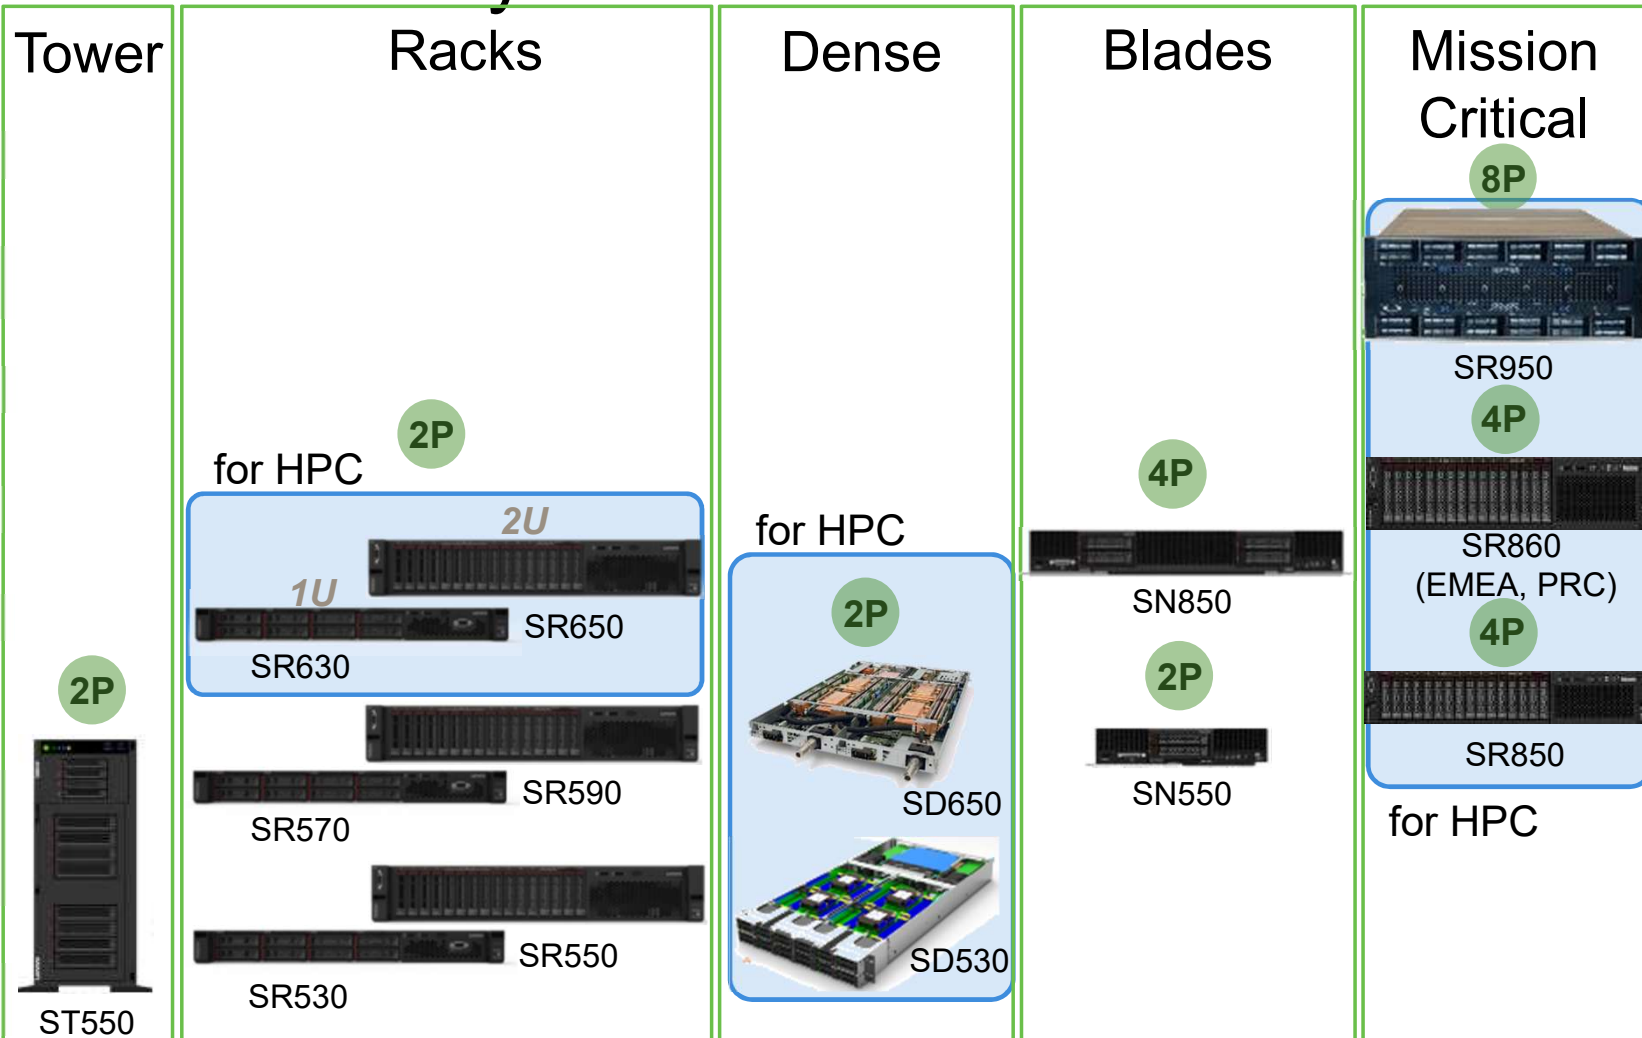

Scalable Infrastructure



Lenovo's HPC Server




+ The Lenovo ThinkSystem Server Portfolio



+ Intel Xeon Server processor history

Release date	Code	Processor	core/chip	TDP(W)	Spec FP	Spec_fp Rate
2006/6/26	Woodcrest	Intel Xeon 5160	2	80	17.7	45.5
2007/11/12	Harpertown	Intel Xeon x5460	4	120	25.4	79.6
2009/3/30	Nehalem	Intel Xeon x5570	4	95	43.8	202
2010/3/16	Westmere-EP	Intel Xeon x5690	6	130	63.7	273
2012/5/1	SandyBridge	Intel Xeon E5-2690	8	135	94.8	507
2014/1/9	IvyBridge	Intel Xeon E5-2697v2	12	130	104	696
2014/9/9	Haswell	Intel Xeon E5-2699v3	18	145	116	949
2015/3/9	Bradwell	Intel Xeon E5-2699v4	22	145	128	1160
2017/7/11	Skylake	Intel Xeon Platinum 8180	28	205	155	1820



Processor performance dependency

- Spec_fp Rate with 2 processor chips/node is growing up 40 times faster past 11 years (2006 – 2017).
- The number of cores on the chip increase 14 times.
- TDP watts increases 2.56 times bigger.
- Current maximum TDP is 205W. Knights Mill Xeon phi processor will be 305 W

Big issue:

For increasing performance will be **processor / server cooling technology.**

+ Moving forward into a decade of dense HPC

Lenovo continues development on both Air and Water cooling system

iDataPlex – Air / DWC

dx360 - Harpertown
dx360M2 - Nehalem
dx360M3 - Westmere
dx360M4 - Sandy Bridge

NeXtScale – Air / DWC

nx360M4 - Sandy Bridge/Ivy Bridge
nx360M5 - Haswell/Broadwell


ThinkSystem – Air

SD530 - Skylake/Icelake

ThinkSystem – DWC

SD650 - Skylake/Icelake

Air Cooling System






2011~

2013~

2017/18~

Direct Water Cooling System



+ Lenovo cooling technology keeps 0.5U/server

2006 HS21 Blade server

Air cooling

Woodcrest 2cores TDP=80W

Blade server E chassis 7U 14servers

(1server = 0.5U)



System

**Air
Cooling**
0.5U
/server

2017 Thinksystem SR650

2U 1server Air cooling

SkyLake Platinum 28cores TDP=205W

(1 server = 2U)




**Air
Cooling**
2U
/server

2018 Thinksystem SD650

1U 2servers DWC



SkyLake Platinum 28cores TDP=205W

(1server = 0.5U)



**Direct
Water
Cooling**
0.5U
/server

Heat sink



+ WCT remove heat by water

Example Node Power Map with 165W CPUs

NextScale Chassis w/ 12 Nodes

Chassis Power 6 kW-7kW capable

Maximum node power w/ standard memory

- ~ 450W w/165W CPUs
- ~540W w/ 205W CPUs

With 3D-Cross Point memory

- ~510W w/ 165W CPUs
- ~600W w/ 205W CPUs

Heat Recovery target with air cooled PSU is 90%+

Scalable water manifold to enable 1-6 chassis in a rack

Chassis water flow 6 – 9 liters/min,

Rack 36 – 54 liters/min for 6 Chassis (72 Nodes)

Direct to Water	Node Power		Estimates				
	Component Specifics (QSPEC, Model, rev, P/N, etc.)		Specification Power (W)	12V Power	Qty	Max. Power (W)	Weighted Power (W)
	CPU	Intel Skylake CPU (165W)	165	175.5	2	351.1	351.1
	Memory	64GB RDIMM (1DPC)	7	7.8	12	93.3	65.3
		256/512GB Apache Pass (2nd DIMM on channel on 12 channels)	18	20.0	0	0.0	0.0
	Chipset	Lewisburg PCH (LWB-2) 2x1Gbe+ 2x10GbE capable	13	14.4	1	14.4	10.1
	Planar	Emulex Pilot 4	4	4.4	1	4.4	3.1
		1GbE Mangement Port	2.2	2.4	1	2.4	1.2
	System component		Item Power (W)	12V Power	Qty	Max. Power (W)	Weighted Power (W)
	I/O	PCIe slot(15W) - General Purpose Dual port infiniband EDR	15	15.8	1	15.8	15.8
Storage TPM	M..2 Module HW RAID	2	2.0	0	2.2	1.6	
	Pluggable module	1	1.0	1	1.1	0.6	
Total						493.4	455.0
Conduction to Water loop							


- Water cooled PSU evaluation in plan for POC effort with Lenovo STIC team to increase heat recovery to 95%+.
- Rack and manifold thermal insulation for external walls could be evaluated.

+ 7. Heat Extraction/Performance Expectations

Compute Tray Configuration (dual compute nodes)	Single Rack Steady State Power (w/Linpack, Turbo ON/OFF)	90% Heat Removal		85% Heat Removal	
		Heat-to-water	Heat-to-air*	Heat-to-water	Heat-to-air*
2 server / tray with TDP=165W CPU, 16x16GB DDR4 Memory, 2x NIC	33.5kW	29.3kW	4.2kW	27.6kW	5.9kW
2 server / tray with TDP=145W CPU, 16x 16GB DDR4 Memory, 2x NIC	30.7kW	26.7kW	4.0kW	25.3kW	5.4kW
2 server / tray with TDP=135W CPU, 16x 16GB DDR4 Memory, 2x NIC	29.4kW	25.6kW	3.8kW	24.1kW	5.3kW
2 server / tray with TDP=120W CPU, 16x 16GB DDR4 Memory, 2x NIC	27.3kW	23.7kW	3.6kW	22.4kW	4.9kW

Assumptions: 36x compute trays (dual node), 72 servers + 6x switches / Rack

- Performance is dependent on many facility variables such as ambient temp, water temp, flow rate, etc... Based on testing, it is reasonable to assume 85-90% heat-to-water given the availability of typical environmental input parameters.



+ ThinkSystem SD530 (Air Cooling)

Feature	SD530
Processors	2 Intel “Purley” Generation processors per node <ul style="list-style-type: none">Up to 165W and 28 Cores
Form factor	1U Half wide tray node / 2U4N Chassis
Memory Slots Max Memory	<ul style="list-style-type: none">12x DDR4 (R/LR) 2667MHz DIMM (up to 64GB)
Storage	<ul style="list-style-type: none">6x SAS/SATA SF HS HDD/SSD or 2x NVMe, 2x M.2 SATA SSD
NIC	2x 10G SFP+ or GBase-T
PCIe	2x x8 or 1x x16 PCIe (HS) in Rear Shuttle 2x x16 internal for expansions
Power	HS/1+1 redundant 1100W/1600W/2000W Platinum
USB ports	1x USB 2.0
Cooling	<ul style="list-style-type: none">5x hot swap redundant fansSupport ASHRAE A2/3/4 (config limitation A3/4)
System MGMT / TPM	Stark Mgmt Module, 1 shared GbE port TPM, Pluggable TCM, Opt. KVM break-out
Expansions	GPU Expansions for 2 GPUs per node
Dimensions	846mm depth, front access w/ rear I/O



x16 PCIe Out

2x10 GbE Out


12 DDR4 DIMMs
+ 4 Apache Pass

2x SATA or PCIe M.2 BootSSD

2S Purley CPU max
~ 165W*

Hardware RAID (Optional)


4 or 6x 2.5 front hot swap drives



*in 4x drive configuration only

+ Front View



Node front view: Six 2.5-inch hot-swap HDD model




Node front view: Four 2.5-inch hot-swap HDD model

Node Storage bay options


- Six 2.5-inch hot-swap SAS/SATA HDD/SSD bays
- Four 2.5-inch hot-swap SAS/SATA HDD/SSD bays
- Four 2.5-inch hot-swap SAS/SATA HDD bays + two 2.5-inch hot-swap SAS/NVMe drive bays




+ The Shuttle – I/O



Eight low-profile PCIe 3.0 x8 cards



Four low-profile PCIe 3.0 x16 cards



+ GPU Tray

Node limitation

- Only support 2x2 Backplane (4 HDD config)
- Must be identical GPU cards in the same GPU tray

Chassis limitation


- Must have 2 x GPU tray in a chassis for air flow balance
- Only support 2000W PSU with GPU tray




GA Q4/2017

Supported Accelerators

NVIDIA Tesla P40 24GB (250 W)



NVIDIA Tesla P100 16GB (250 W)




NVIDIA Tesla M10 32GB (250 W)




Intel Xeon Phi 7240P 16GB, 1.3 GHz, 68 Core (275 W)



NVIDIA Tesla V100 16GB (300 W)



NVIDIA Tesla M60



+ SD530 Embedded Management Engine

XClarity Controller delivers modernized deployment and inter-operability for data centers who put a premium on efficiency and precision

All-new management engine

- Common in all ThinkSystem servers

Elegant User Experience

- Fresh, uncluttered GUI

Modernized inter-operability

- Built on Redfish-compliant APIs

Vastly Improved Performance

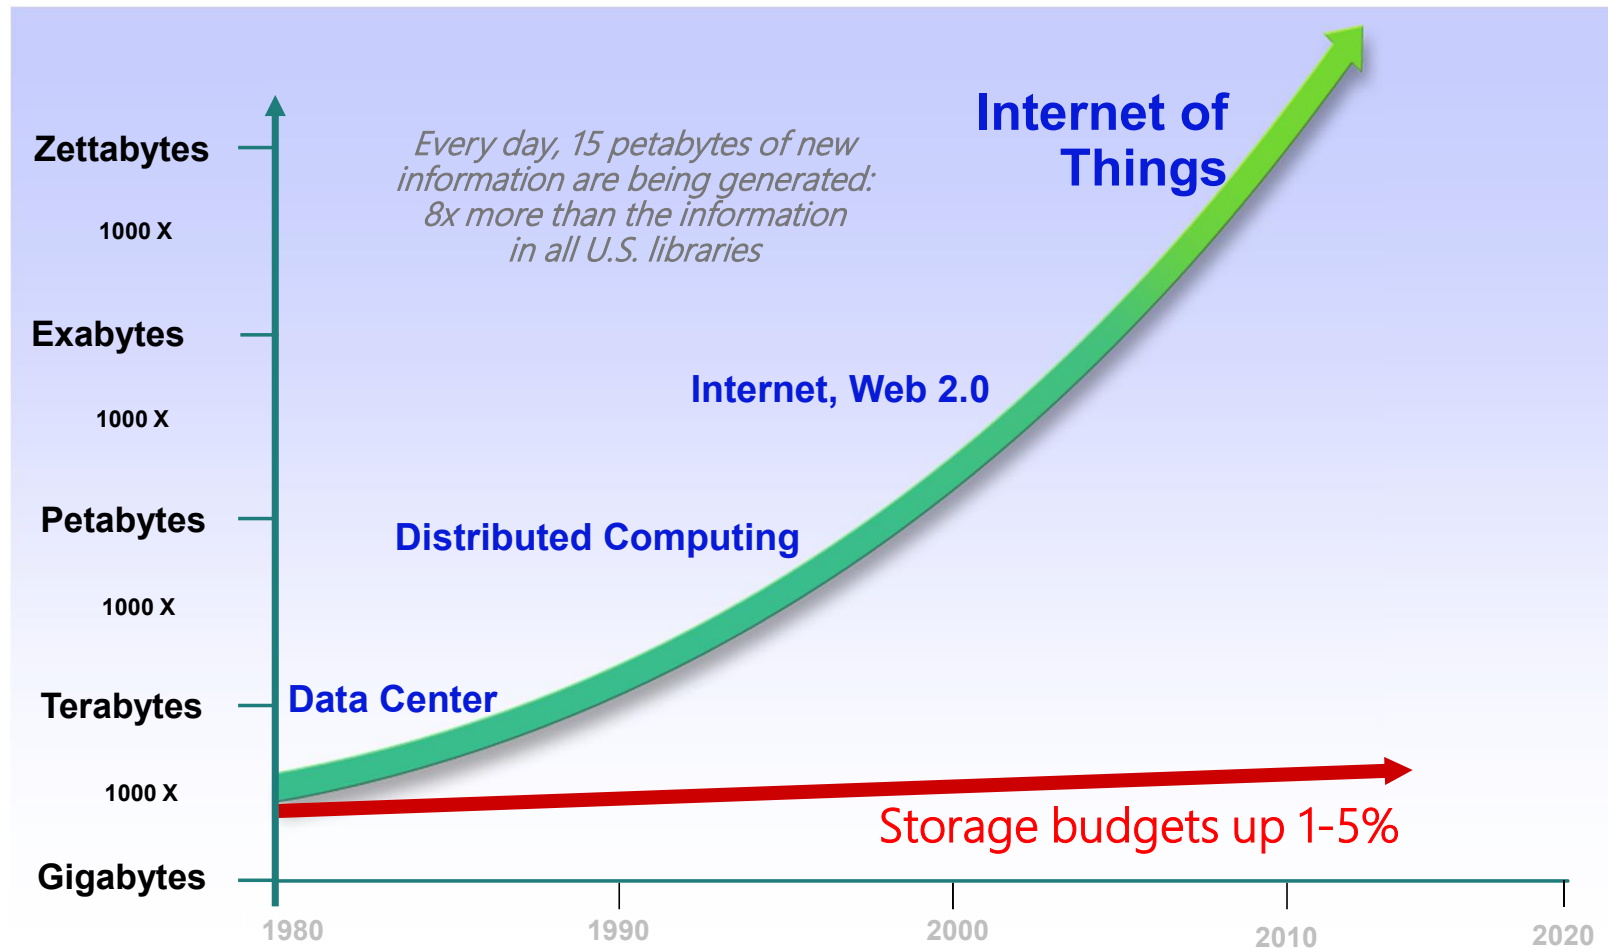
- 2x faster boot time to OS screen
- 6x faster firmware updates



Lenovo's HPC Storage & Solution




+ Storage Requirements Devouring Resources



+ Storage is Essential and We Have the Tools

Taking our off-the-shelf server and storage portfolio
marrying it with leading HPC Storage Software



Distributed Storage
Solution
for
IBM Spectrum Scale

DSS - G

Defined **Solution** especially
for large capacity, high
performance workloads in
HPC environments

Distributed Storage
Solution
for
SUSE Enterprise Storage

DSS - C

Defined **Solution** especially
for interaction with Lenovo
scale-out HANA solutions.*

Distributed Storage
Architecture
for
SUSE Enterprise Storage
/ Red Hat Ceph Storage

DSA - C

Tested **architecture** as
entry point and mid range
CEPH offering in HPC
environments.

Distributed Storage
Architecture
for
Intel Lustre EE

DSA - L

Tested **architecture** as
entry point and mid range
Lustre offering in HPC
environments.

2017 April Announcement

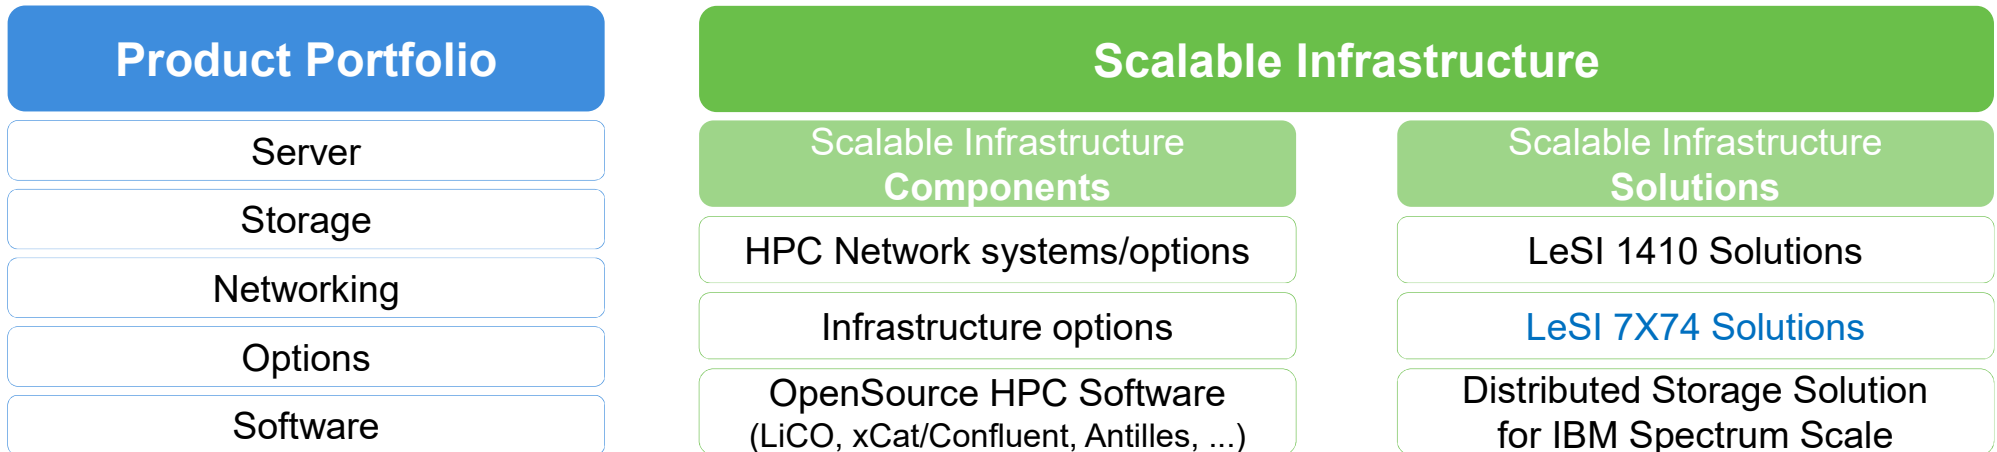
Future Aspirations

*formerly published as „ThinkStorage for SAP HANA TDI“

+ Lenovo Scalable Infrastructure (LeSI)

Lenovo Scalable Infrastructure (LeSI) is a framework for development, configuration, build, delivery and support of integrated data center solutions

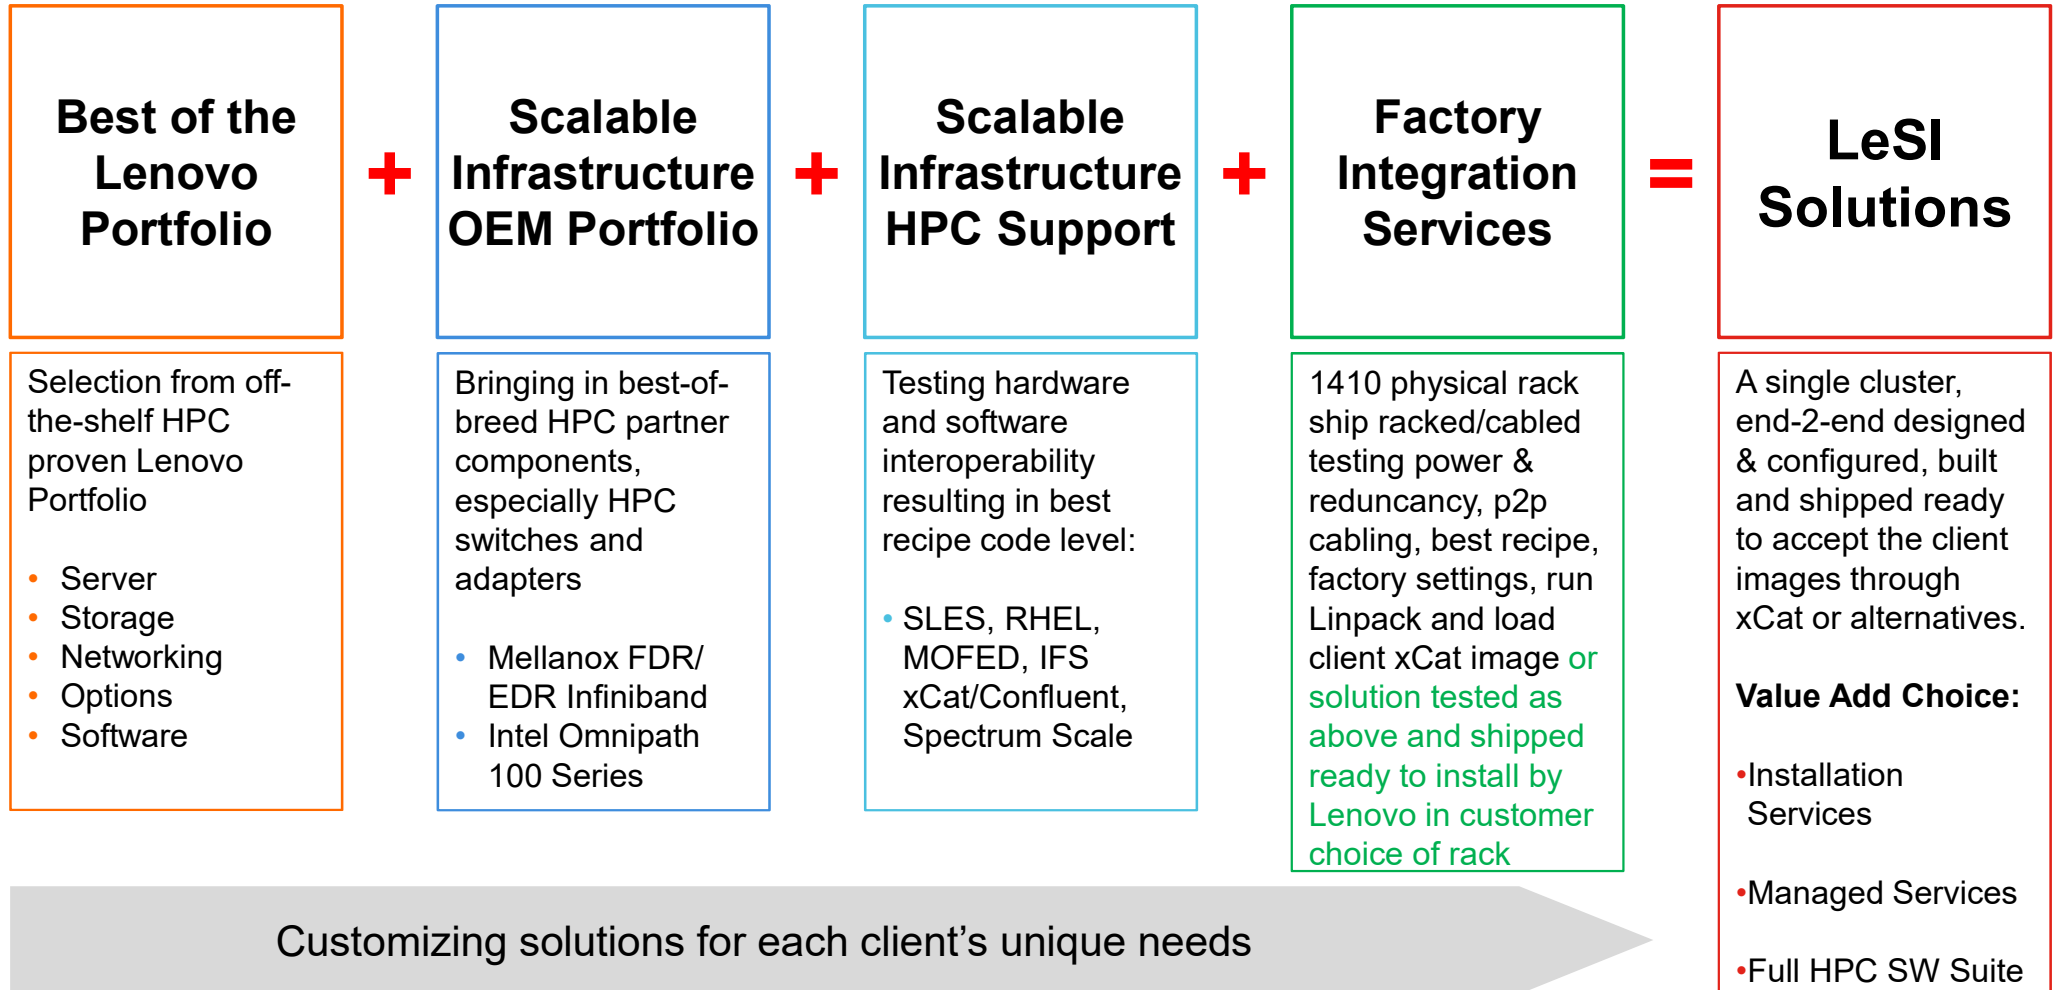
- Complete HPC data center portfolio with the best-of-breed partner technology
- Collaborate on OpenSource HPC software in true commitment to Openess
- End-to-end expert-designed, tested, integrated and supported HPC solutions



+ LeSI Solutions

From Complete Portfolio to Results Driven Designs

Effective
11/21/2017
for DSS-G



+ Lenovo Distributed Storage Solutions : **DSS-G** for IBM Spectrum Scale

Lenovo DSS-G

Our first organic HPC Storage Offering



Lenovo Distributed Storage Solution for IBM Spectrum Scale

- Designed for extreme use cases
- High storage density and I/O performance with superior availability, reliability and resiliency
- HPC, BigData and Cloud

DSS-G integrates many components, but has **No storage controller!**

Servers:



Storage Enclosures/Disks:

D1224 (2U24)
24 x 2.5" SAS/SSD



Disk Options: Many!

- GNR MD – (2) 12Gb 2.5" SAS, 3DWD: 400GB
- 15K, 12Gb 2.5" SAS: 300GB, 600GB
- 10K, 12Gb 2.5" SAS: 600GB, 900GB, 1.2TB, 1.8TB
- 7.2K, 12Gb 2.5" SAS: 1TB, 2TB
- 3DWD, 12Gb 2.5" SSD: 400GB, 800GB, 1.6TB, 800GB SED (FIPS)
- 1DWD, 12Gb 2.5" SED: 3.84TB

Network:

DSS-G will work with any LAN supported by LeSI (e.g., EDR, OPA, Ethernet). While adapters are included with the solution, switches must be purchased separately; this allows customers to use existing fabrics.



Disk Options:

- GNR MD – (2) 12Gb 2.5" SAS, 3DWD: 400GB
- Data – 12Gbit NL-SAS: 4TB, 6TB, 8TB, 10TB

Lenovo Distributed Storage Solution for IBM Spectrum Scale

- Acceptable short names
 - Lenovo DSS for Spectrum Scale
 - DSS-G (G= GPFS/Spectrum Scale)
- DSS-Gxyz
 - X = the number of DSS-G servers (always 2)
 - Y = the number of large storage enclosures (D3284)
 - Ex.: DSS-G220 = 2 servers & 2 D3284 enclosures
 - Z = the number of small storage enclosures (D1224)
 - Ex.: DSS-G204 = 2 servers & 4 D1224 enclosures

+ What is DSS-G?

DSS-G refers to the *Lenovo Distributed Storage Solution for Spectrum Scale™* (aka GPFS). It is a storage *solution* based on Spectrum Scale RAID (aka, GPFS Native RAID or GNR). DSS-G integrates many components as part of a *complete storage solution*.

DSS G202
SSD or SAS

DSS G206
SAS

DSS G240
NL-SAS

DSS-G Head node

2 x D1224

6 x D1224

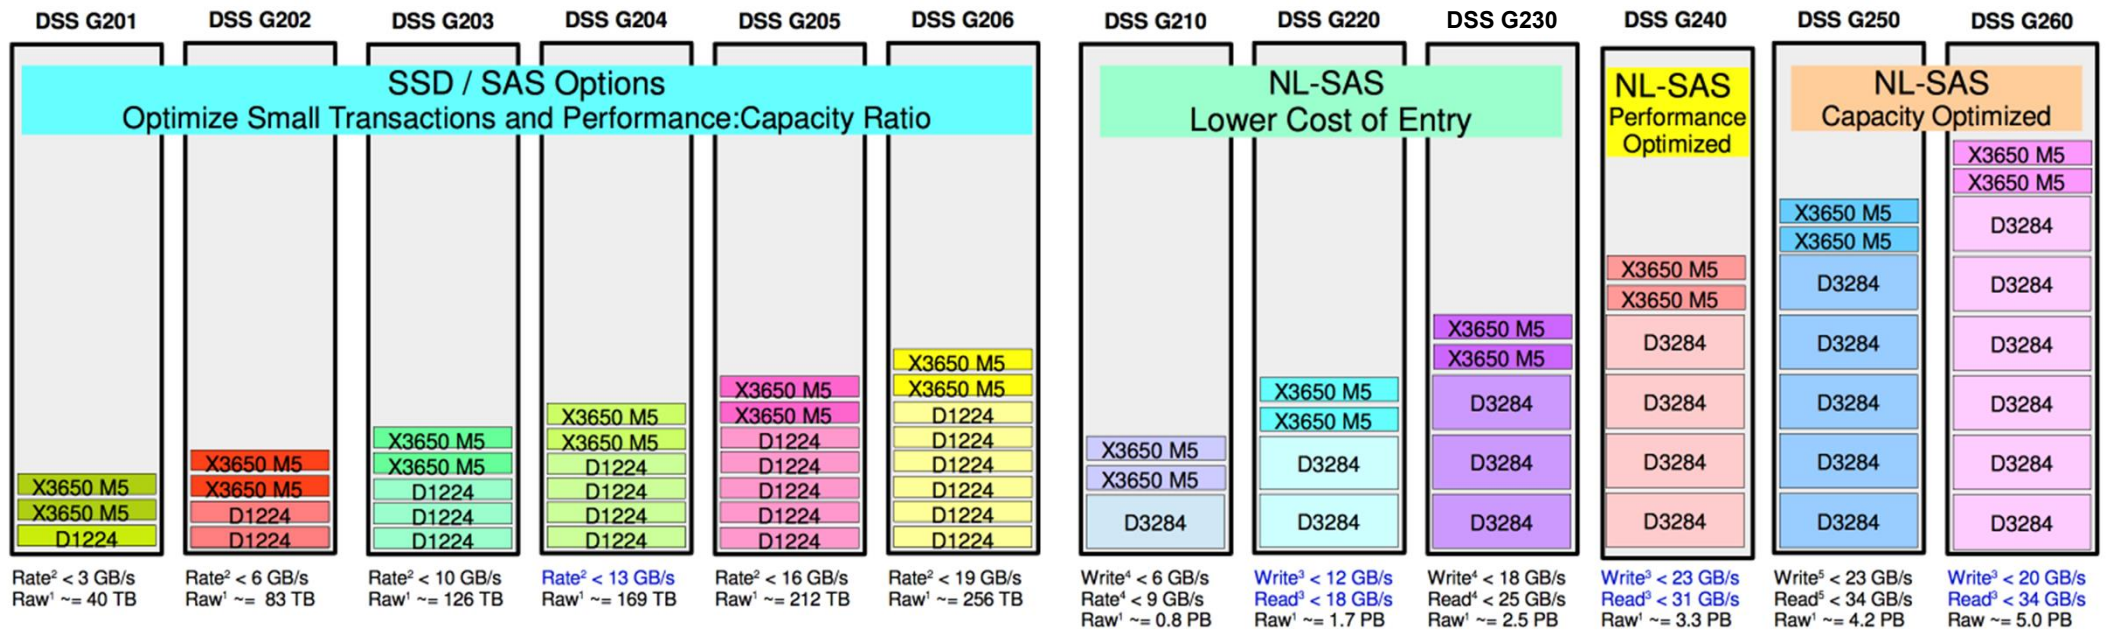
Lenovo Storage D1224 enclosures/disk
2U rack mount enclosure SAS direct-attached storage
Holds 24x 2.5-inch hot-swap drive bays
Selectable May capacity SAS HDD / SSD

DSS-G
JBOD Storage

Lenovo Storage D3284 Drive enclosures
5U rack mount enclosure with 12 Gbps SAS direct-attached storage
Holds 84x 3.5-inch hot-swap drive bays in two drawers.

+ DSS-G line up


The complete set of configurations supported in later releases this year.



Comments:

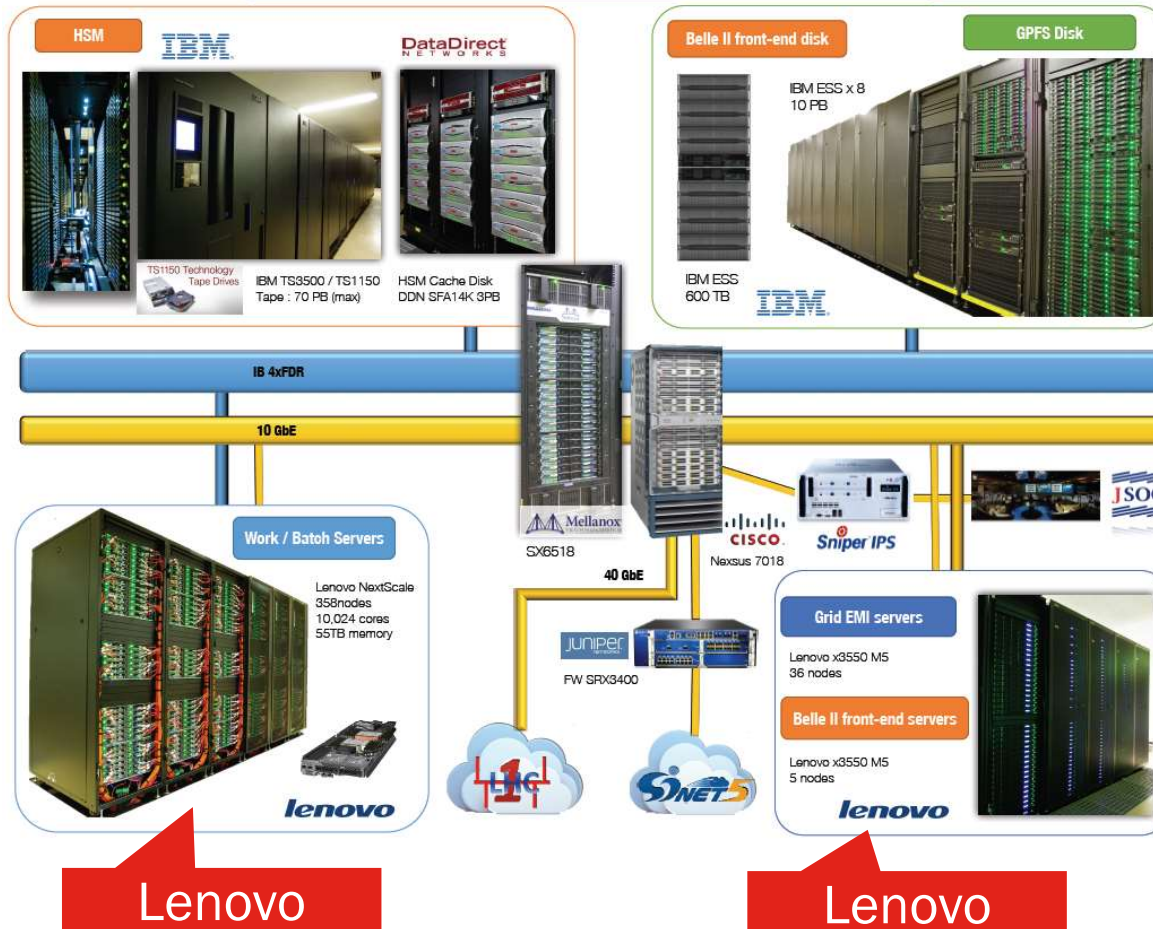
- Capacity for D1224 is based on 1.8TB, 2.5", 10KRPM SAS.
Capacity for D3284 is based on 10TB NL-SAS
- Performance is an **estimate** based on raw disk tests(n.b., no filesystem) on G204(1.2TB, 10KRPM, 2.5" SAS)
- Performance based on **actual** G220 and G240 benchmarks
- Performance is an **estimate** based G220 and G240 benchmarks
- Performance is an **estimate** based G220 and G260 benchmarks

Lenovo's Event information



+ Please join **KEK Facility Tours**

KEKCC 2016



Lenovo provided “KEK Central Computer System : **Work/Computing Server**”
Lenovo NeXtScale nx360 M5 10,000 cores

The Data Analysis System provides high-performance computing servers and a large-capacity storage system for accumulation and analysis of large amounts of experimental data.

The computing servers, which comprise approximately 10,000 CPU cores, can be used for both interactive and batch analyses.

Ref. <https://www.kek.jp/en/PublicRelations/DigitalLibrary/2017/01/17/Eng.pdf>

Please join KEK Facility Tours.

Friday after lunch,
October 20th 13:00 - 14:30

+ Please Visit Lenovo at **Super Computing 2017**

A dark blue rectangular graphic with a pattern of yellow and green dots at the bottom. It contains text about visiting the Lenovo Data Center Group at SC17. Logos for Intel, ThinkAgile, ThinkSystem, and a vertical Lenovo logo are also present.

Solve humanity's greatest challenges
with a future-defined data center.

Visit Lenovo Data Center Group at SC17

November 12-17, 2017
Denver, CO

Booth #1353

intel

ThinkAgile
ThinkSystem

Lenovo

REMIND ME TO VISIT LENOVO ►

<http://eblasts.cleandesign.com/lenovo/2017-10-06-super-computing/Visit-Lenovo-Booth-1353.ics>

BOOK A BRIEFING ►

<https://esg.formstack.com/forms/sc17>

Lenovo helps our customers solve the greatest challenges of today's world with our HPC solutions.

As the *fastest-growing supercomputing and HPC company and research*, Lenovo is uniquely positioned to help you accomplish your research and development goals.

At SC17, visit Lenovo at **Booth #1353** to see exclusive demos, meet with Lenovo Data Center experts, and envision your future-defined data center.

thanks.

Different is better

Lenovo™