



# Beyond Storage 2021

Data & AI에 특화된 스토리지 솔루션

2021. May 13  
한국IBM 이상호 차장



IBM Spectrum Scale



**Beyond Storage 2021**

Data & AI에 특화된 스토리지 솔루션



업계 최강의 병렬파일시스템 어플라이언스  
**Elastic Storage System 3200**

**Beyond Storage 2021**

Data & AI에 특화된 스토리지 솔루션



*Solution*

**+20years**

Evolution & Verification

*AI & Big data*

**#1 in HPC**

NVIDIA POD RA

GDS(GPU Direct Storage) Support

*Professional*

**+170PB**

Various Use Cases  
Professional Support

~ 2021년 국내사례

3 **Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, 1,572,480  
 NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM /  
 NVIDIA / Mellanox  
 DOE/NNSA/LLNL  
 United States



Home » DOE/NNSA/LLNL » Sierra - IBM Power System AC922, IBM POWER9 22C 3.1GH...

## SIERRA - IBM POWER SYSTEM AC922, IBM POWER9 22C 3.1GHZ, NVIDIA VOLTA GV100, DUAL-RAIL MELLANOX EDR INFINIBAND

Site:	DOE/NNSA/LLNL
System URL:	<a href="https://hpc.llnl.gov/hardware/platforms/sierra">https://hpc.llnl.gov/hardware/platforms/sierra</a>
Manufacturer:	IBM / NVIDIA / Mellanox
Cores:	1,572,480
Memory:	1,382,400 GB
Processor:	IBM POWER9 22C 3.1GHz
Interconnect:	Dual-rail Mellanox EDR Infiniband
<b>Performance</b>	
Linpack Performance (Rmax)	94,640 TFlop/s
Theoretical Peak (Rpeak)	125,712 TFlop/s
Nmax	11,902,464
HPCG [TFlop/s]	1,795.67

**Tweets** by @top500supercomp

**TOP500**  
 @top500supercomp  
 Don't forget to vote ...  
 Or it becomes "The Rise of Vectors"  
<https://twitter.com/top500supercomp/status/1329560336614125568>

Nov 21, 2020

TOP500 Retweeted

**SC21**  
 @Supercomputing  
 Congratulations to the TOP500 Supercomputers as announced at #SC20! Other honors:

- 1 HPCG | HPL-AI: Fugako
- 2 HPCG | HPL-AI: Summit
- 3 HPCG: Sierra
- 4 HPL-AI: Selene

JUWELS (Germany) is No. 1 in Europe

**Beyond Storage 2021**

Data & AI에 특화된 스토리지 솔루션

**IBM**



IBM Elastic Storage System 3200



*Performance*

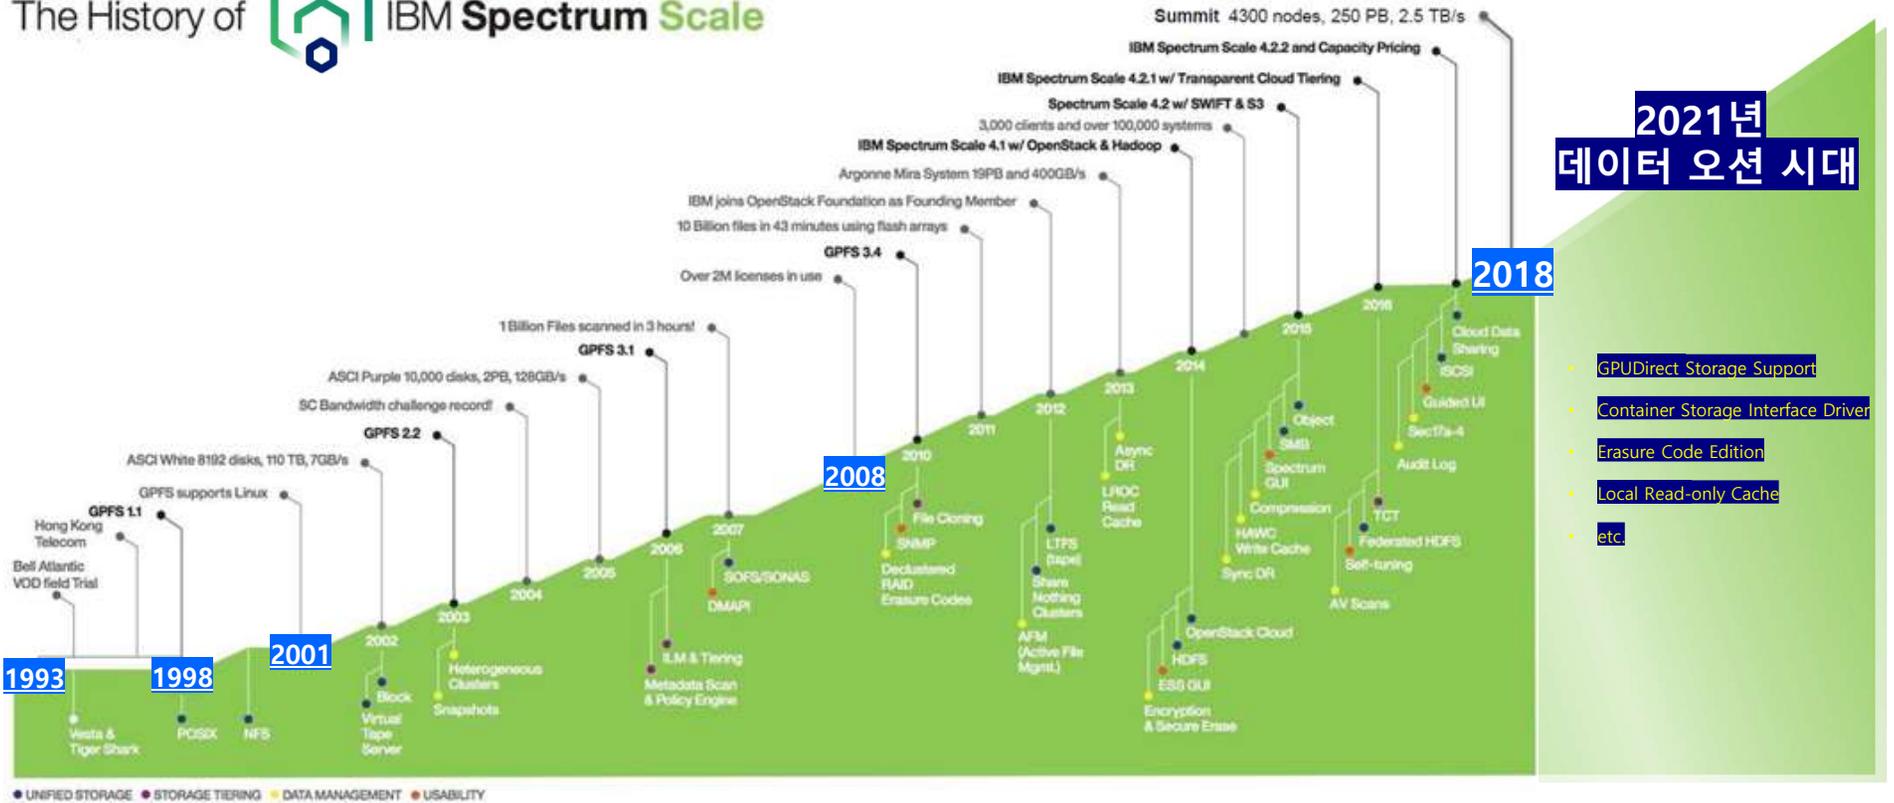
단일 노드당

**+80GB/s**

NVMe, PCIe G4 & HDR

## 검증된 솔루션

### The History of IBM Spectrum Scale



# Beyond Storage 2021

Data & AI에 특화된 스토리지 솔루션



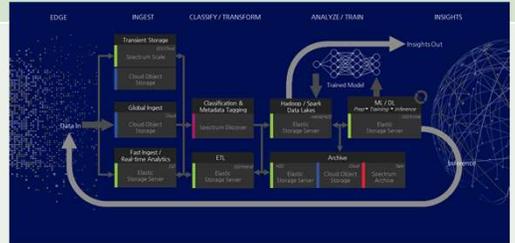
# AI & 빅데이터 최적화

	Summit@US
Workload	AI Deep Learning for Science
Supercomputer	#2 in Top500 : 148.6 PFLOPs
IBM Storage	IBM ESS : 250PB
# of GPUs	27,648 V100 GPUs
Reference	 <a href="https://www.top500.org/system/179397/">https://www.top500.org/system/179397/</a> <a href="https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/">https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/</a> <a href="https://www.ibm.com/blogs/nordic-mpsp/ibm-supercomputer-summit-attacks-coronavirus/">https://www.ibm.com/blogs/nordic-mpsp/ibm-supercomputer-summit-attacks-coronavirus/</a>

	DGX SuperPOD@US
Workload	AI Deep Learning for Science
Supercomputer	#26 in Top500 : 9.4PFLOPs
IBM Storage	IBM ESS : 552TB
# of GPUs	1,536 V100 GPUs
Reference	 <a href="https://www.top500.org/system/179691/">https://www.top500.org/system/179691/</a> <a href="https://developer.nvidia.com/blog/dgx-superpod-world-record-supercomputing-enterprise/">https://developer.nvidia.com/blog/dgx-superpod-world-record-supercomputing-enterprise/</a>

## End to End AI data flow support

IBM Spectrum Scale Elastic Storage System

- **Fast Ingest / Real-time Analytics** from EDGE
- **GPUDirect Storage** Support for Training
- **Hadoop transparency** with Data Lakes
- Easy deployment for AI applications with **CSI**

Beyond Storage 2021

Data & AI에 특화된 스토리지 솔루션



# 축적된 노하우

체크포인트 #1 : 다양한 구축 사례 확보

체크포인트 #2 : 전문 기술 지원 조직 확보





회사 소개   가치 창출   기업시민의식   파트너십



홈 / 뉴스 & 미디어 ▾ / 뉴스 / 아람코와 stc, 담맘 7(Dammam 7) 슈퍼컴퓨터 공개

# 아람코와 stc, 담맘 7 (Dammam 7) 슈퍼컴퓨터 공개

다란, 사우디아라비아, 2021/01/19



아민 나세르 (Amin H. Nasser) 아람코 대표이사 사장이 회사 관계자들과 단체 사진을 찍고 있다.



# Beyond Storage 2021

Data & AI에 특화된 스토리지 솔루션



## 업계 최고 성능



New ESS Gen3 Model  
IBM ESS 3200



### Performance Tier



ESS 3000  
4Q 2019

NVMe All-Flash 전용

### Capacity Tier



SLx SCx

ESS 5000  
3Q 2020

대용량 SAS Disk 전용

"고성능 병렬 처리 솔루션"



# Spectrum Scale

+

"최신의 하드웨어 컴포넌트"

**nvm**  
EXPRESS

# H·D·R

Infiniband 200G

PCI Express  
4.0

업계 최고의 성능 제공

# 80GB/s

per 2U System

# Beyond Storage 2021

Data & AI에 특화된 스토리지 솔루션



## 업계 최고 성능



### IBM Elastic Storage System 3200

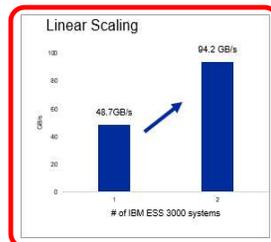


## WW 성능 테스트 결과

### IBM Storage performance with NVIDIA DGX A100 2/4/8 node POD configurations

#### IBM ESS 3000 results

- ESS 3000 demonstrated linear scaling
- Highest performing 2/4/8 node reference architecture



#### Typical customer recommended AI configurations

- One ESS 3000 in 2U for up to four DGX A100 nodes
- Two ESS 3000s in 4U for five to eight DGXA100 nodes with 94 GB/s
- Mellanox InfiniBand recommended for forward scalability
- Simple upgrades with SDS for more capacity and bandwidth
- Provides performance scalability, simple growth, and enterprise features for putting AI into production

Vendor	2 x DGX A100 Read perf	4 x DGX A100 Read perf	8 x DGX A100 Read perf
IBM - 2 x ESS 3000	42 GB/s	78 GB/s	94 GB/s
NetApp AFF A800 Pair	18 GB/s	25 GB/s	45 GB/s
Pure Storage	15 GB/s	25 GB/s	
Weka IO	34 GB/s	62 GB/s	
EMC Isilon F800	10 GB/s	19 GB/s	

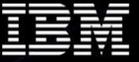
#### DDN

- No 2/4/8 node DGX System scaling results
- One DGX A100 with One AI400x
  - 49 GB/s read; 35 GB/s write
- One DGX A100 with Two AI400x = 99 GB/s read
- 10 DDN AI400x in SuperPOD
  - ~500 GB/s max IOR read



## Beyond Storage 2021

Data & AI에 특화된 스토리지 솔루션



“IBM ESS3200 은 AI 및 빅데이터 환경을 위한 최고 성능의 스토리지”

*Solution*



IBM Elastic Storage System 3200



*AI & Big data*

*Professional*

The image features the IBM logo centered on a dark blue background. The background is composed of various geometric shapes, primarily triangles and polygons, some of which are highlighted with glowing blue lines. The IBM logo itself is rendered in white, using its characteristic eight horizontal stripes. The overall aesthetic is modern and technological.

IBM