AI, 빅데이터 시대에 기업의 비즈니스 혁신을 위한 데이터 플랫폼

송성환 부장 / 퓨어스토리지코리아

PURESTORAGE®

4차 산업혁명의 시대

BIG DATA, AI DRIVING CHANGE IN EVERY INDUSTRY

1st Revolution
1760-1820's
Steam Power
Rural to Industrial

2nd Revolution
1870-1914
Electricity
Industrial to Mass Production

3rd Revolution
1980-2010
PC
Mass production to Digital

4th Revolution
2010-now
AI, Big Data & IoT
Digital to Intelligence

새로운 데이터 활용 전략

“데이터 허브”

PURESTORAGE®

데이터의 종속성
IN A SPRAWL OF SILOS

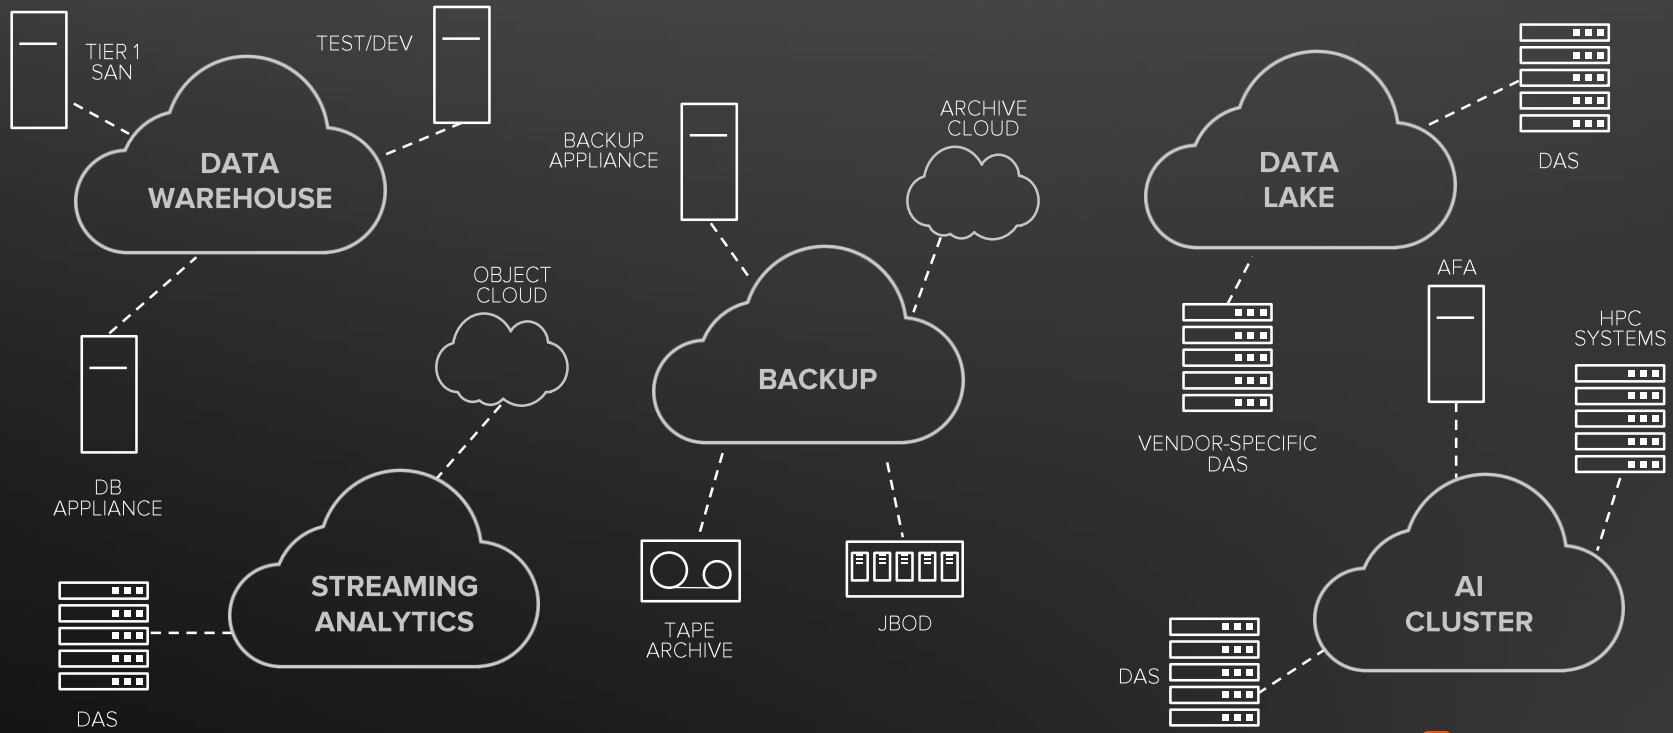

PURESTORAGE®

4가지 분리된 IT 환경
IN THE WORLD OF ANALYTICS

DATA WAREHOUSE

MySQL · ORACLE · Microsoft SQL Server · TERADATA

1990

DATA LAKE

APACHE HBASE · APACHE HIVE · hadoop Map Reduce · APACHE Spark

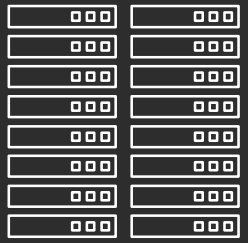

2000

STREAMING ANALYTICS

elastic · docker · kafka · logstash

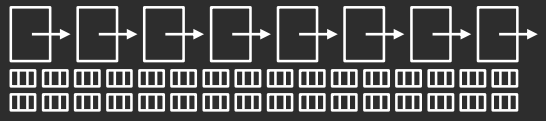

2010

AI CLUSTER

Caffe2 · mxnet · PYTORCH · TensorFlow

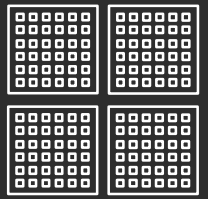

2015

PURESTORAGE®

데이터 통합의 어려움
HERE'S WHY

	BACKUP & DATA WAREHOUSE	DATA LAKE	STREAMING ANALYTICS	AI CLUSTER
	 APPLIANCES	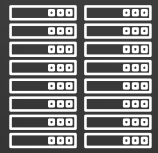 DAS	 DISAGGREGATED	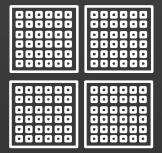 HPC SYSTEM W/ GPU
DATA TYPE	Structured	Unstructured	Unstructured	Unstructured
PROCESSING MODE	Batch	Batch	Micro-batch / real-time	Real-time
I/O TYPE	Random Read	Sequential	Random	Sequential to random
ARCHITECTURE	Scale-up	Scale-out	Multi-dimensional	Massively parallel

클라우드로의 확산 또는 전환의 고민

CLOUD-NATIVE APPS REQUIRE DISAGGREGATED ARCHITECTURE WITH OBJECT STORE

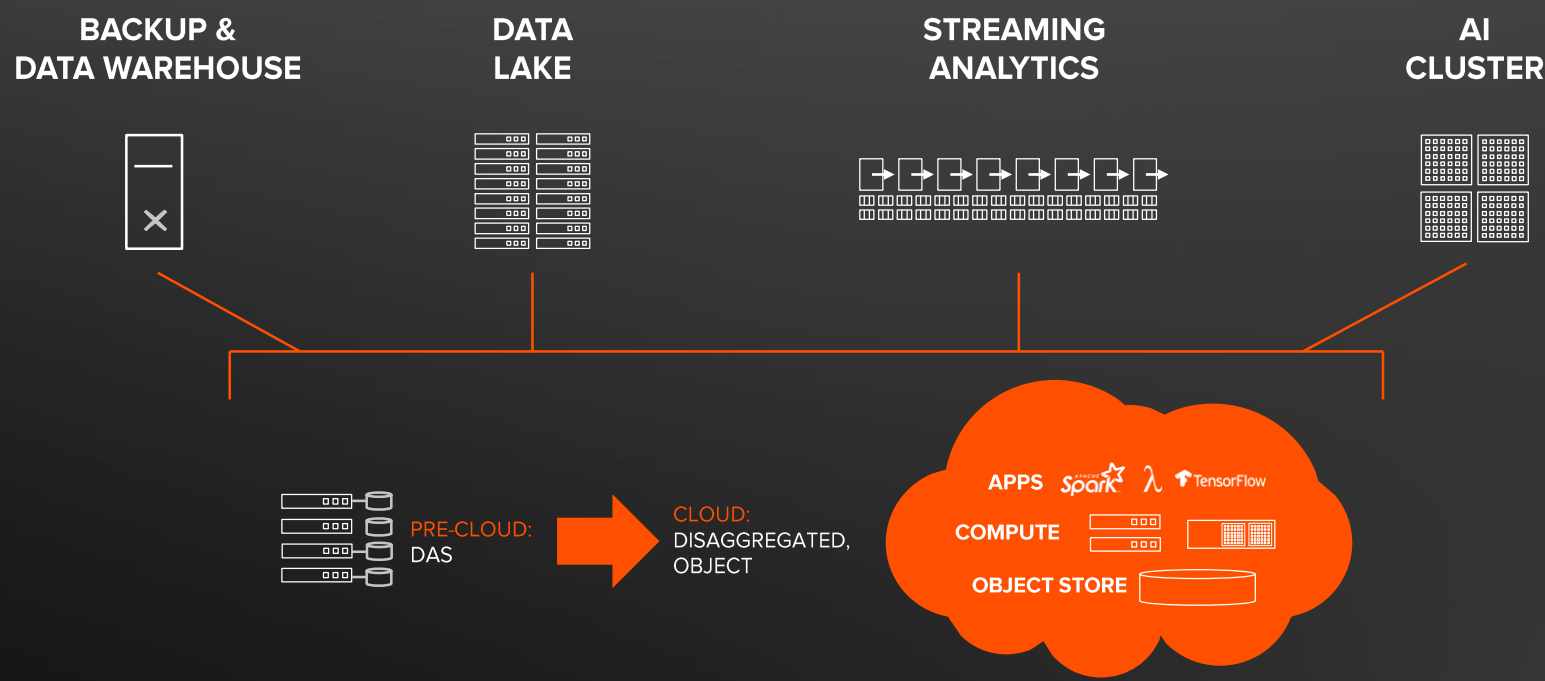

PURESTORAGE®

데이터 통합을 위한 필요 사항

FOUR ESSENTIAL FEATURES OF EACH SILO MUST BE BUILT INTO A SINGLE PLATFORM

	BACKUP & DATA WAREHOUSE	DATA LAKE	STREAMING ANALYTICS	AI CLUSTER
	 Appliances	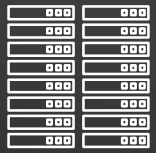 DAS	 Disaggregated	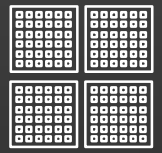 HPC System w/ GPU
KEY REQUIREMENT FOR STORAGE	 High Throughput File & Object	 Seamless, Native Scale-Out	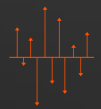 Multi-dimensional Performance	 Massively Parallel

PURESTORAGE®

업계 최초 데이터 허브 플랫폼 - FLASHBLADE™

THE INDUSTRY'S FIRST DATA HUB

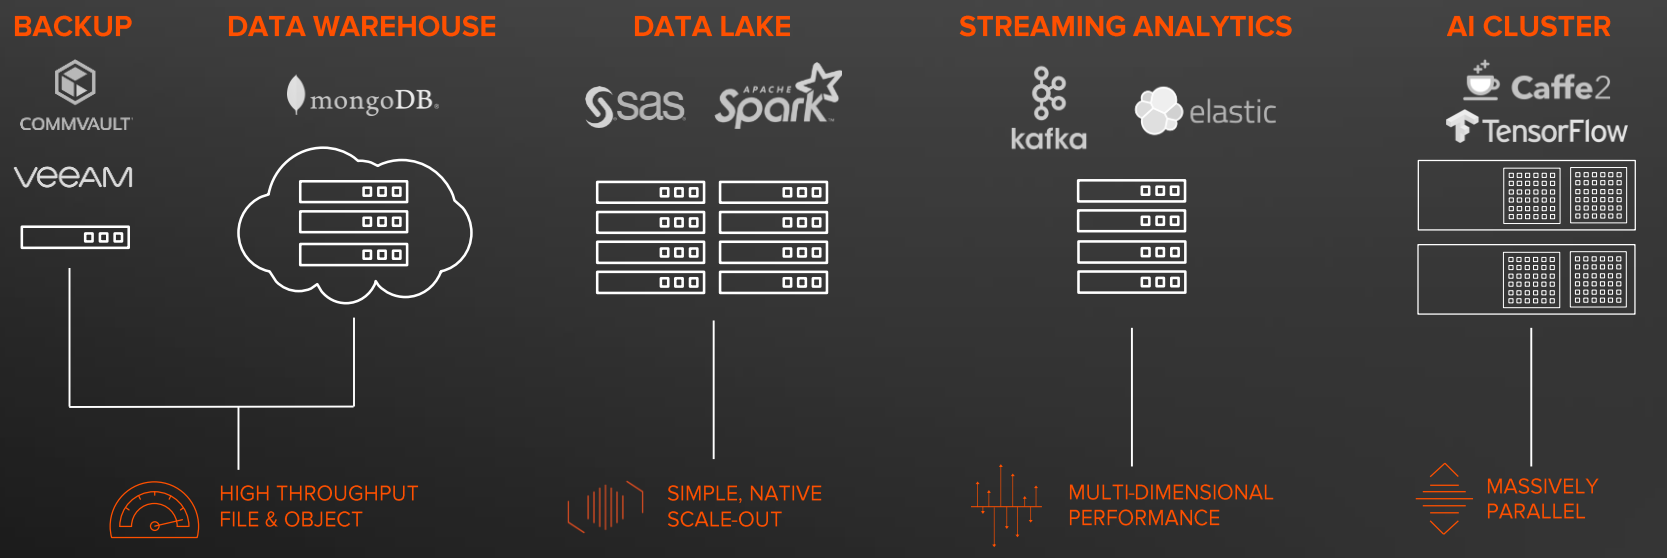

PURESTORAGE®

FlashBlade is industry's first data hub for storage infrastructure

새로운 **AI** 전략

“AI Ready Infrastructure”

PURESTORAGE®

AI - 모든 산업의 변화

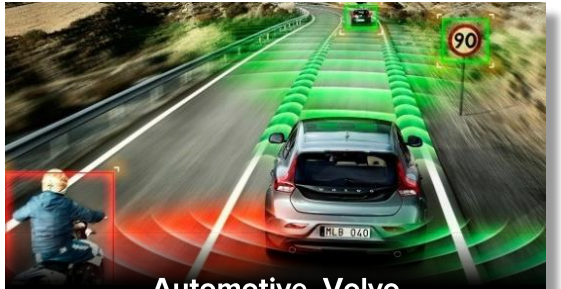

Automotive- Volvo

Volvo and its subsidiary, Zenuity, are racing to build fully autonomous cars by 2021

Farming- Blue River

10% of lettuce in the US is harvested by LettuceBot, using AI to maximize crop yield & minimize chemicals

Kitchen- LG

LG Instaview refrigerator with AI-powered Alexa helps owners shop for groceries with their voice

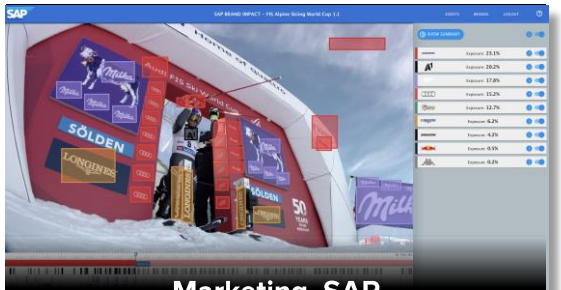

Marketing- SAP

SAP's Brand Impact tool measures ROI of brand sponsorships using video analytics and AI

Healthcare- Mayo Clinic

Finds genetic markers in images to avoid surgery for tumor samples & recommend treatments

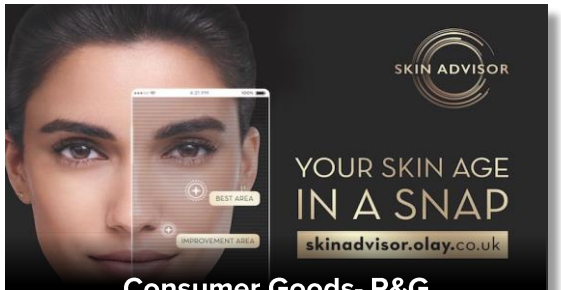

Consumer Goods- P&G

Proctor & Gamble's Olay is using AI to inspect skin & improve trouble spots, all with a mobile phone

“We don’t have better algorithms. We just have more data.”

PETER NORVIG
Google Research Director

PURESTORAGE®

기존 AI 프로세스

INGEST

From sensors, machines, & user generated

CLEAN & TRANSFORM

Label, anomaly detection, ETL, prep, stage

CPU Servers

EXPLORE

Quickly iterate to converge on models

GPU Server

TRAIN

Run for hours to days in production cluster

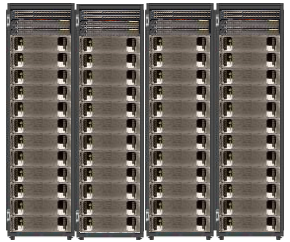

GPU Production Cluster

PURESTORAGE®

"DO-IT-YOURSELF" IS OFTEN THE ONLY OPTION

NEVER-ENDING CYCLES OF COMPILING & TUNING OPEN SOURCE SOFTWARE

MONTHS OF SYSTEM BUILDING AND TUNING, CONSTANT MAINTENANCE

YET LEGACY SOLUTIONS FULL OF DATA BOTTLENECKS, FROM STORAGE TO GPU TO APPS

AI환경의 IO Flow

	INGEST	CLEAN & TRANSFORM	EXPLORE	TRAIN
	From sensors & machines	CPU Servers	GPU Server	GPU Production Cluster
Access Pattern	sequential	sequential or random	random	random
Access Type	write	read & write	read	read
File Size	metadata is small data is small to large	small to large	small to large	small to large
Concurrency	depends on # of sources	high	low	high

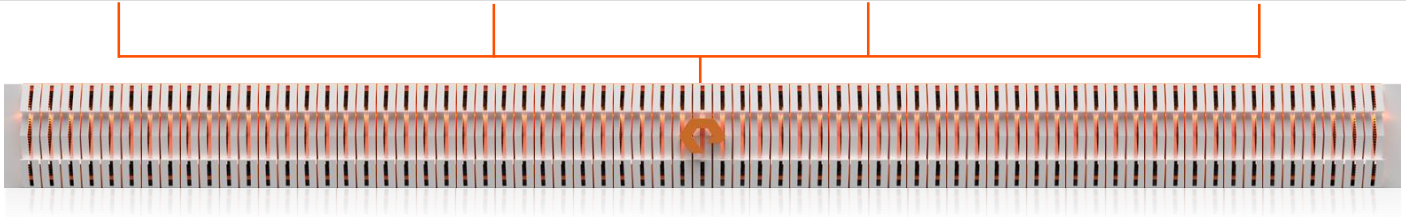

AI 인프라를 위한 3가지 구성 요소

FRAMEWORKS & EXPERTISE

Caffe2

mxnet

PYTORCH

TensorFlow

COMPUTE
FROM CPU TO GPU SERVERS

NVIDIA

STORAGE
POWER ENTIRE AI PIPELINE

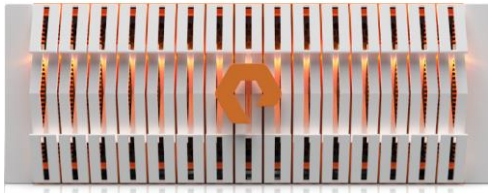

PURESTORAGE®

SIMPLICITY FABRIC DESIGN

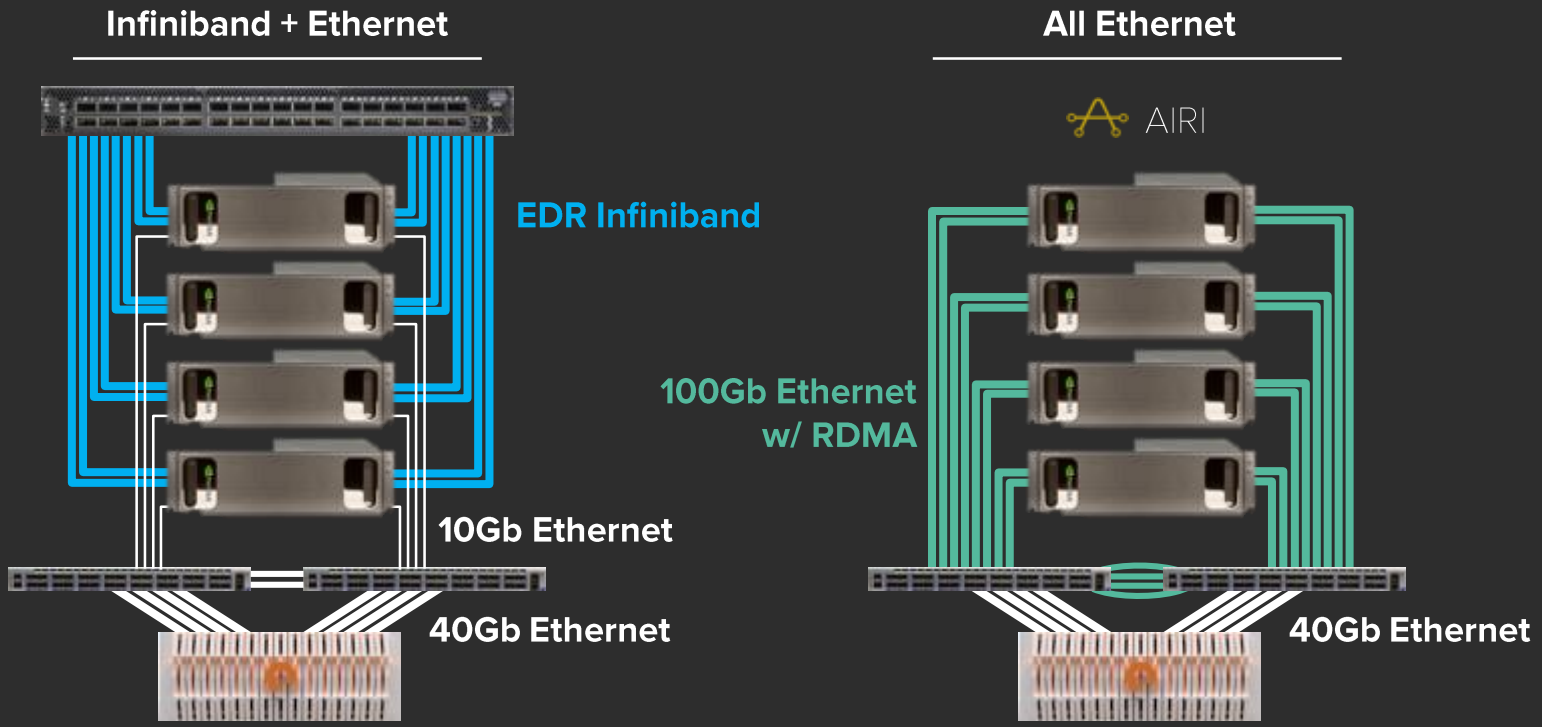

PURESTORAGE®

AIRI

THE INDUSTRY'S FIRST

COMPLETE AI-READY INFRASTRUCTURE

HARDWARE

AIRI
NVIDIA® DGX-1™ | 4x DGX-1 Systems | 4 PFLOPS of DL Performance
PURE FLASHBLADE™ | 15x 17TB Blades | 1.5M IOPS
CISCO or ARISTA | 2x 100Gb Ethernet Switches with RDMA

AIRI "mini"
NVIDIA® DGX-1™ | 2x DGX-1 Systems | 2 PFLOPS of DL Performance
PURE FLASHBLADE™ | 7x 17TB Blades | 700K IOPS
CISCO or ARISTA | 2x 100Gb Ethernet Switches with RDMA

SOFTWARE

NVIDIA GPU CLOUD DEEP LEARNING STACK | NVIDIA Optimized Frameworks
AIRI SCALING TOOLKIT | Multi-node Training Made Simple

100% NVMe방식의 고성능 오브젝트 스토리지

FLASHBLADE

4U

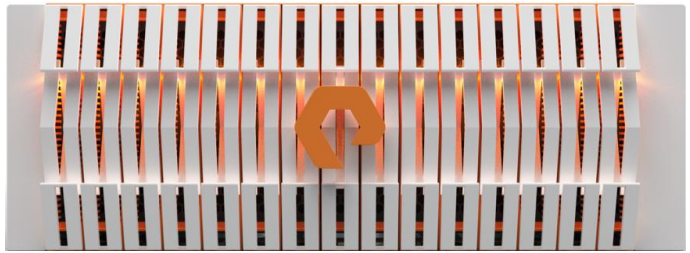

1PB

BLADE

SCALE-OUT
PROCESSING + FLASH

ELASTIC Fabric Module

LOW-LATENCY,
SW-DEFINED ETHERNET
INTERCONNECT

Purity//FB2

SCALE-OUT STORAGE
SOFTWARE

PURESTORAGE®

전용 FPGA및 SW기반 데이터 콘트롤

BLADE

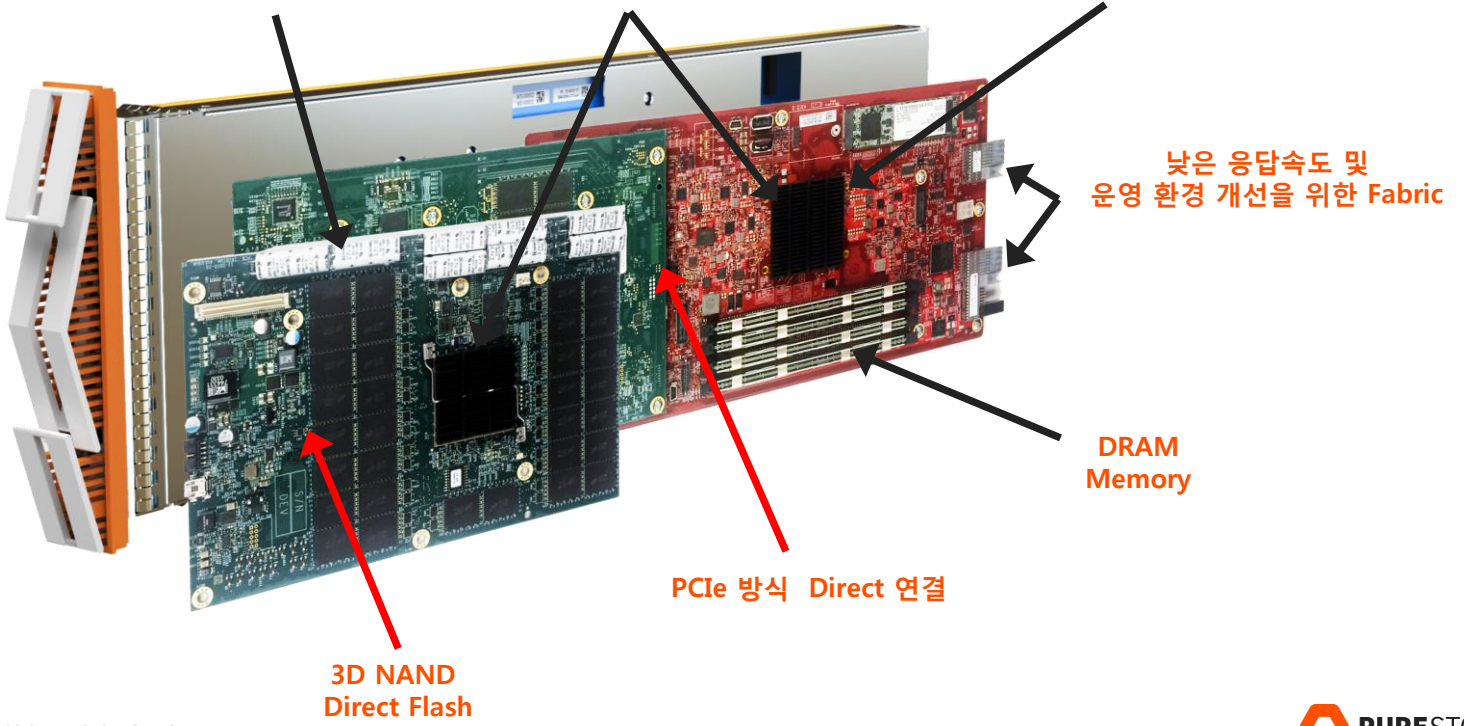

PURESTORAGE®

고밀도 확장성

BLADE CHASSIS

FLASHBLADE CHASSIS

- Up to 15 Blades
- 4RU Height
- N+2 Redundant

FABRIC MODULE

- Dual switched mid-plane
- 30 x Internal 10Gbit ports
- 8 x 40GbE External ports
- Each 40GbE can be broken into 4x 10GbE

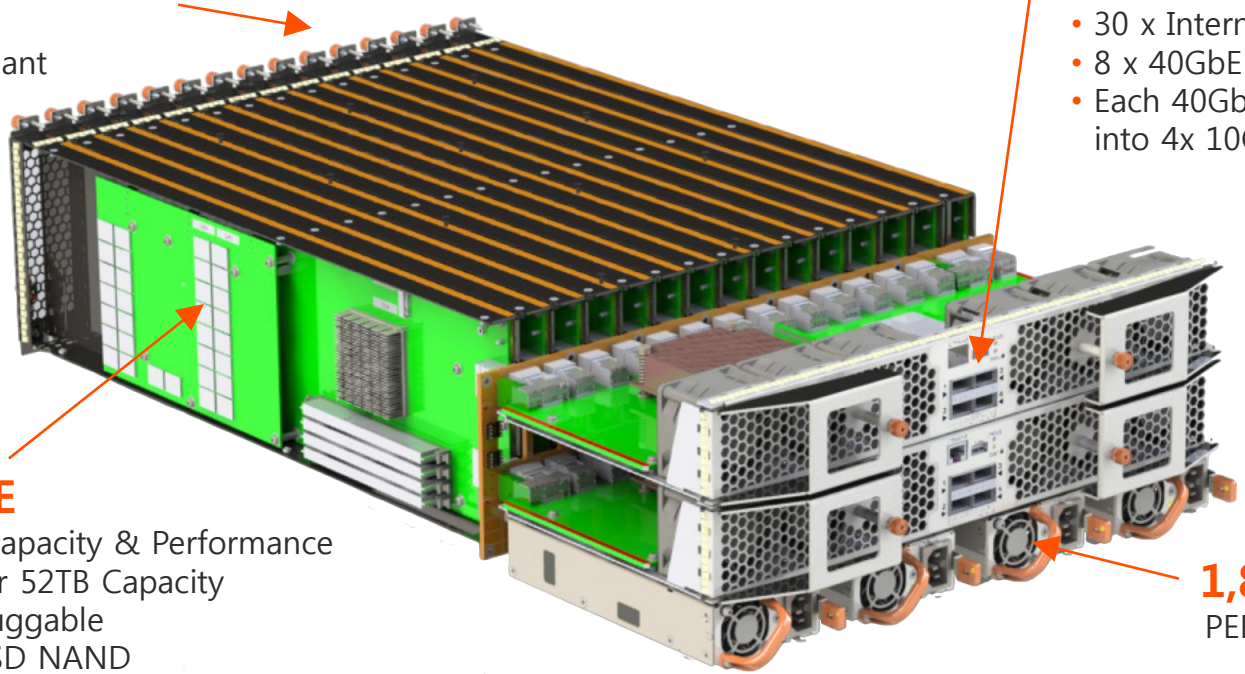

BLADE

- Adds Capacity & Performance
- 17TB or 52TB Capacity
- Hot Pluggable
- Non SSD NAND
- Embedded NVRAM

1,850 WATT

PER PETABYTE USABLE

PURESTORAGE®

DELIVERING ELASTIC SCALE AT EVERY POINT

SCALE-OUT ARCHITECTURE BUILT TO GROW PERFECTLY WITH YOUR AI JOURNEY

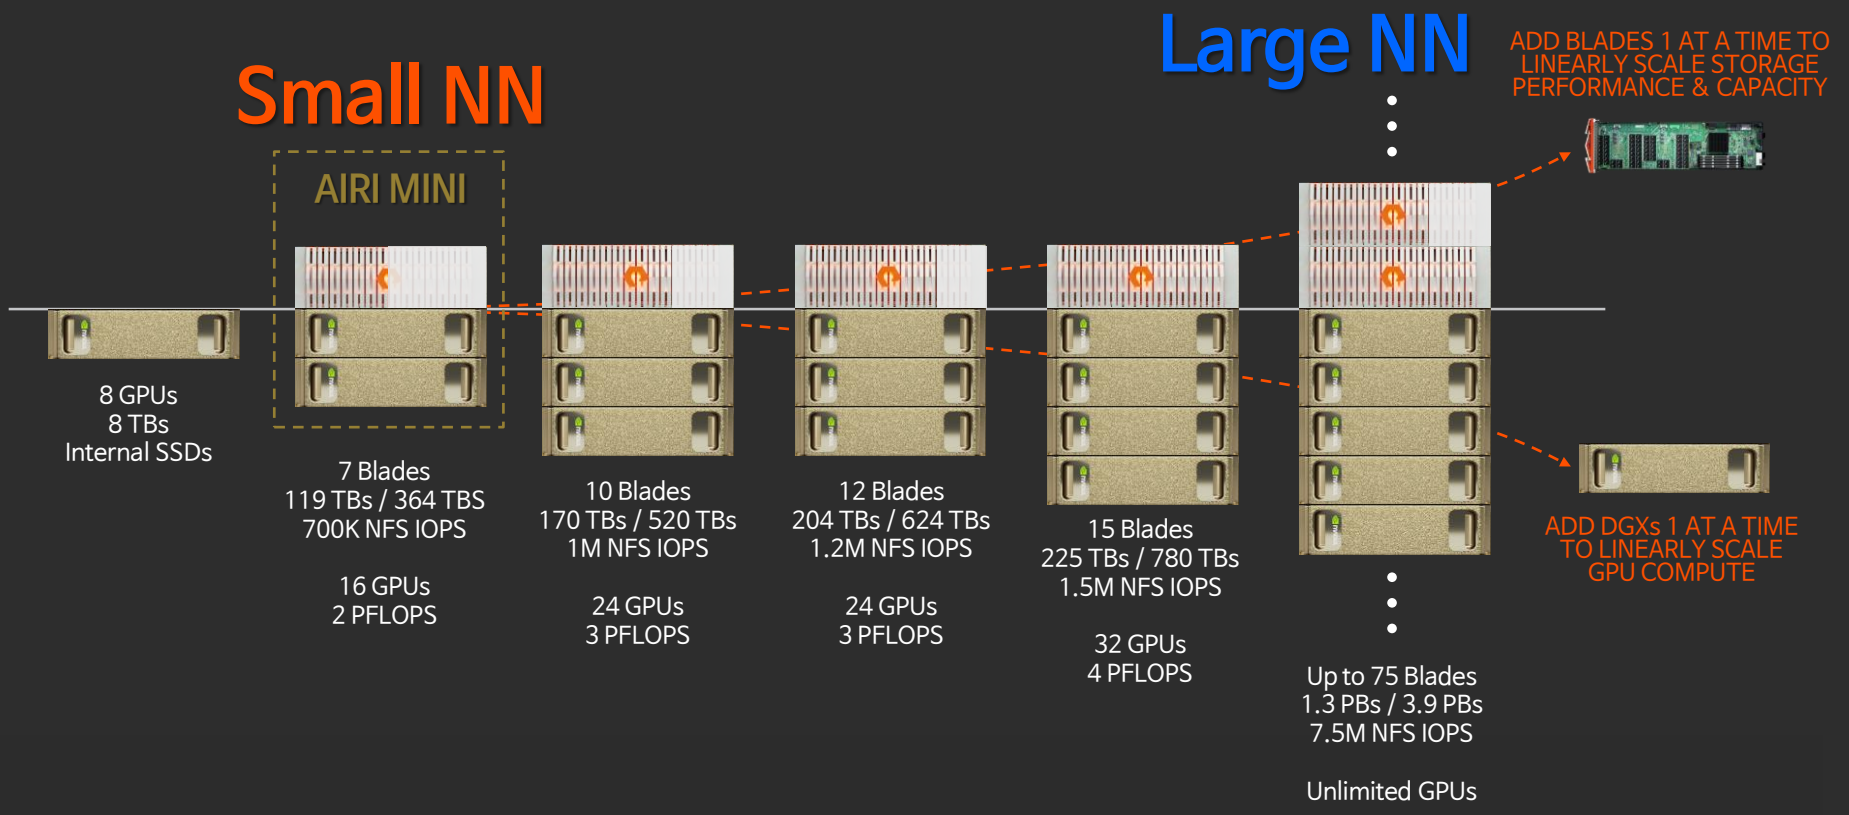

Why FlashBlade For AI?

Only FlashBlade Value

FlashBlade의 병목 제거 기술

- 클라이언트 로드밸런싱 : 소프트웨어 기반 분산 네트워킹
- 스케일아웃 DB코어 : 확장에 따른 성능을 보장하기 위한 분할 데이터베이스
- DirectFlash/NVMe 기반의 병렬 미디어 엑세스
- Object 기반의 Reference Table

Core 아키텍처

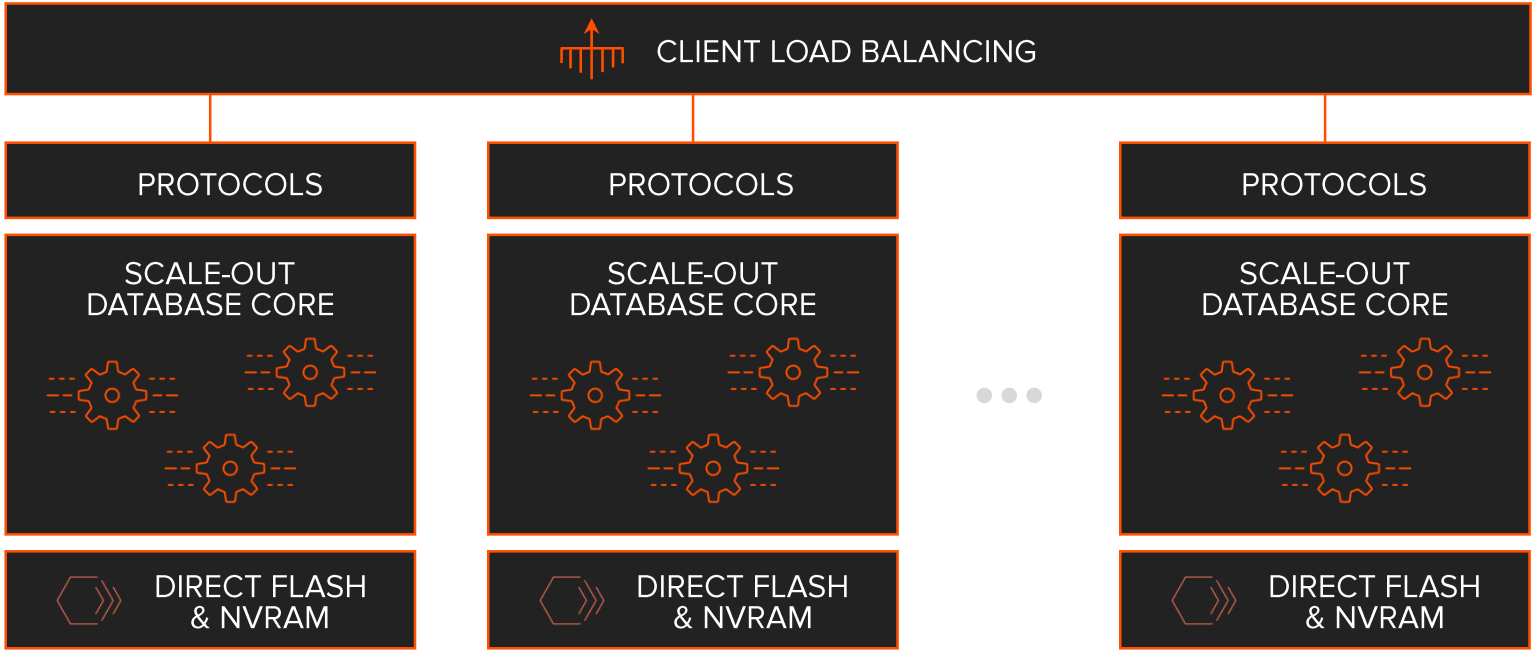

대규모 분산 병렬 처리 구조

NVMe, Direct Flash

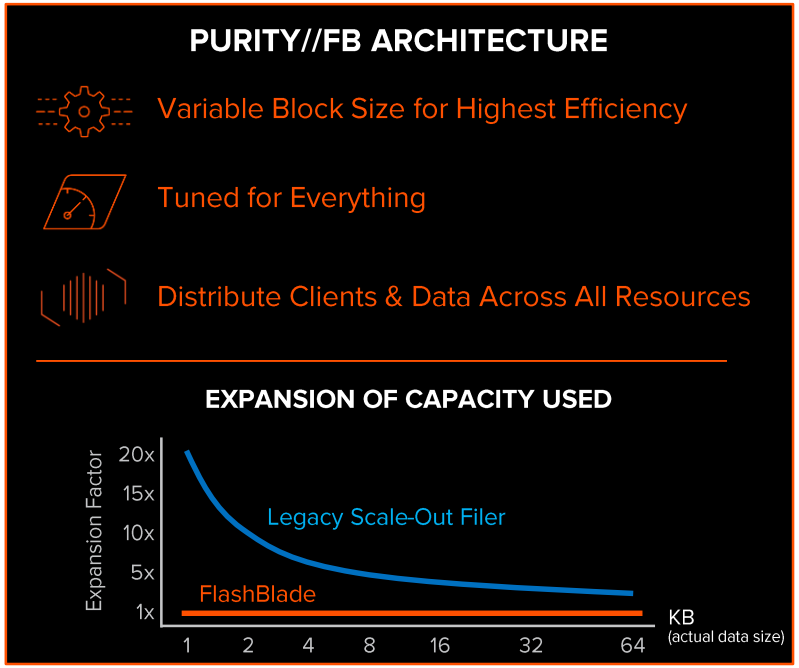

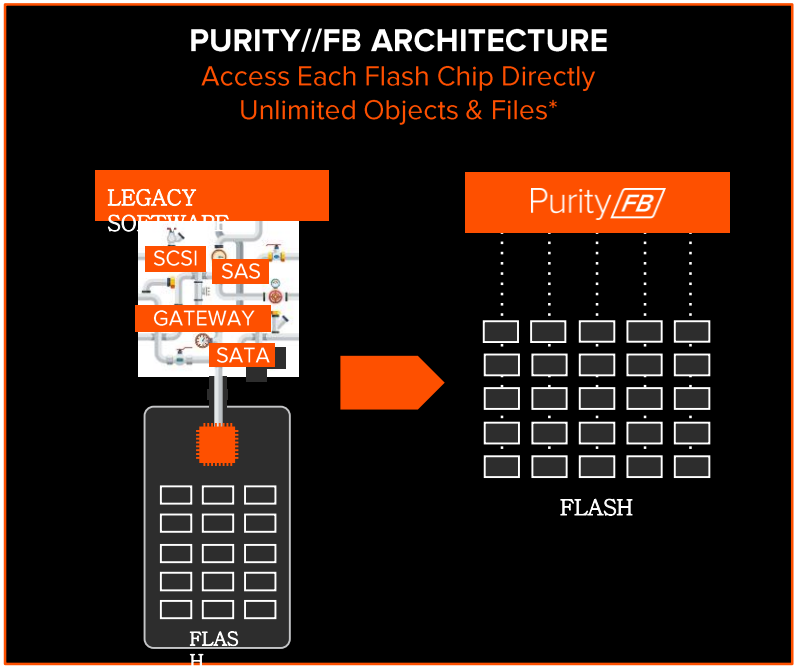

PURESTORAGE®

File System Search 문제점

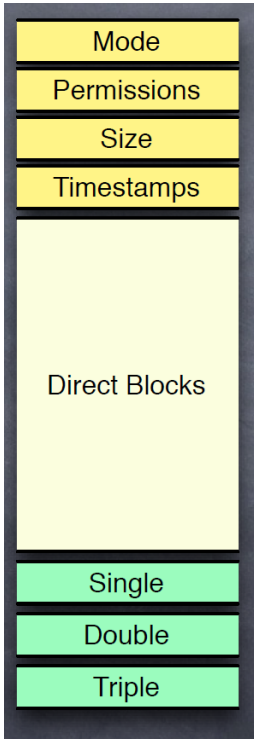

File개수/File크기/Directory Depth에 따라 성능 저하

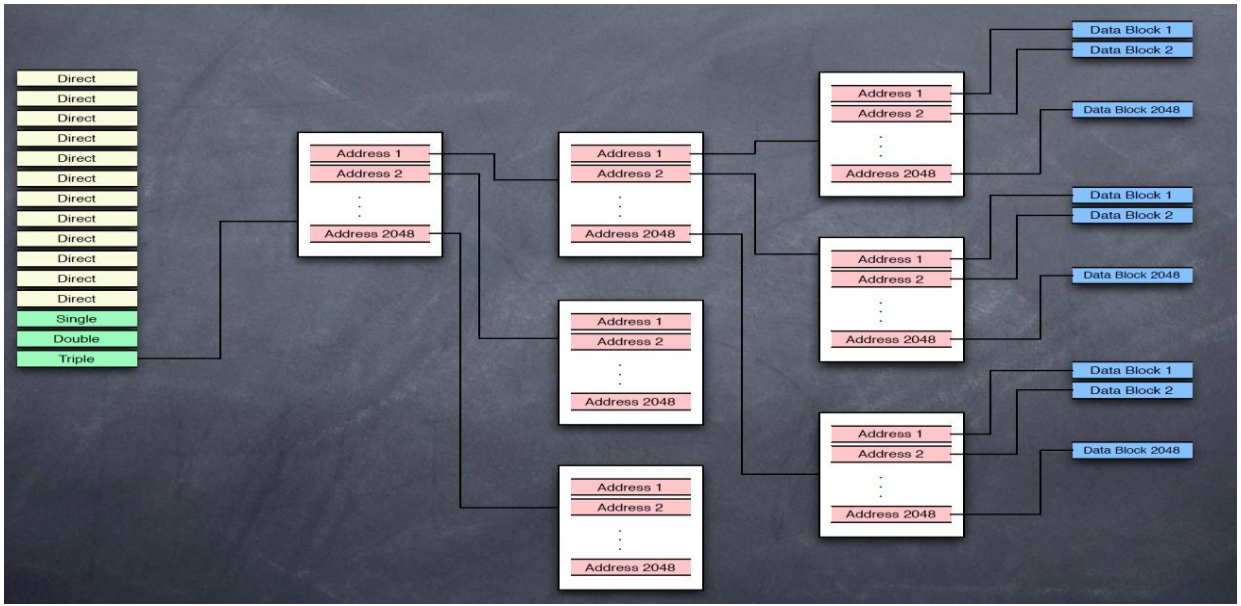

PURESTORAGE®

Object 기반 Reference Table

/home/test/labeled/12345.jpg

Get (directory, segment)

Get (segment, Blade)

Get (Blade, Object)

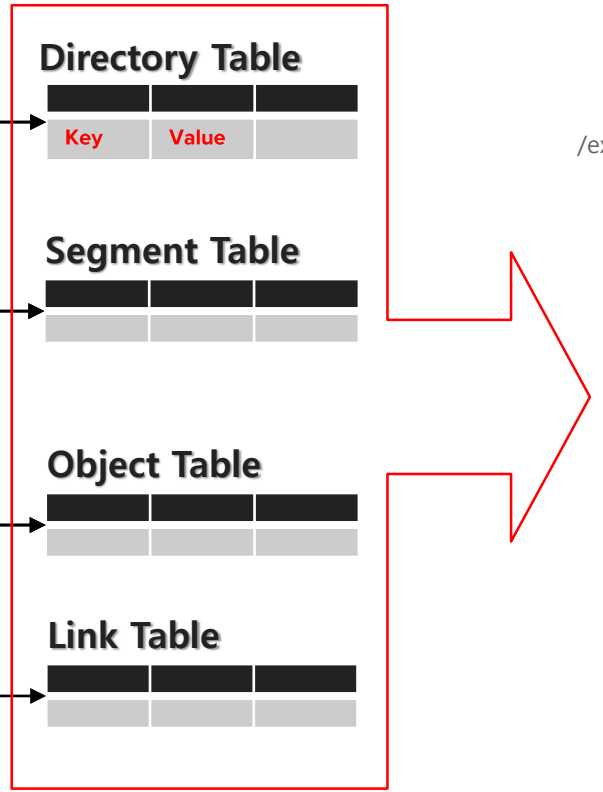

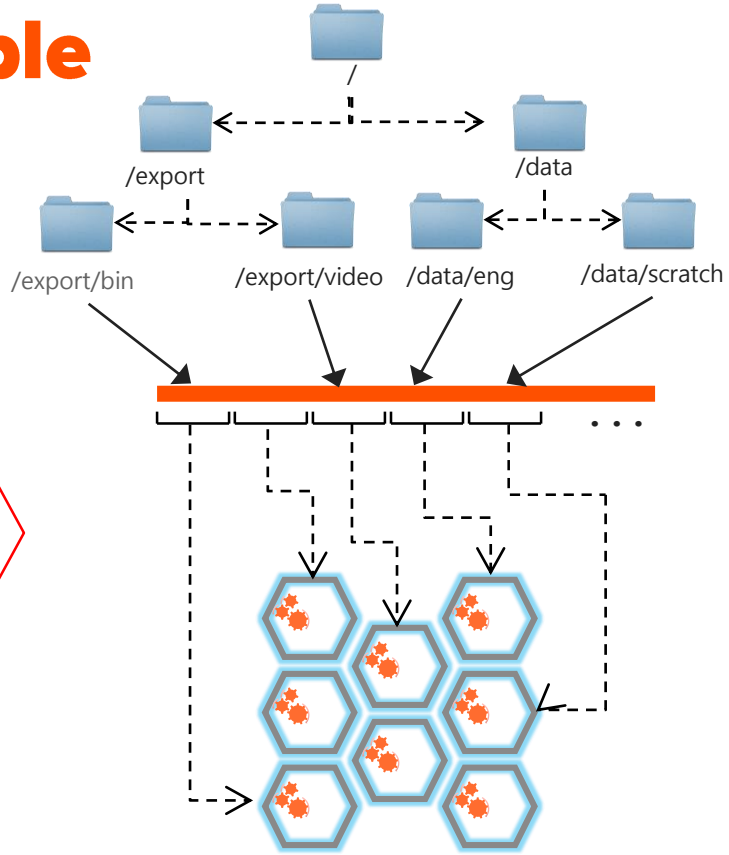

PURESTORAGE®

AI 성공 사례

AI/ML/DL 성공사례

고객 요구 사항

- 스케일아웃이 용이해야 하며, 확장 시 성능 저하가 없어야 함
- 성능 저하를 최소화 할 수 있는 아키텍처

용도

- GPU 서버 클러스터 60여대, Linux 서버 40여대
- 텍스트, 이미지, 동영상 등 다양한 데이터 처리 및 훈련

도입 효과

- 데이터 재가공 시간 대폭 단축 / 데이터 활용도 증가
- NVMe 기반의 고성능 아키텍처 업무 생산성 향상

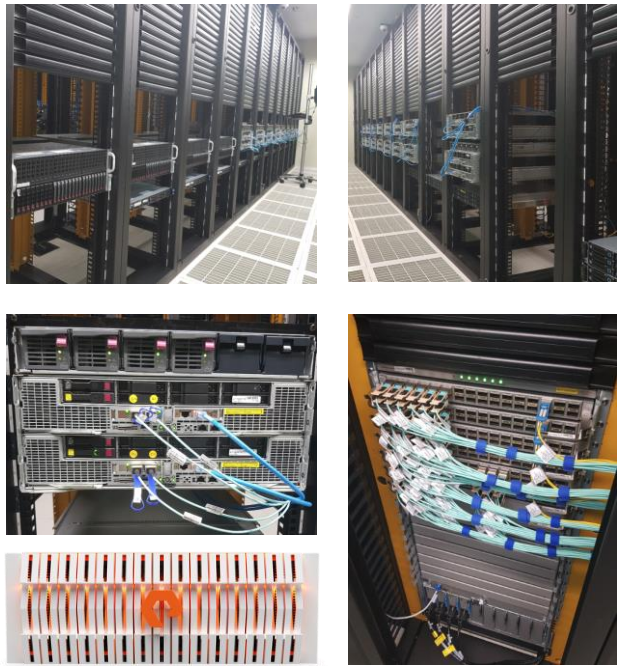

AI/ML/DL 성공사례

10대의 FlashBlade 클러스터와 2PB 이상의 데이터 운영

고객 요구 사항

- 고성능 AI 훈련 인프라 필요
- 요구성능 : random read & write 성능 - 약 70 & 10GB/s 이상

용도

- 128여대 DGX-1 노드와 70여대 리눅스 서버에서 병렬 처리를 위한 고성능 스토리지
- 분당 14만 건의 이미지 업로드, 35만 건의 자료 업데이트를 처리하고 이를 이용하여 AI관련 처리 목표
- 문자 분석, 안면 인식, 타겟 광고, AI 어플리케이션 디자인, 예측 분석 등

도입 효과

- 고객의 요구 성능 충족, 향후 용량 및 성능 확장 유연성 확보
- 상면, 전력, 냉각 비용 절감, 실시간 데이터 공유