인공지능, 퓨어스토리지를 만나다.

2018.01.09

퓨어스토리지 홍준혁

PURESTORAGE®

Artificial

"**Artificial** is Made or produced by human beings rather than occurring naturally, especially as a copy of something natural."

Intelligence

"**Intelligence** is the ability to think, reason, and understand instead of doing things automatically or by instinct."

PURESTORAGE®

A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence

August 31, 1955

John McCarthy, Marvin L. Minsky,
Nathaniel Rochester,
and Claude E. Shannon

Google Trend (2012 ~ 현재)

발췌 : Google Trend


언어이해
감정표현
음성인식
이미지 인식
언어표현

인공지능 로봇 (Human-Like) 프로세스

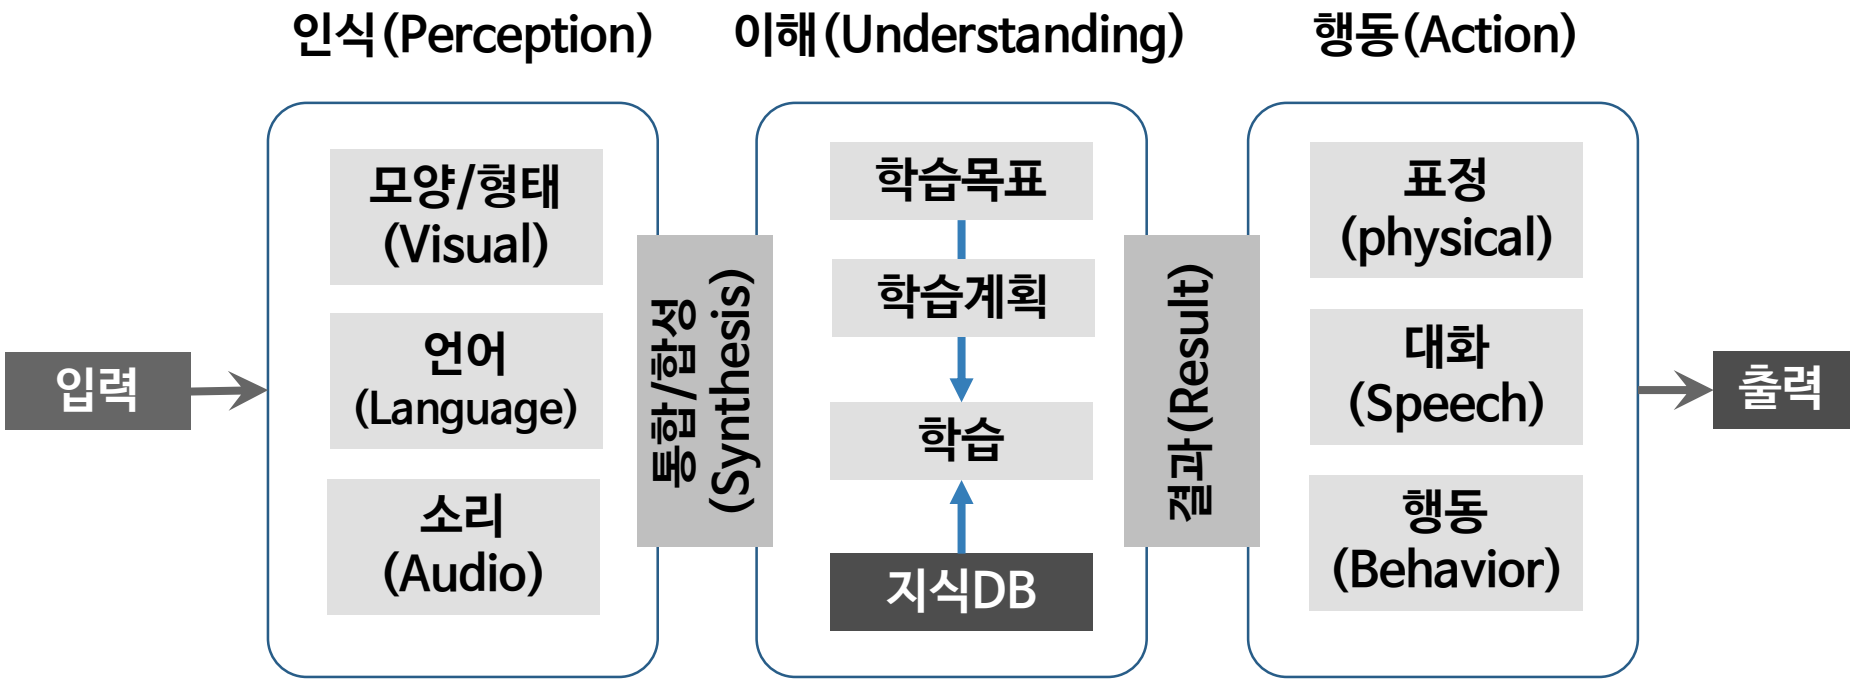

Learning 과정

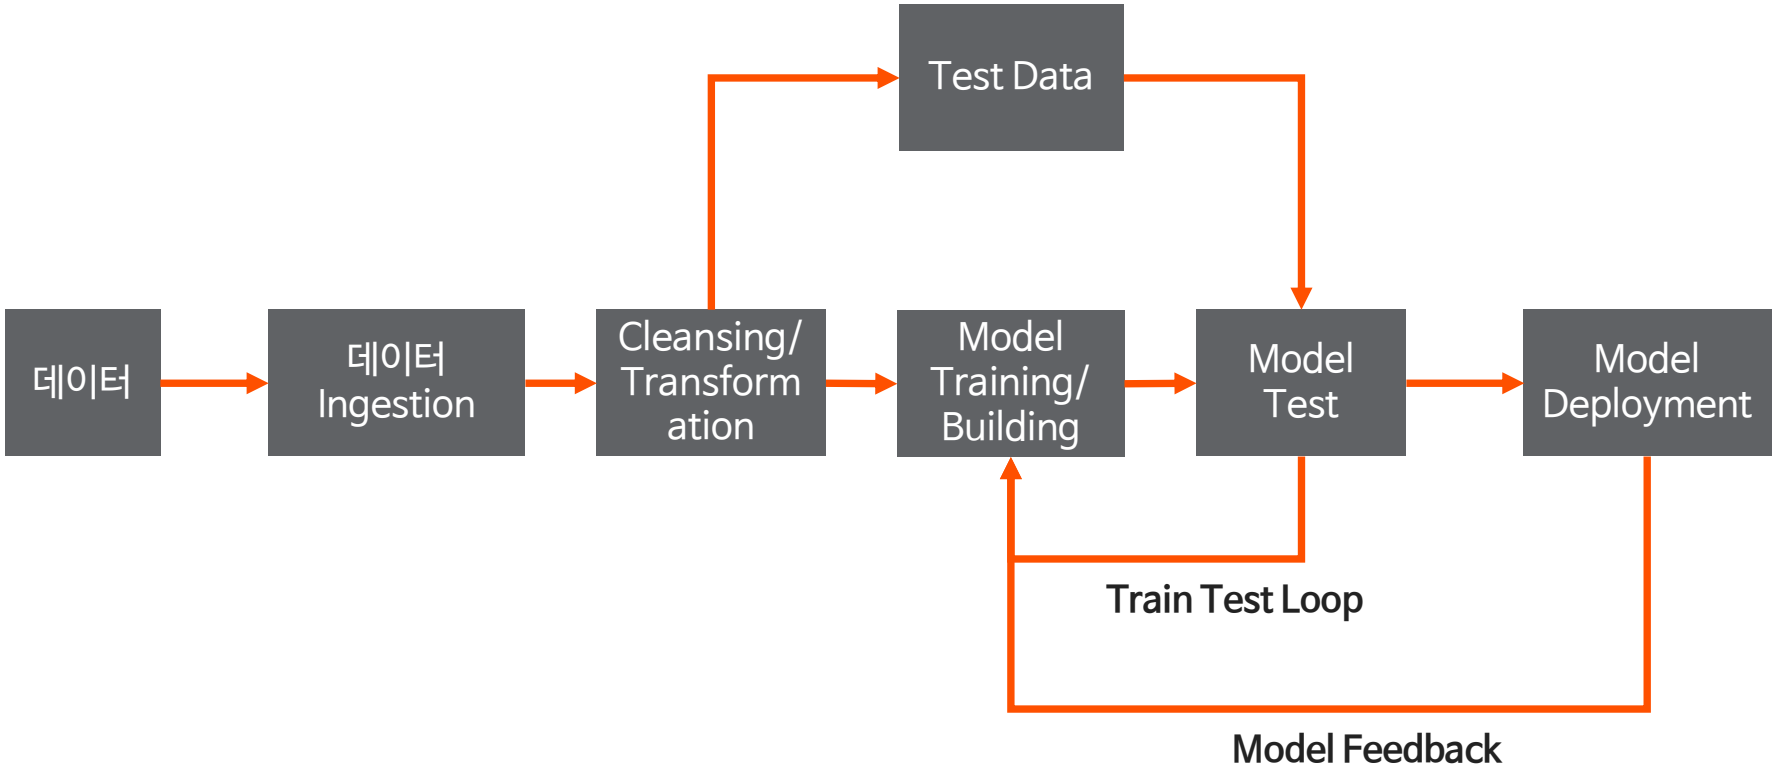

Engineering vs. Data

다양한 방식/방법 Learning Algorithm

- Memory Based
- Winnow
- Perception
- SVM

...

> “ It’s not who has the best algorithm that wins. It’s who has the most data. ”

Andrew Ng

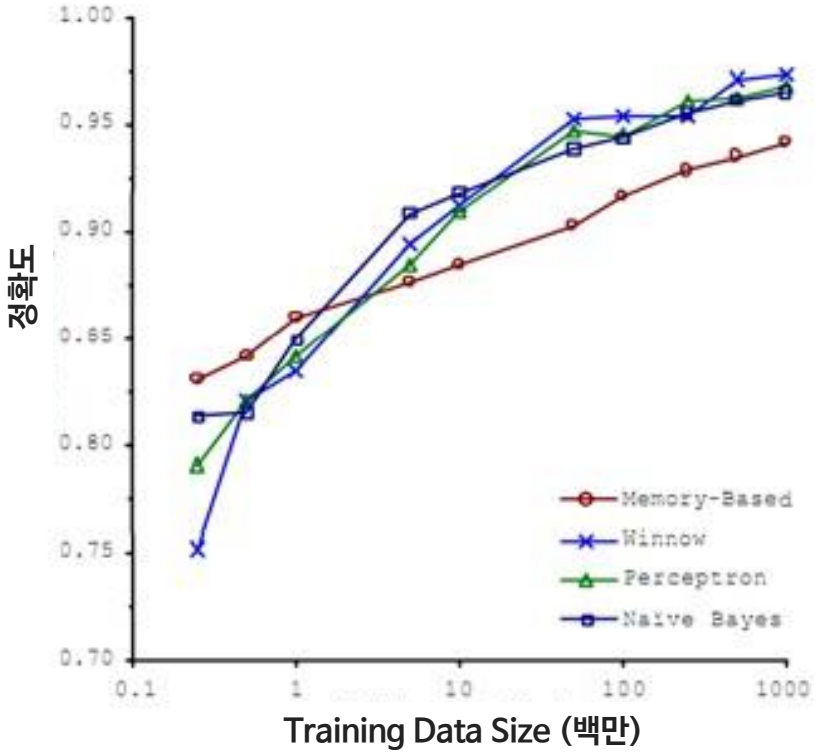

“We don’t have better algorithms. We just have more data.”

PETER NORVIG
Google Research Director

PURESTORAGE®

데이터 활용증가에 따라 데이터플랫폼 중요성

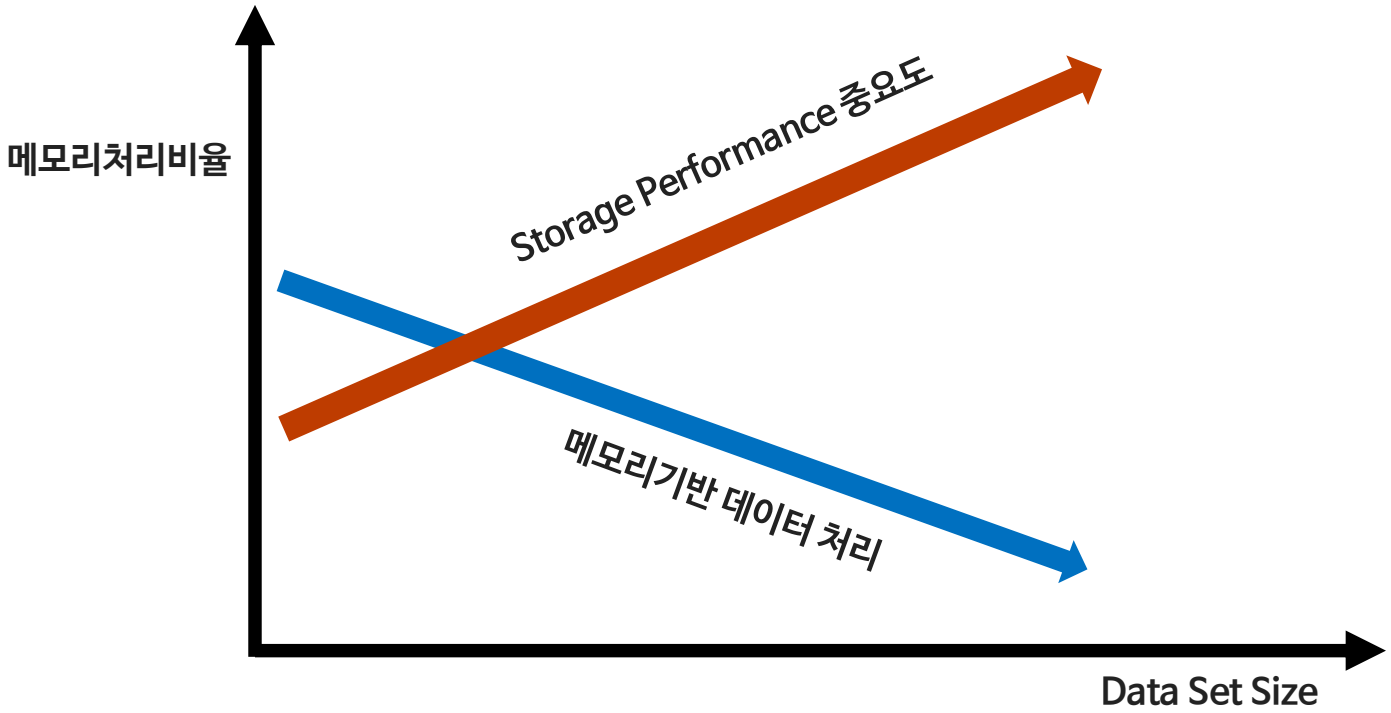

PURESTORAGE®

인공지능 데이터 프로세싱 과정별 특징

데이터 접근형태	sequential	sequential or random	random	random
데이터 접근방법	write	read & write	read	read
데이터 요청/처리 사이즈	metadata is small data is small to large	small to large	small to large	small to large
프로세스 활용	**IO Bound**	**CPU**	**GPU**	**CPU**
동시 사용 프로세스 수	**depends on # of sources**	**high**	**High**	**Depends on # of target**

수집,처리 단계의 이슈

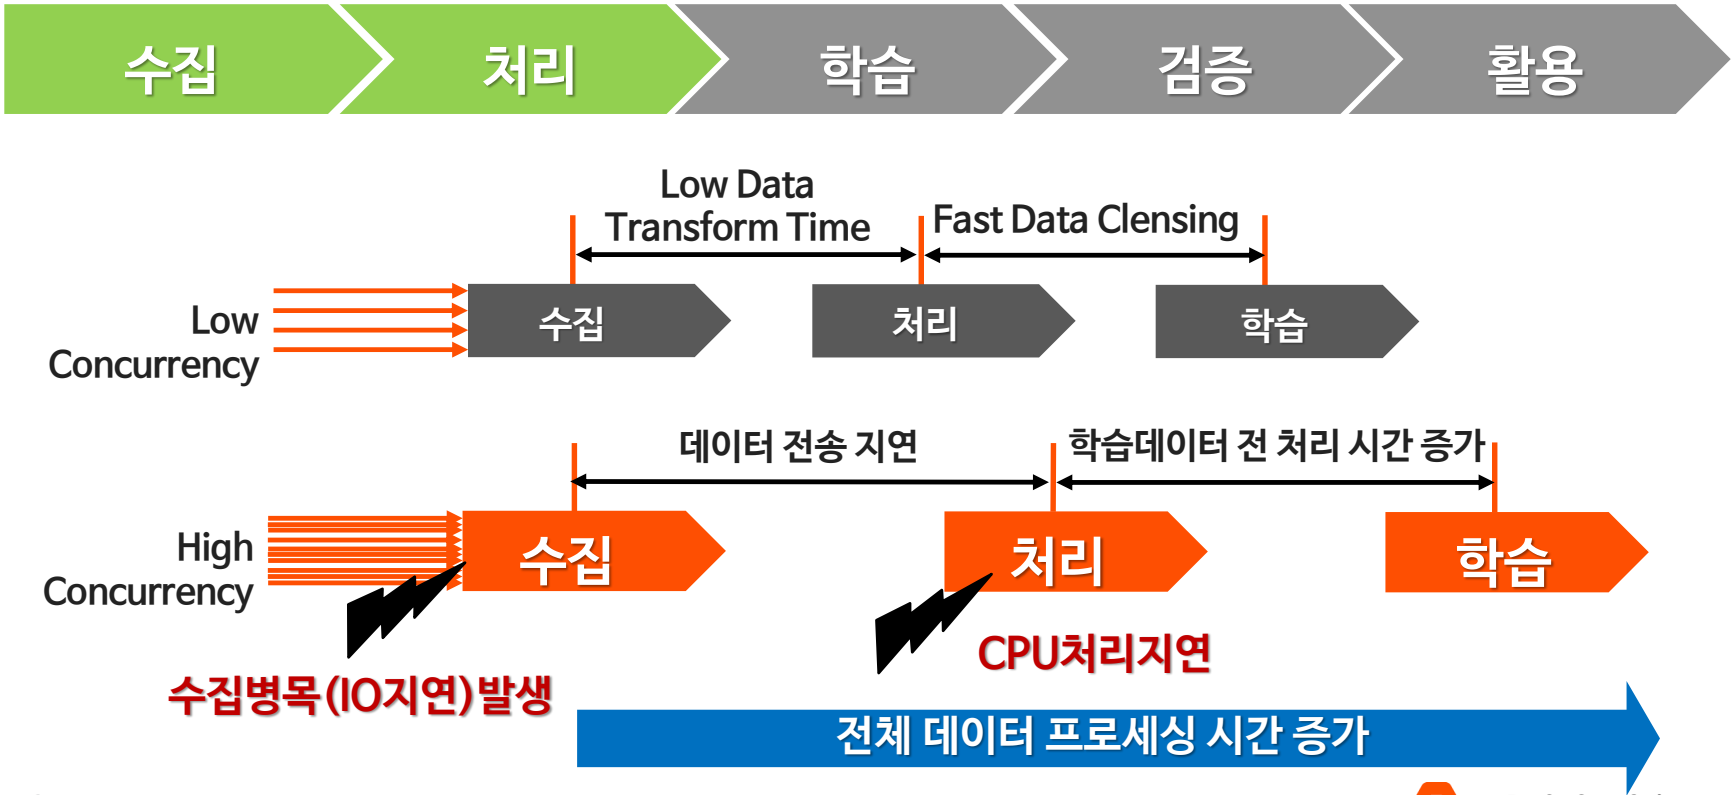

PURESTORAGE®

수집,처리 단계 개선 – 새로운 데이터 관리 기술

안정적인 응답속도 보장 (Stability)

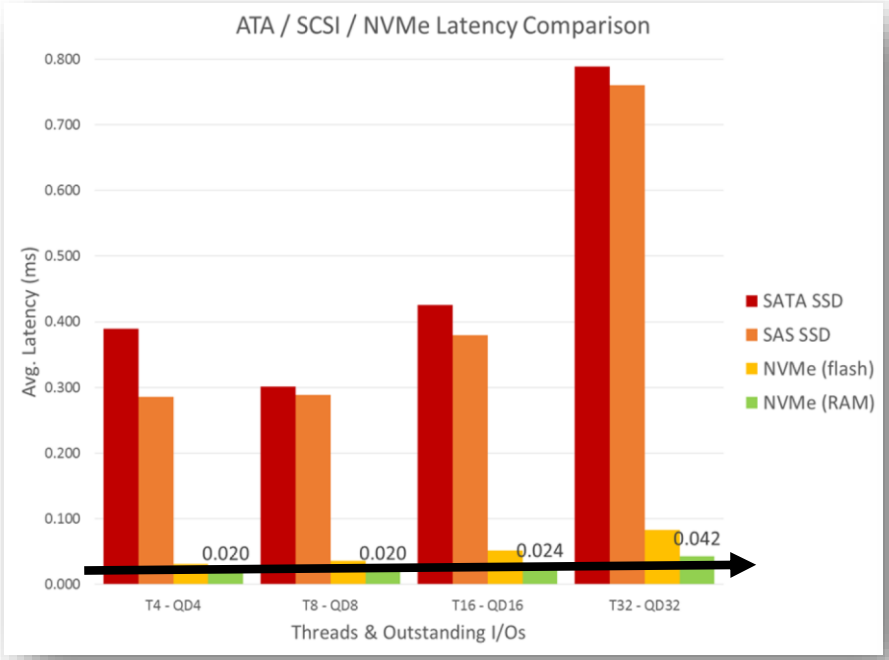

데이터 처리능력 확보 (Performance)

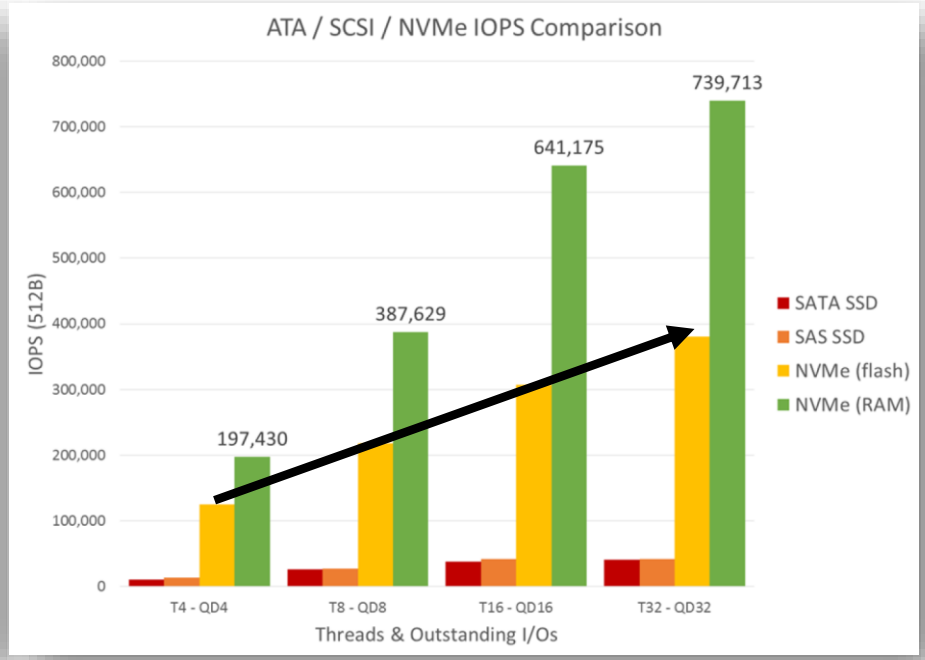

PURESTORAGE®

학습,검증 단계 이슈

수집 → 처리 → 학습 → 검증 → 활용

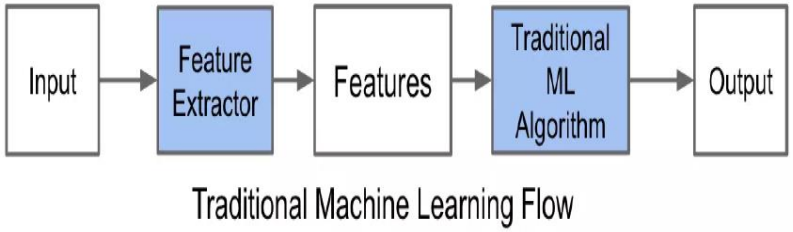

Traditional Machine Learning Flow

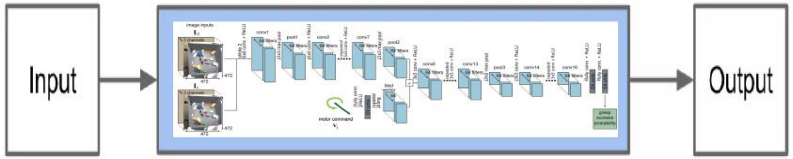

Deep Learning Flow

CPU 대용량 연산 병목

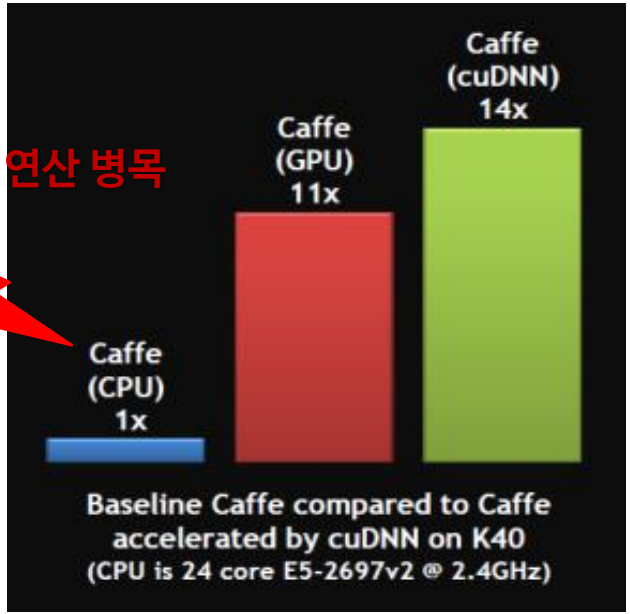

PURESTORAGE®

빠른 학습,검증을 위한 GPU처리 능력 충족

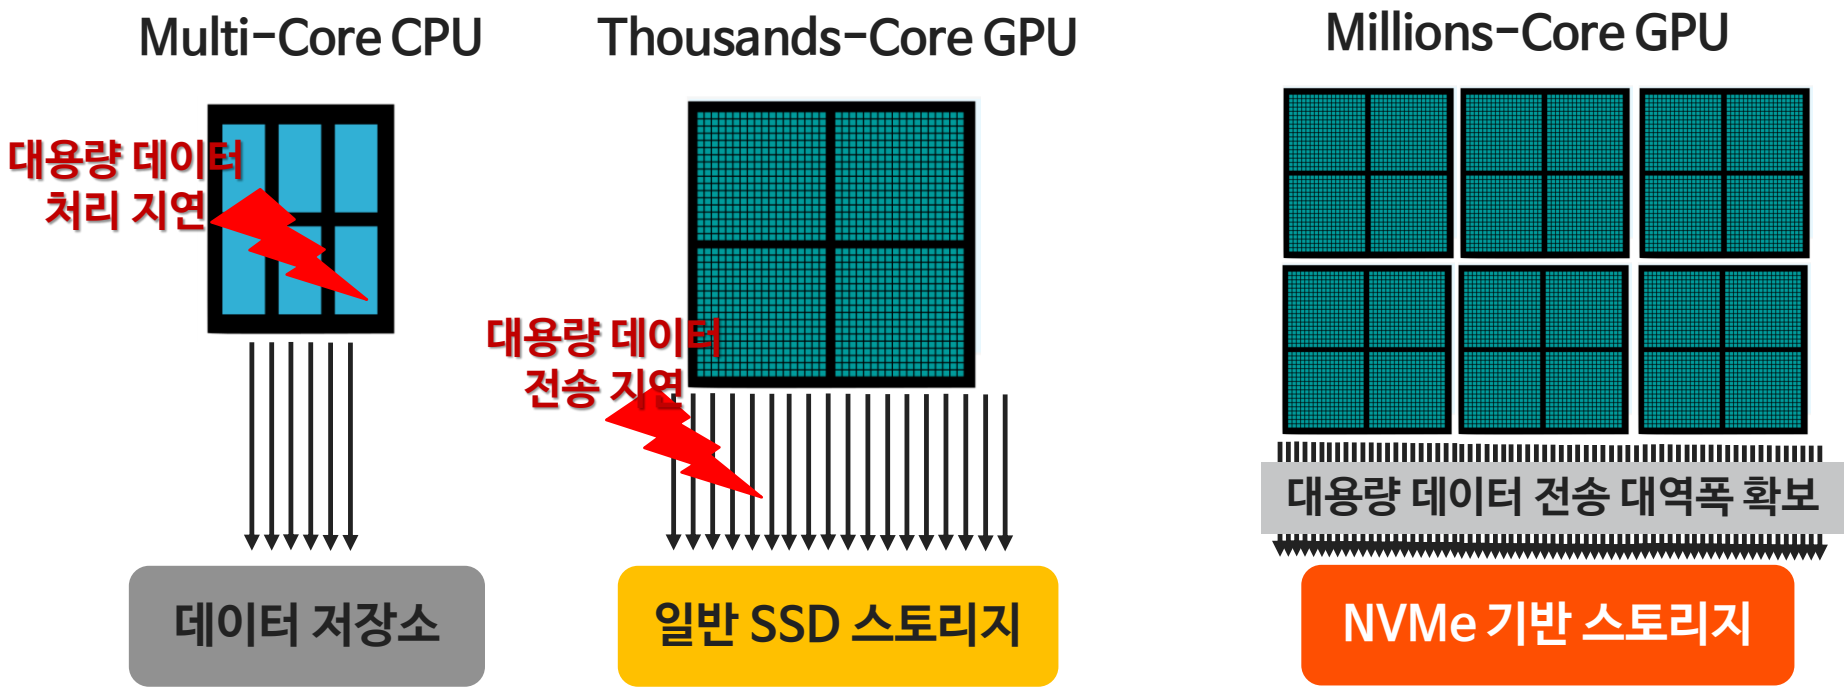

PURESTORAGE®

인공지능을 위한 차세대 데이터플랫폼은
BOTH A TRUCK AND A RACE CAR

CAPACITY, THROUGHPUT,
SEQUENTIAL ACCESS

CONCURRENCY, LATENCY,
RANDOM ACCESS

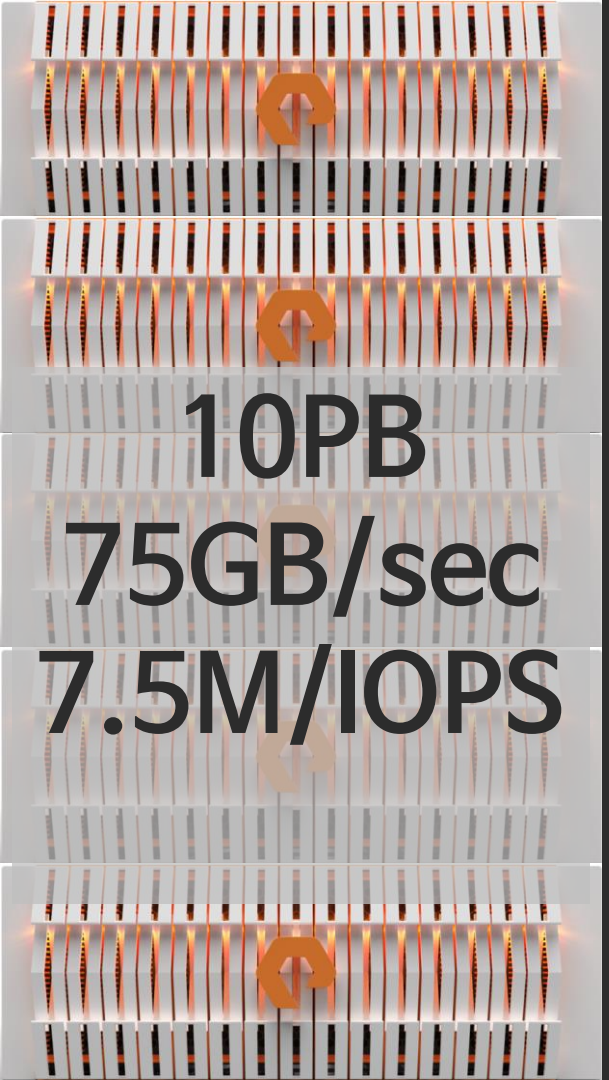

FLASHBLADE

INDUSTRY'S FIRST FLASH PURPOSE-BUILT FOR MODERT DATA ANALYTICS

FAST

Elastic Performance Up to 75 GB/s
Always-Fast, Small to Large Files
Massively Parallel from SW to Flash

BIG

10's of Thousands of Clients
10's of Billions of Objects & Files
8 Petabytes with Single IP

SIMPLE

Evergreen
No Manual Tuning
Just Add Blades for Performance

퓨어스토리지를 활용한 AI 환경 구축 사례

기존 신차개발/출시 5~6 Year

인공지능기반 자율주행 자동차 신기술 적용

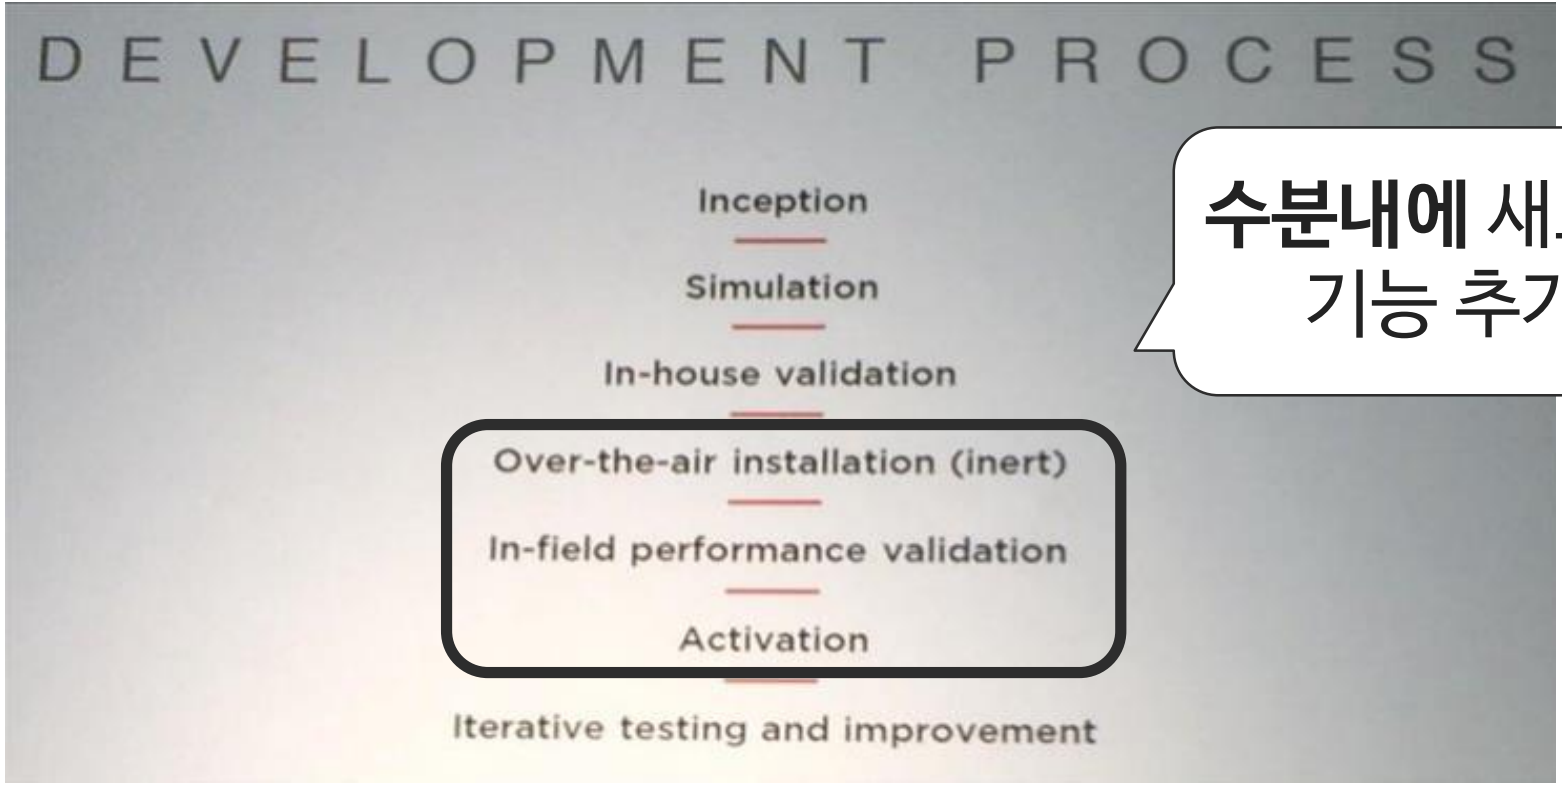

참조 – electrek.co, tesla-autopilot-billion-miles-data-self-driving-program

기능 개선화를 위한 개발 속도 요구

CAR Firmware Build

- Hundreds of SW engineers writing code
- Individual modules must be compiled together to form a full image
- Traditional Storage(NAS) systems are challenged

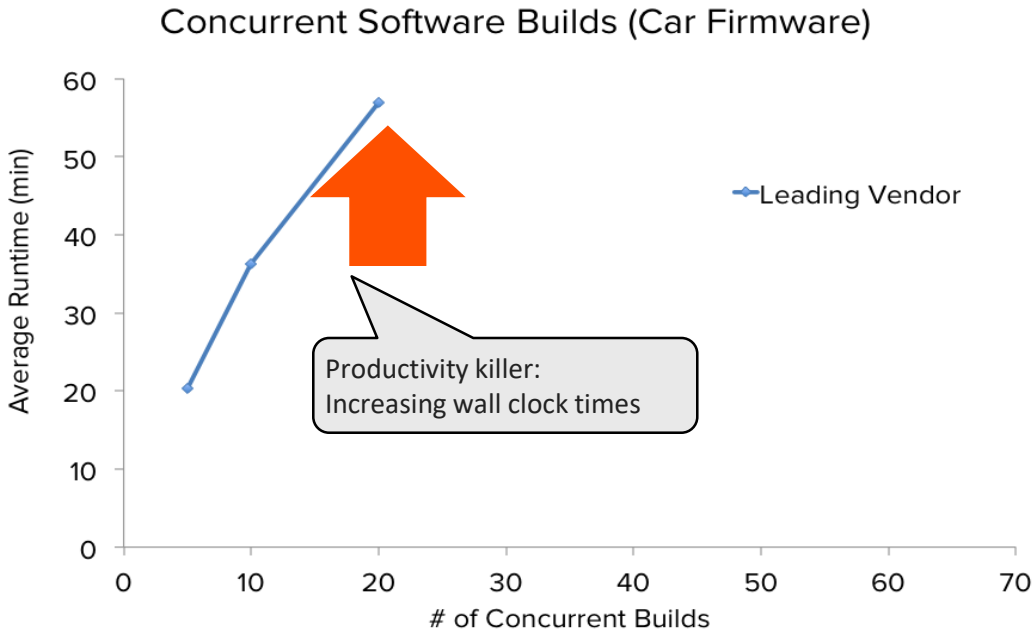

퓨어스토리지를 활용한 개발 속도 향상

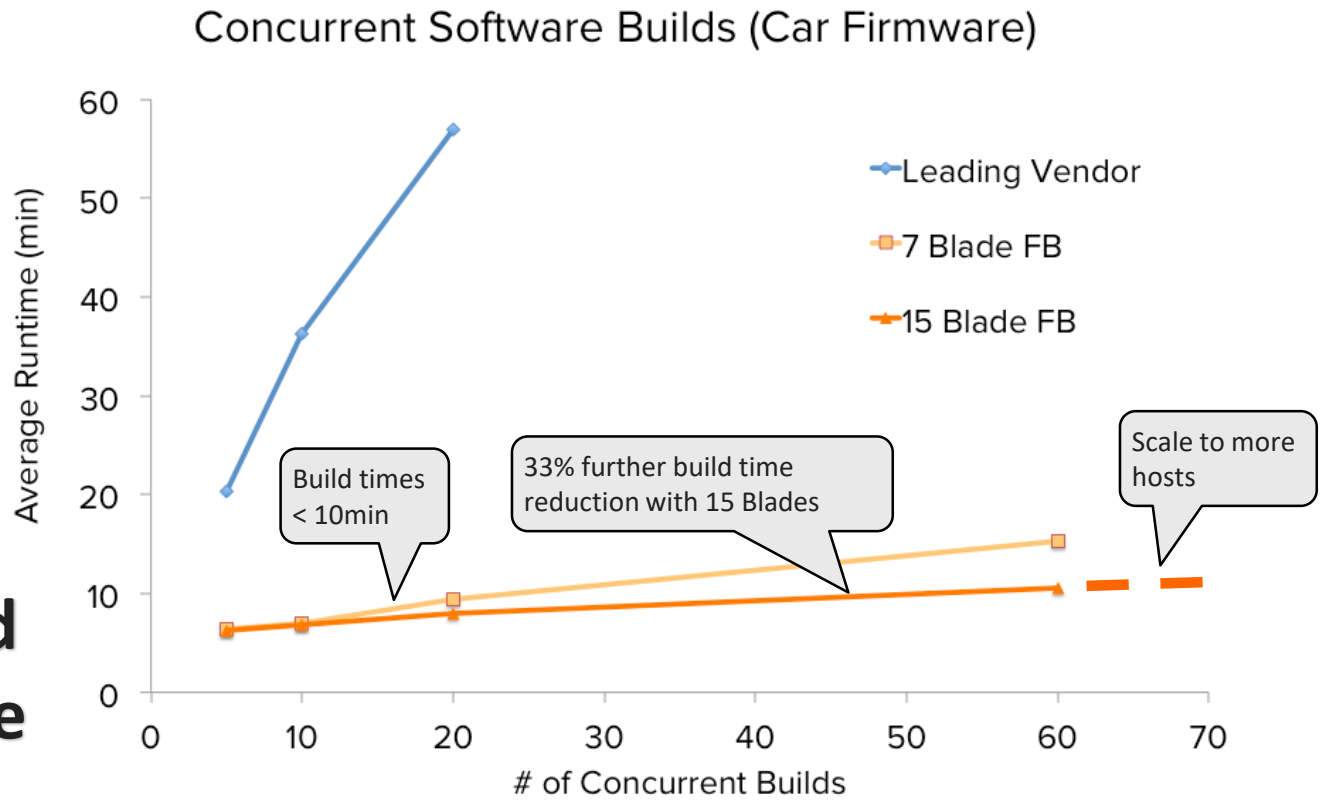

80GB/sec Read

55GB/sec Write

PURESTORAGE®

THE COMING FLOOD OF DATA IN AUTONOMOUS VEHICLES

다양한 데이터, 데이터 사일로(SILO)

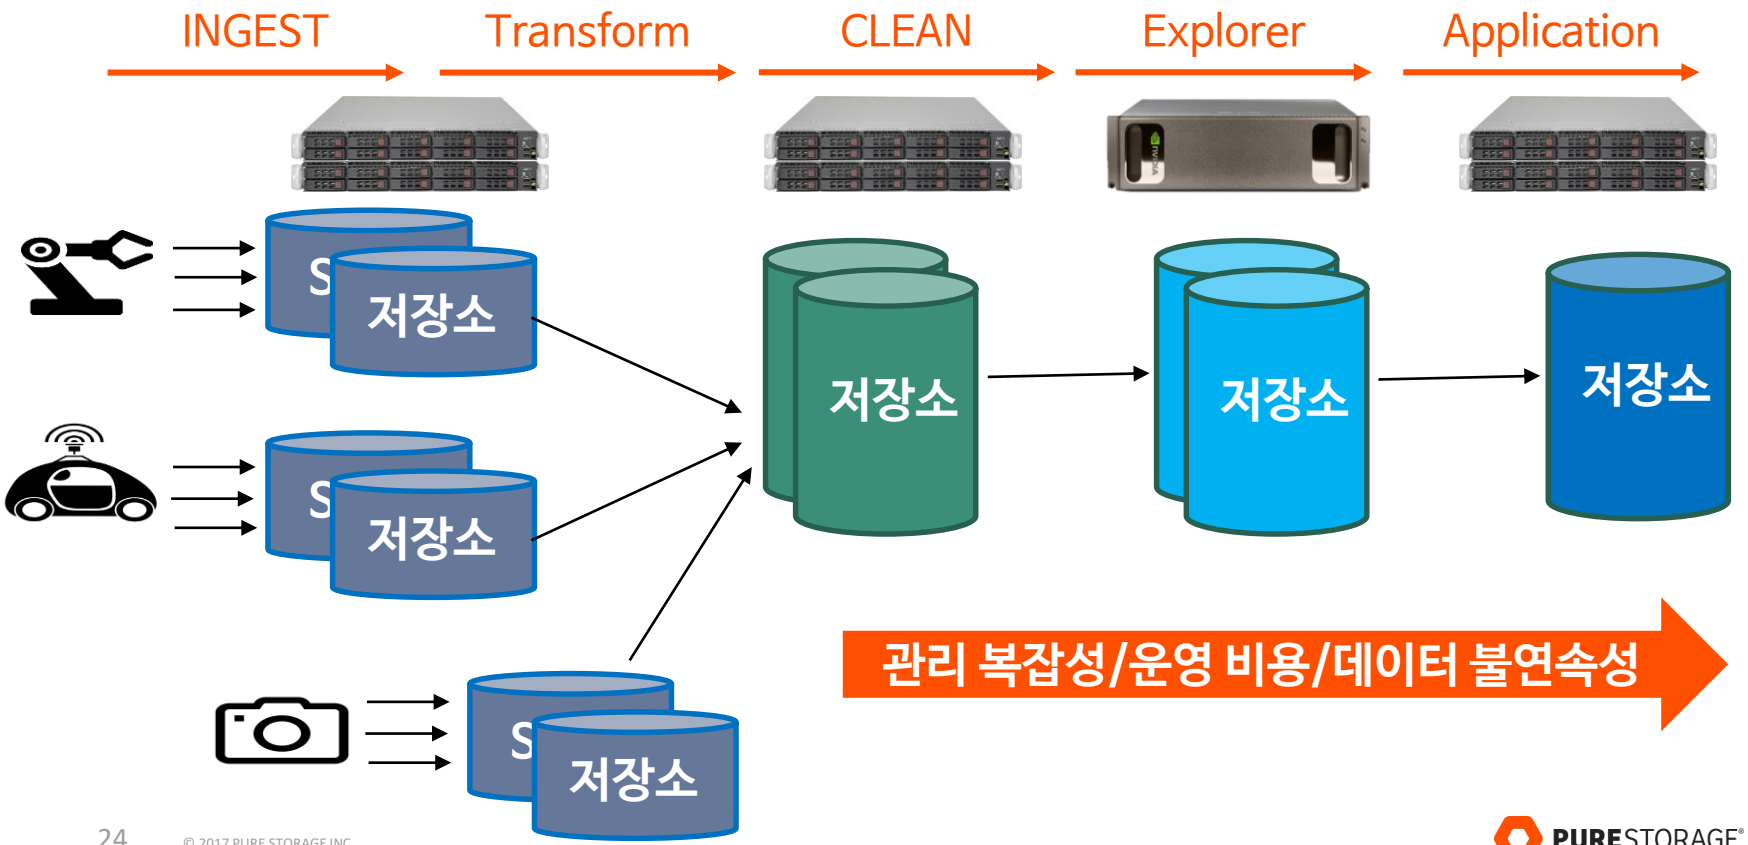

PURESTORAGE®

전문가들의 조언!!!

BUILDING A DATA PIPELINE IS CRITICAL

"Building and managing data pipelines is typically one of the most costly pieces of a complete machine learning solution."

Jeremy Hermann & Mike Del Balso
Uber Machine Learning Platform
https://eng.uber.com/michelangelo/

"If your boss asks you, tell them that I said [to] build a unified data warehouse."

Andrew Ng
Former head of Baidu AI/Google Brain
Nuts and Bolts of Applying Deep Learning

OLD WAY

12 Hours per Car

단일 데이터파이프라인

2.5 Hours per Car

TensorFlow-Base Product 플랫폼

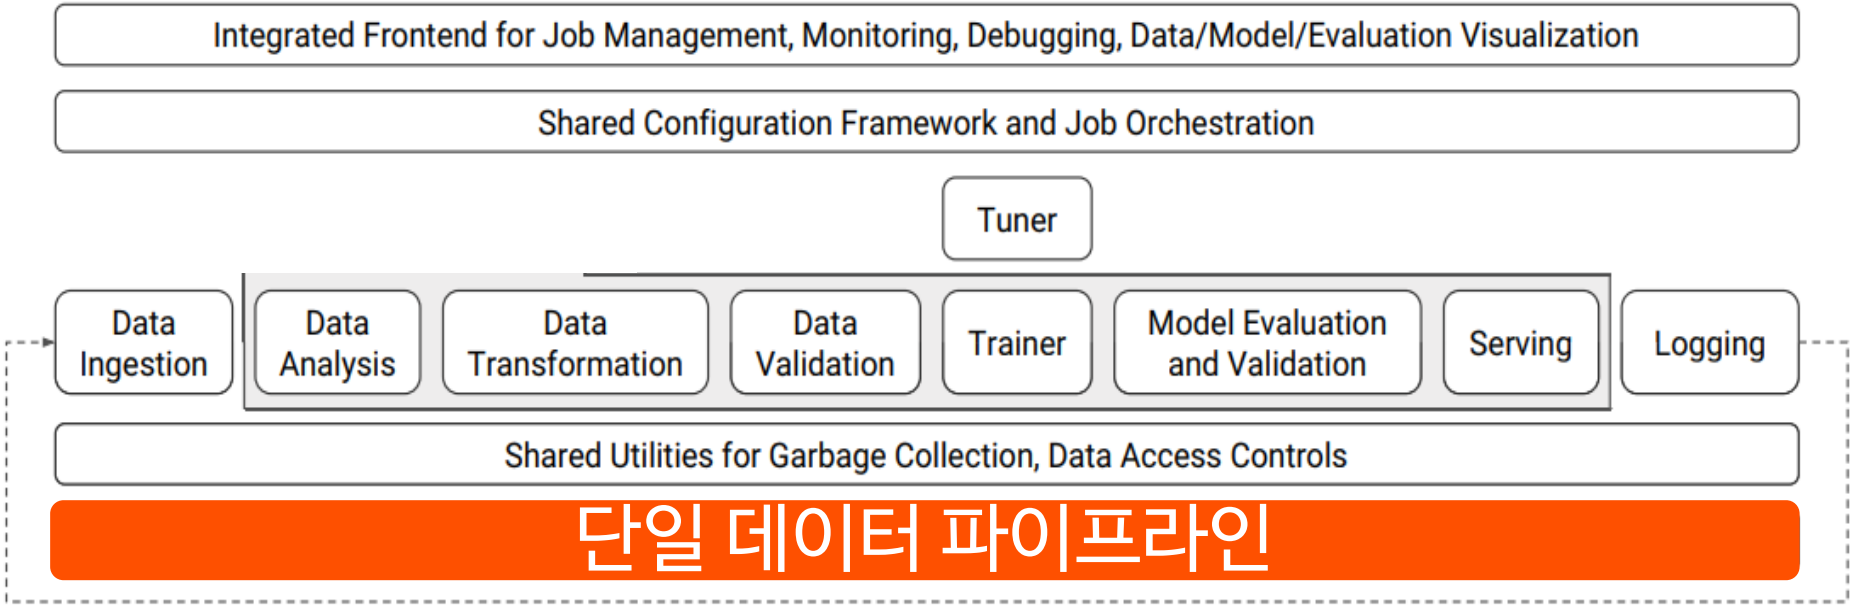

참조 - TFX: A TensorFlow-Based Production-Scale Machine Learning Platform

데이터 파이프라인 단순화

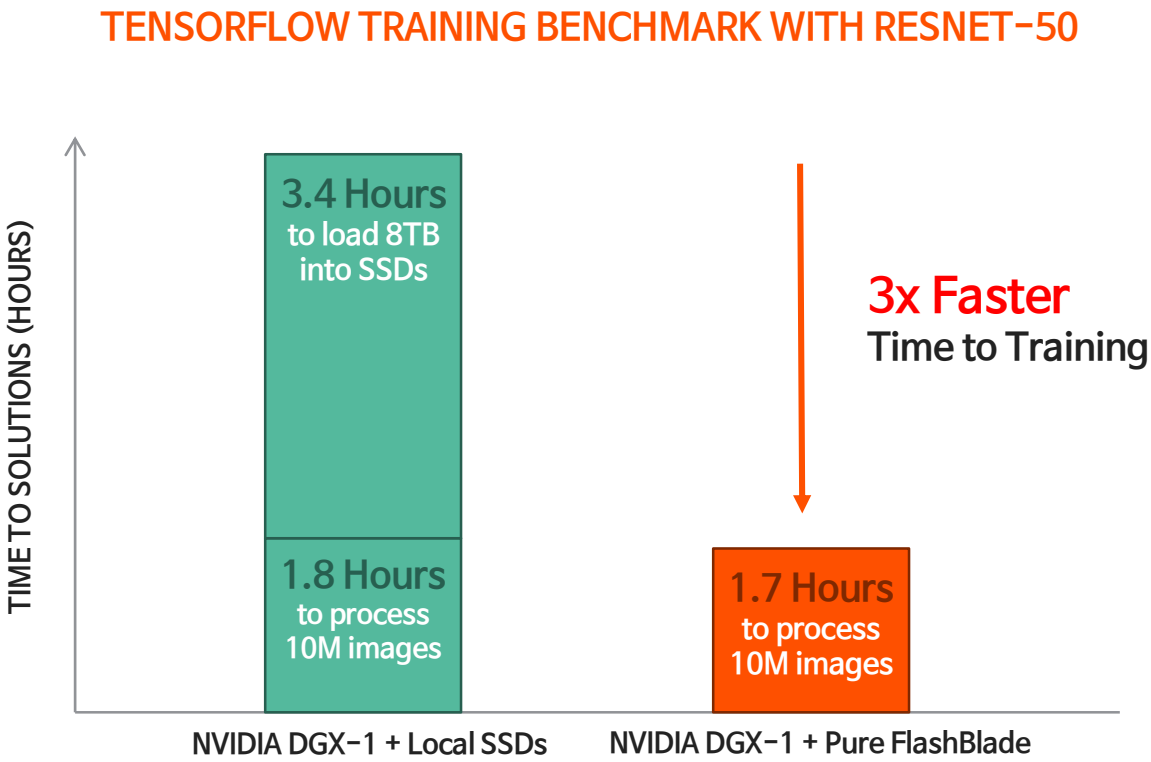


Make it real.
MAKING AUTONOMOUS CARS POSSIBLE BY 2021
Zenuity, a joint venture of Volvo and Autoliv, chose NVIDIA DGX-1 and Pure FlashBlade systems for their deep learning infrastructure.

빅데이터플랫폼, 리소스 증설 한계 발생

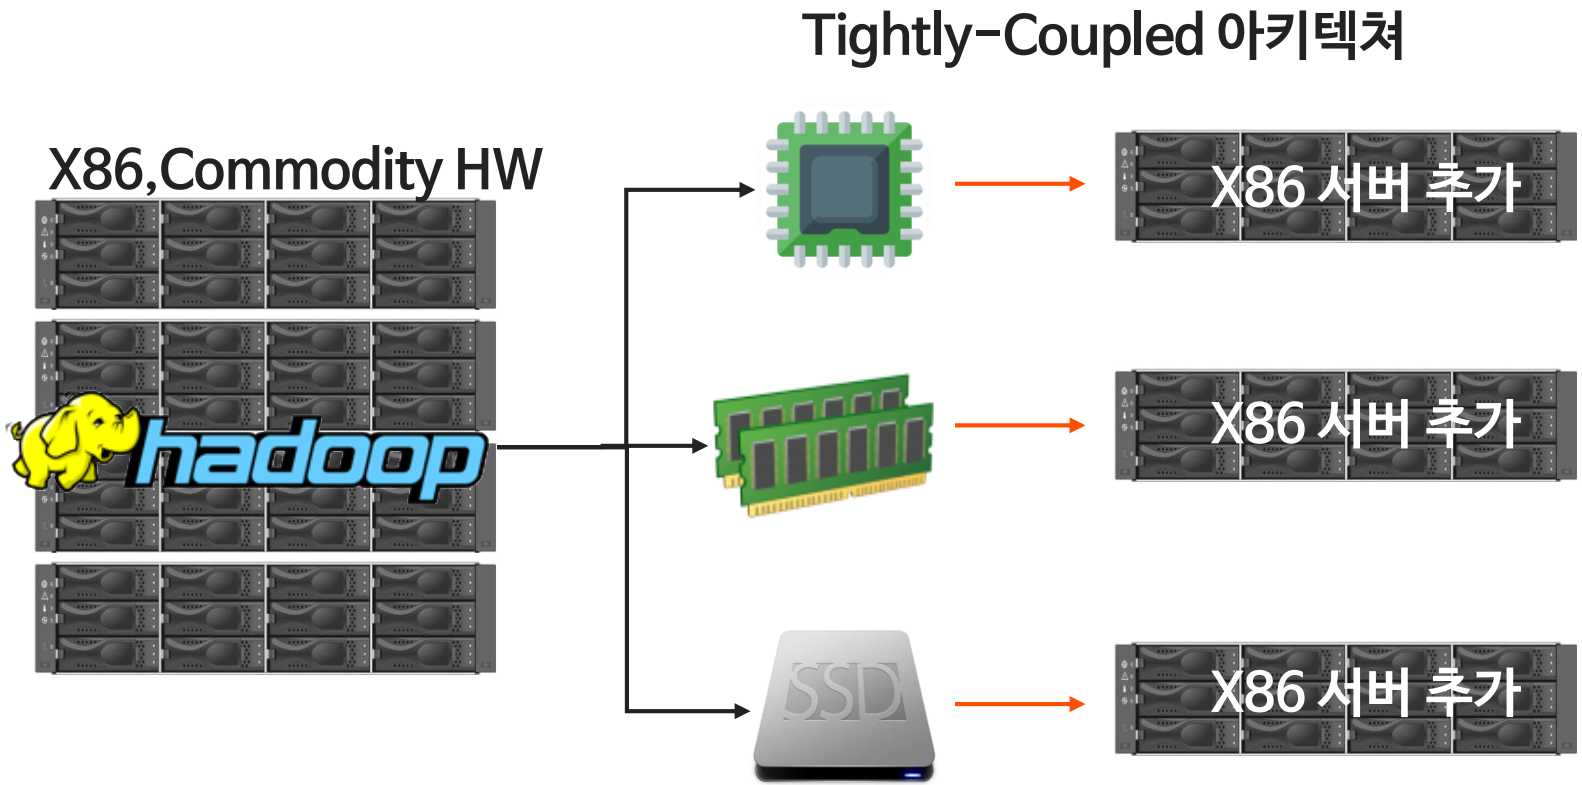

빅데이터플랫폼, 클러스터 관리 복잡

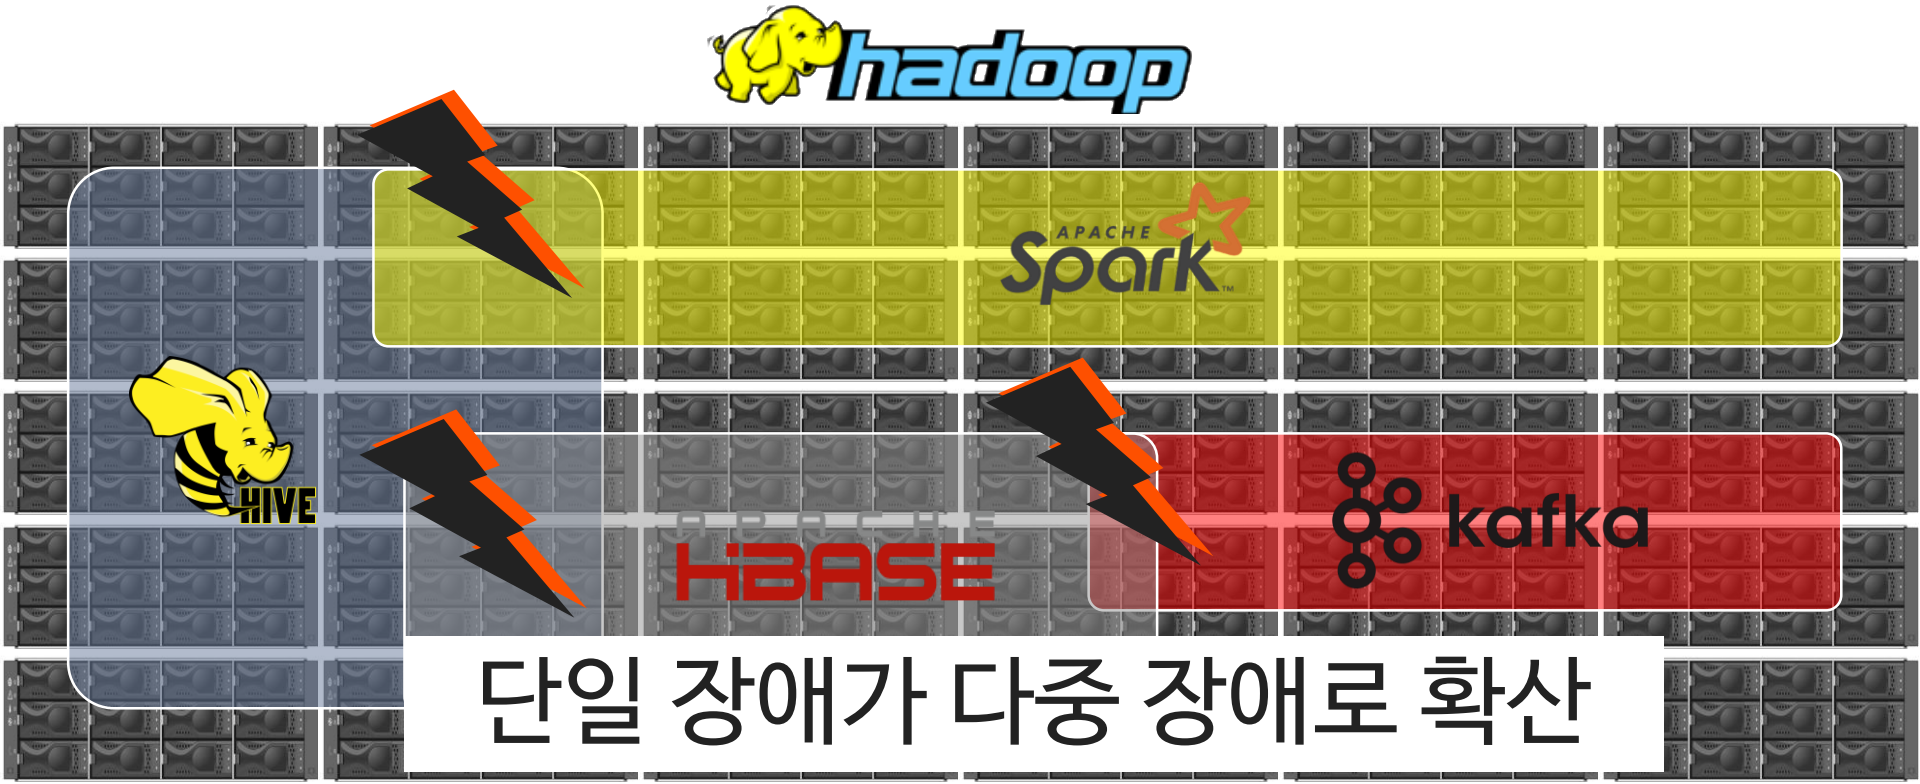

단일 장애가 다중 장애로 확산

결론, 핵심역량에 집중

PURESTORAGE®

FLASHBLADE, HDFS 성능 및 안정성 향상

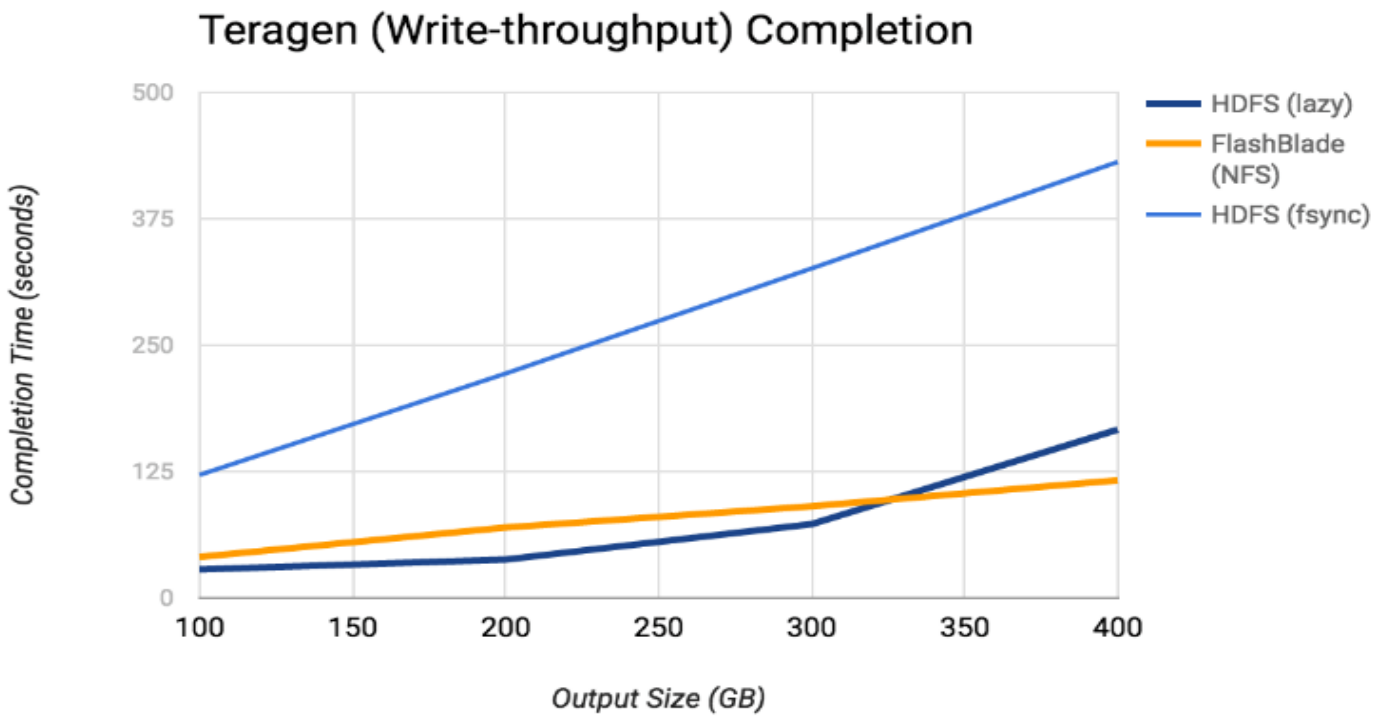

10X FASTER INVESTMENT DECISIONS WITH PURE FLASHBLADE

"Our quants want to test a model, get the results, and then test another one – all day long."

Gary Collier, co-CTO, Man AHL

람다 아키텍쳐 (실시간+배치) 분석 플랫폼

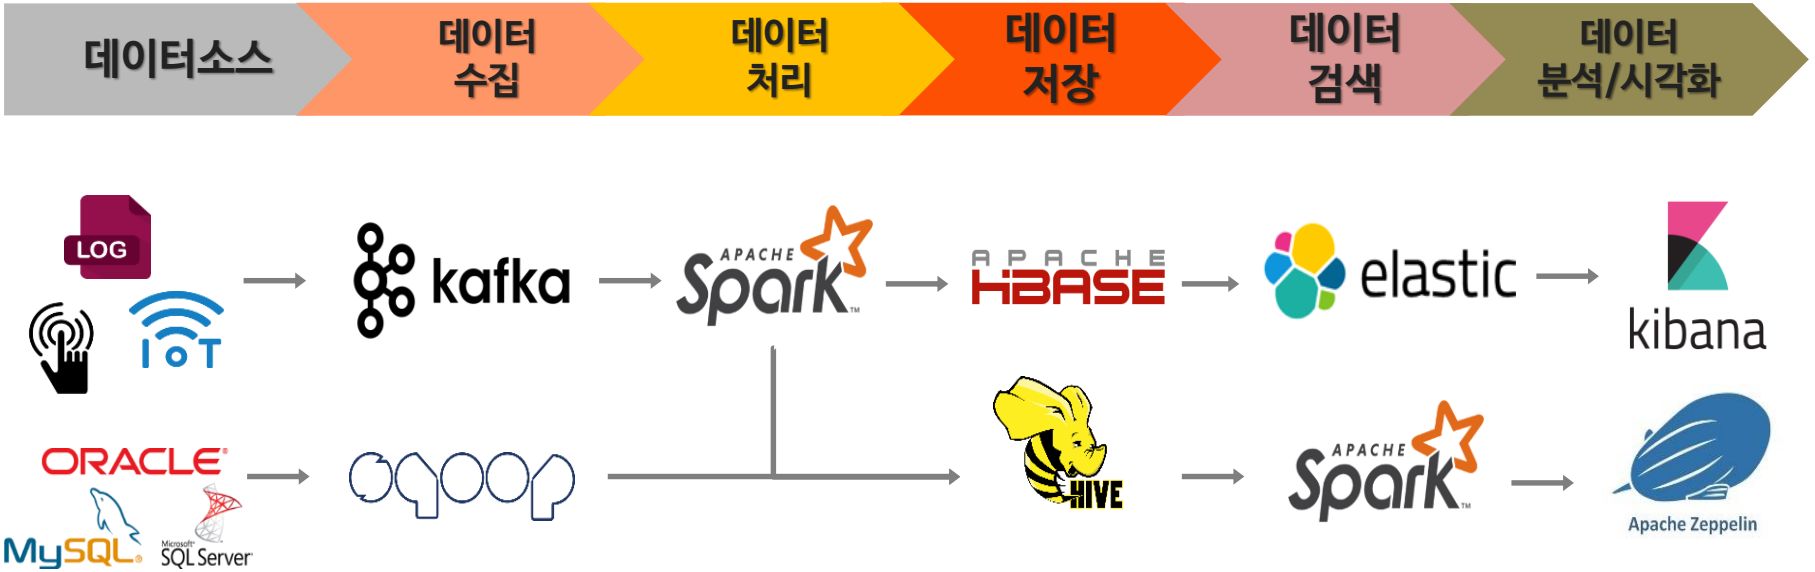

공유 파일 시스템 (NFS) 기반 고가용성, 단일 관리 HDFS구성

hadoop

EDA, HPC기반 고성능 시스템 필요

- ✓ 수 백대 혹은 수 천대 서버
- ✓ 분산 컴퓨팅
- ✓ 대용량 네트워크 대역폭
- ✓ 동시 접속 사용자
- ✓ 목적에 따른 다양한 툴

High **T**hroughput **C**omputing

High **P**erformance **C**omputing

High **C**oncurrency **C**omputing

복잡한 데이터 아키텍쳐

Lustre GPFS Ceph

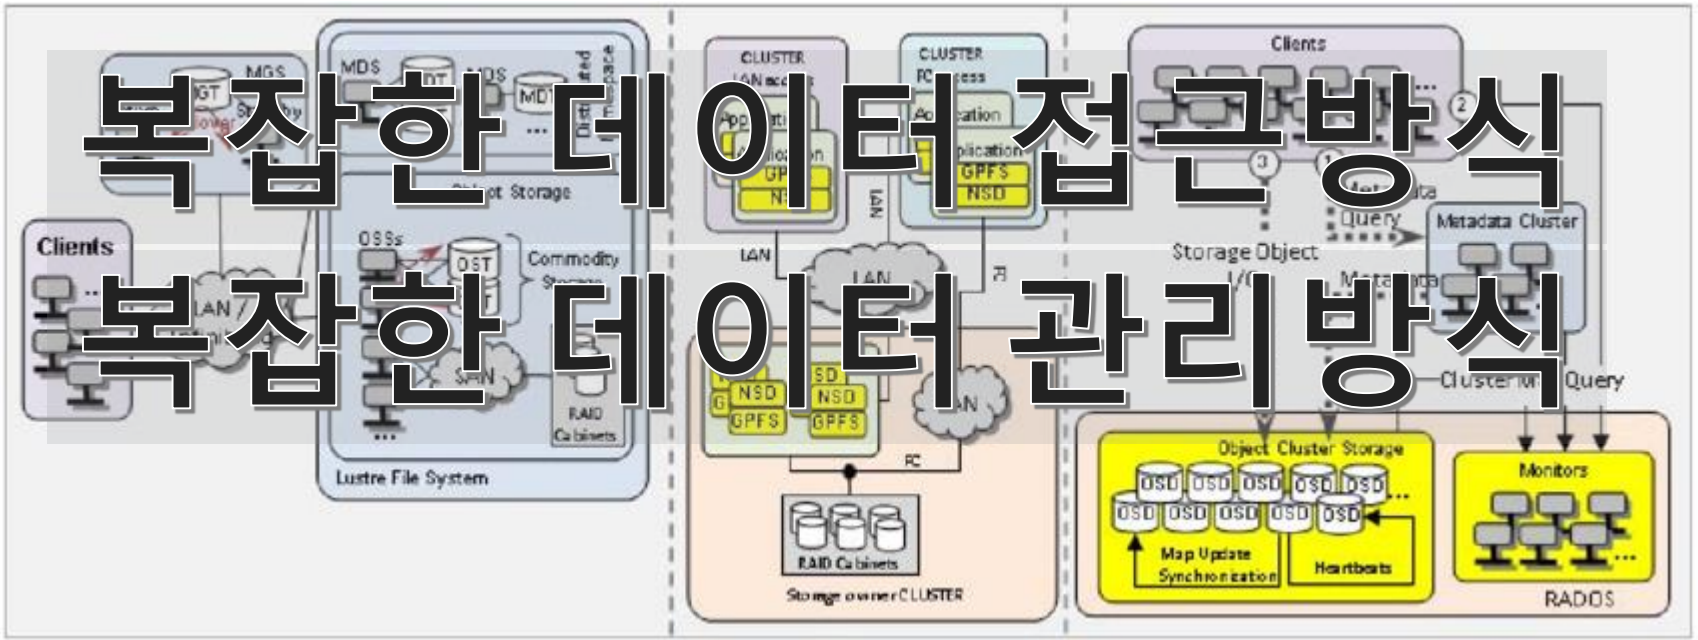

복잡한 데이터 접근방식

복잡한 데이터 관리방식

https://www.researchgate.net/figure/278847153_fig1_Fig-6-a-Lustre-Architecture-b-GPFS-Architecture-c-Ceph-Architecture

성능, 관리를 위한 데이터 접근 단순화 필요

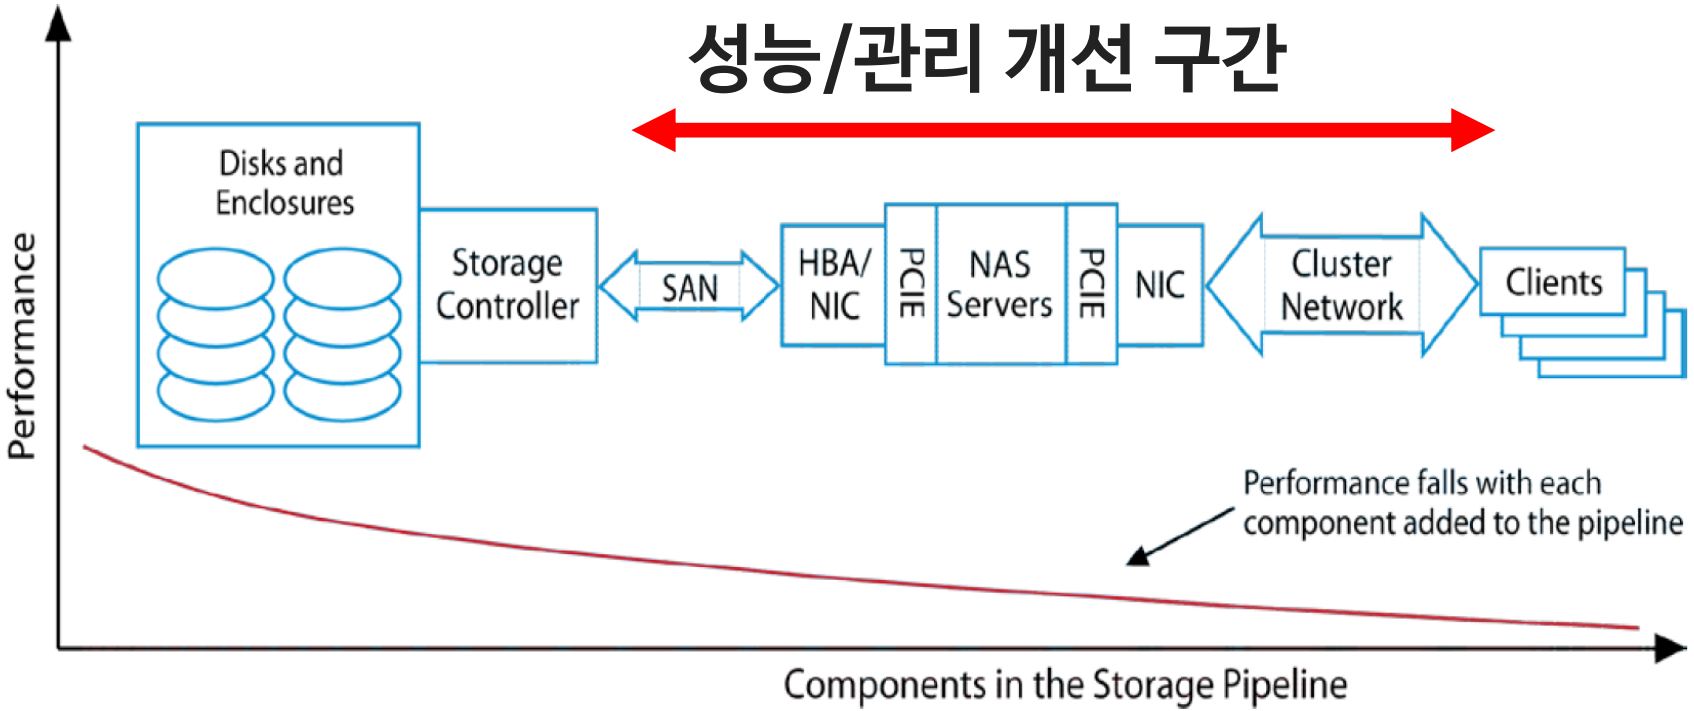

고성능 컴퓨팅 스토리지 3가지 요구 조건

1 Linear Scalability보장

- The ability to scale out in a clustered architecture to cater for increased I/O and capacity requirements.

2 Object 기반 데이터 관리

- The use of object storage technology, which organises files as objects in a flat directory rather than a traditional tree hierarchy. This boosts performance for very large collections of files because the storage system does not have to navigate a tree structure to locate a file but instead consults an index, often held entirely in memory.

3 IO처리과정 단순화

- Use of parallel file systems across multiple nodes and in some cases driven by custom silicon, which allows more rapid access to large files than when using a single node, which can become a bottleneck.


WORLD'S FASTEST
RACE CARS
DESIGNED WITH PureStorage
PURESTORAGE
QUALCOMM
PIRELLI
PETRONAS
SYNTIUM

[3.2 x 10^9 bp x 2 bit]

* http://sandwalk.blogspot.kr/2011/03/how-big-is-human-genome.html
* https://www.technologyreview.com/s/427720/bases-to-bytes/

다양한 형태의 유전자 분석 데이터를 저장

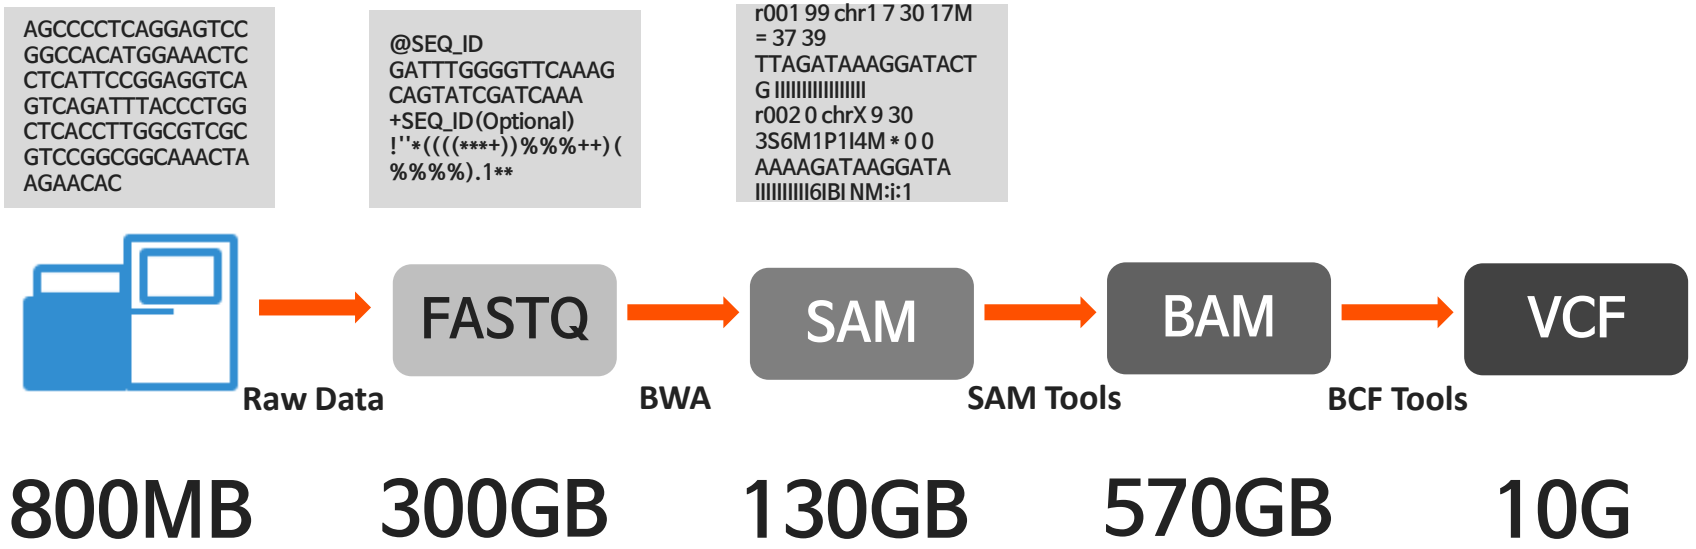

* Workload characteristics of DNA Sequence Analysis : from storage system's perspective 논문에서 발췌

고성능 유전자 데이터 처리능력 필요

FASTQ->VCF생성

Table 3: NGS Analysis Pipeline Configuration

	NGS sequence mapping		SNP Calling						
Application	bwa		samtools						bcftools
Option	aln	sampe	view | sort	index	flagstat	merge	index	mpileup	view
Input	.fastq	.sai, fastq	.sam	.bam	.bam	.bam	.bam	.merged .bam	.bcf
Output	.sai	.sam	.bam	.bam.bai	.flagstat	.bam	.bam.bai	.bcf	.vcf
Term	bwa aln	bwa sampe	samtools sort			samtools merge		samtools mpileup	

Total Data Size : 약 1100GB

Table 4: # of File Operation and Write/Delete File Size

	# of Files Created/Deleted/Opend	Write size (GByte)	Delete size (GByte)
bwa aln	14/0/2979	91.0	0
bwa sampe	7/0/1647	320.0	91.0
samtools sort	556/535/12	227.0	0
samtools merge	2/0/10	69.4	0
samtools mpileup	52/0/108	323.4	0
total	631/535/4756	1031.7	91.4

1인 분석을 위한 필요 스토리지 성능

최대 1GB/sec

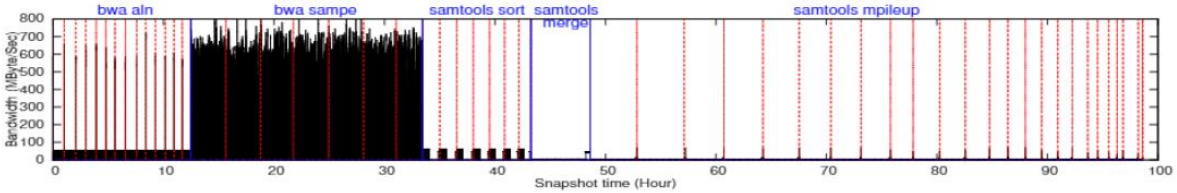

(a) Read Bandwidth

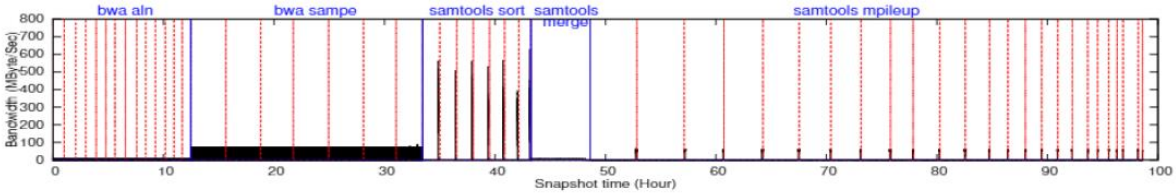

(b) Write Bandwidth

NGS 다양한 워크로드 충족

PRIMARY ANALYSIS	SECONDARY ANALYSIS	INFORMATICS
Create Base Pairs And Quality Scores	Align Reads to Assemble Sequence, Map to Reference Genome and Variant Calling	Interpretation, Discovery and Reporting
SMB/CIFS 100TB+ 100GB – 200GB	NFS 10PB+ 200GB – 500GB	NFS/Objects 100PB+ GB - TB
Medium Performance Fast writes, Medium BW Scratch Single Server FASTQ files	Very High Performance High IOPS/Metadata, Low Latency, High Aggregate BW Scratch –〉 Enterprise Large Server farms - 100K cores BAM files, VCF files	Medium Performance Med IOPs/BW Enterprise, Hybrid/Cloud Enterprise Servers Word/PDF/proprietary

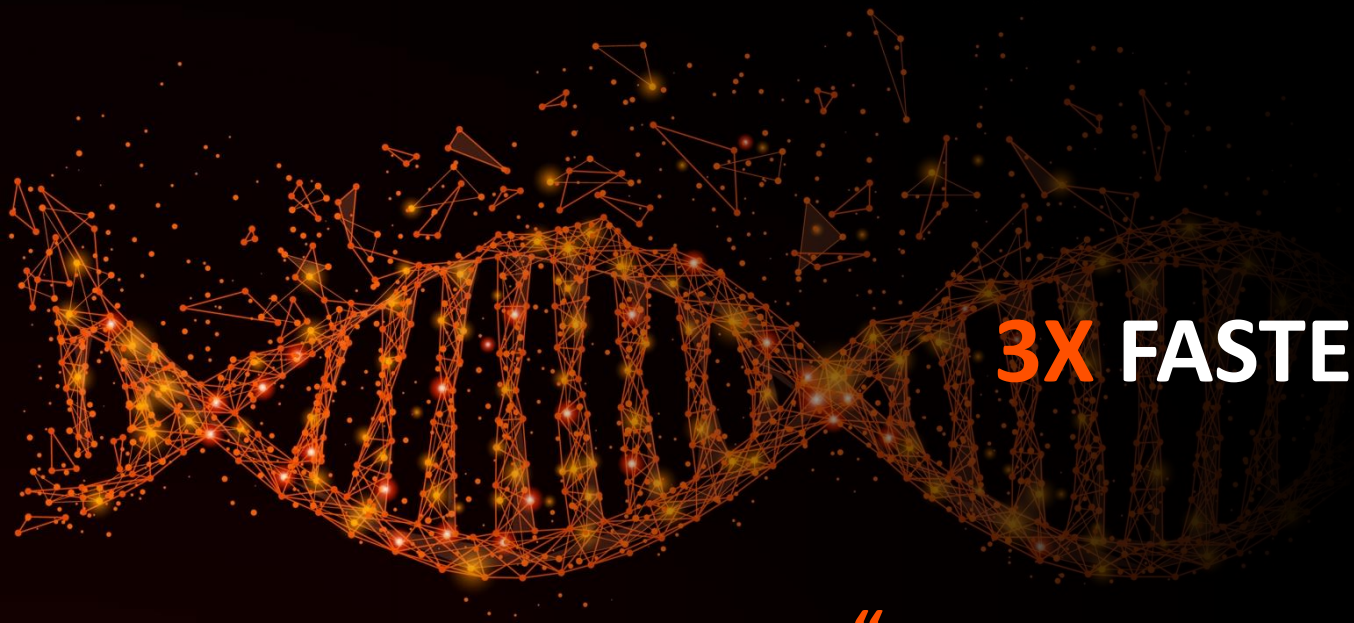

riselab
UC Berkeley

3X FASTER SPARK FOR GENOMICS

"In our HDFS system, our cluster typically runs at 10-30% utilization of compute capacity, but at 80-85% of storage capacity… trying to **increase storage capacity built out around HDFS would have been impractical**.

We've been able to take some workloads that **we couldn't run on our existing hardware, move them to FlashBlade and run them with no problem**. While that's huge for our team, it's even more important to the lives we might impact in the future."

Frank Austin Nothaft, UC Berkeley RISELab

The AIconics Awards 2017

Best Innovation in AI Hardware

WINNER

PURESTORAGE®

The AI Summit®
SAN FRANCISCO

PURESTORAGE®