



빅데이터를 위한 X86기반 오픈 인프라 전략

HP Korea

February, 2013

Legacy Data Warehouse Architecture



Alfa Romeo's stylish 8C supercar

What is Big Data? – Finding out at the Junk Yard

What is meaningful in the garbage?



What is Big Data? – Finding out more and more again...



What is Big Data?

Big Data의 속성

- 굉장히 크고 복잡한 data sets
- 기존의 DBMS나 data 처리 tool등을 이용해서는 처리하기 불가능한 data sets

Big Data를 위한 요구조건

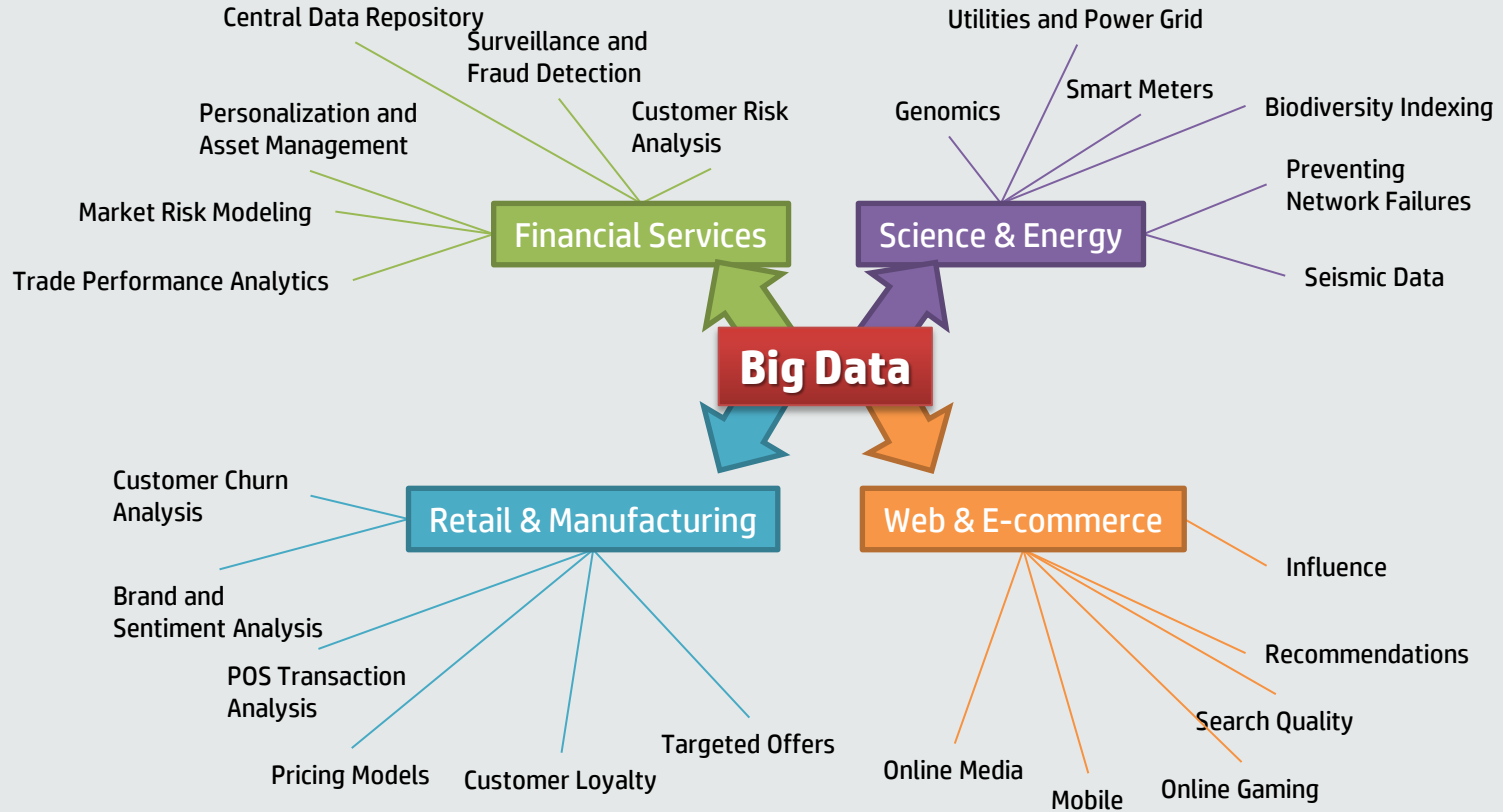
- data sets의 포착, 장기간의 보관, 저장공간, 찾기 (search), 공유(Sharing), 분석 (Analysis), 그리고 시각화 (Visualization)

Big Data가 주목 받는 이유

- 같은 크기의 data sets를 분석할 때 여러개의 작은 data sets로 나누어진 data보다 하나로 뭉쳐있는 data에서 더 많은 유익한 정보가 나오는것이 확인



Big data Use Cases for Key Industries



Why it is difficult?

Data의 폭발적 증가

Every 60 seconds



98,000+ tweets



695,000 status updates



11 million instant messages



698,445 Google searches



168 million+ emails sent



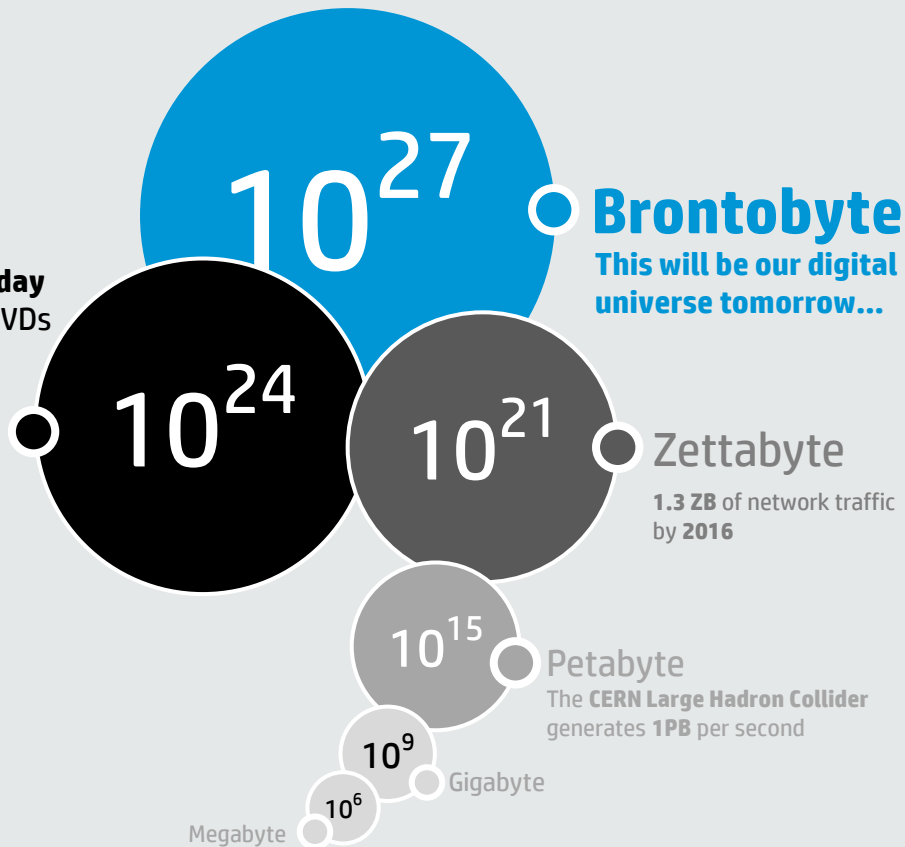
1,820TB of data created



217 new mobile web users

Yottabyte

This is our digital universe today
= 250 trillion of DVDs



1 EB of data is created on the internet each day = 250 million DVDs worth of information.
The proposed **Square Kilometer Array telescope** will generated an EB of data per day



What is Apache Hadoop?

오픈 소스 리눅스 기반의
데이터 저장공간과 처리를 위한 플랫폼

- ✓ Scalable
- ✓ Fault tolerant
- ✓ Distributed

어떤 종류의 데이터라도
저장하고 분석가능한 유연성

- 기존에 접근하기 힘들었던 정형/비정형 데이터에 대한 Query
- 단일 스키마에 종속되지 않음


복잡한 데이터를
신속하게 처리

- 복수개의 노드를 이용하여 워크로드를 분산하는 스케일 아웃 아키텍처
- ETL 병목을 제거하는 유연한 파일 시스템

경제적인
스케일 확장

- X86 기반의 상대적으로 저렴한 하드웨어를 이용하여 구축 가능
- 벤더 종속적이지 않은 오픈 소스 플랫폼으로 구성

Core HADOOP system components



Concept

- “In pioneer days they used oxen for heavy pulling, and when one ox couldn’t budge a log, we didn’t try to grow a larger ox.”
 - Grace Hopper, early advocate of distributed computing



Scale-up

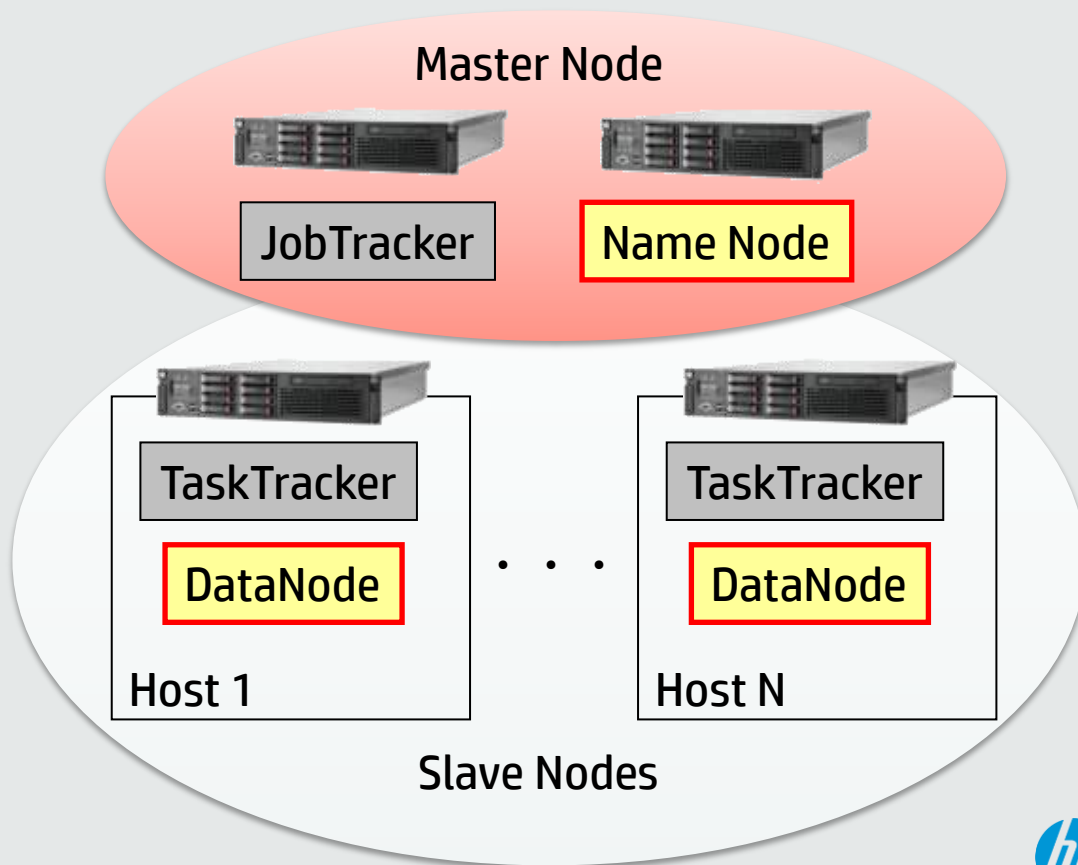


Scale-out

Hadoop Core

전형적인 Master & Slave 구조

- HDFS & MapReduce
- **HDFS**
Name Node + Data Node
- **MapReduce**
Job Tracker + Task Tracker



The Reality of Hadoop

파일럿에서 실제 프로덕션까지 진행하는 데 있어서 많은 시행착오 발생

Hadoop 은 복잡하다 ?

- **Deployment & Management**

- 플랫폼과 애플리케이션의 복잡성
- 생각만큼 쉽지 않은 사이징, 그리고 구현

- **Support is limited**

- 지원은 아직 초기 단계

- **Data integration is challenging**

- 기존 데이터를 Hadoop 과 연계함에 있어 따르는 제약

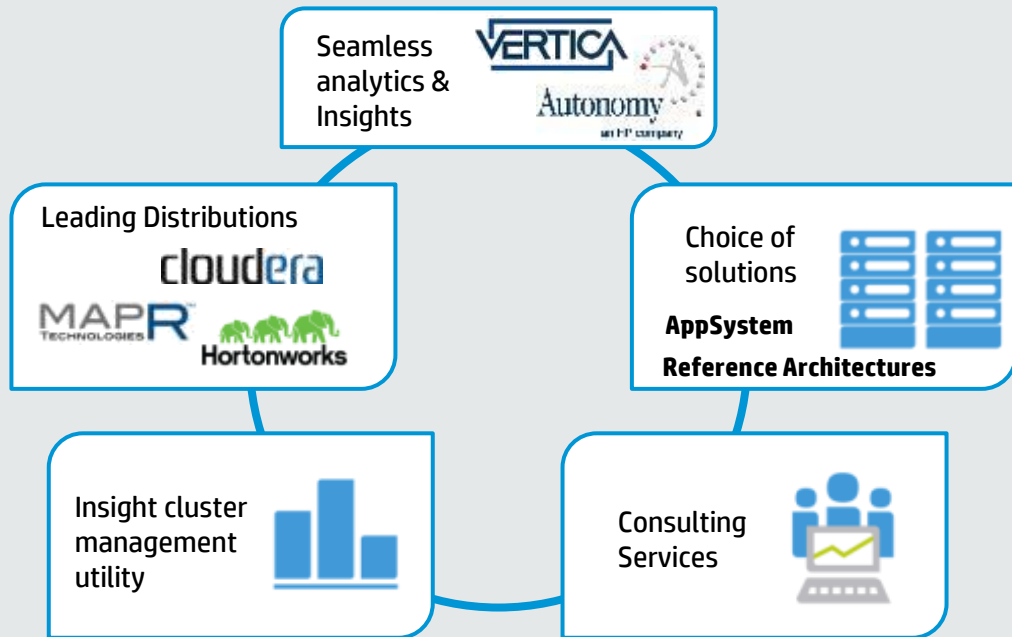


HP offers shortest route to Hadoop success

HP 는 “Best of Breed” 솔루션과 서비스를 제공

Benefit

- 기술 습득의 용이
- 신속한 구축
- 단순화 된 관리
- 리스크를 최소화하는 확장성
- **Hadoop** 에 단점을 보완하는 분석 솔루션



“We just want to use Hadoop and we are spending way too much time on Hadoop operations”

“우린 단지 hadoop 을 쓰고 싶을 뿐입니다.
하지만, hadoop 관리를 위해 너무 많은 시간이
필요합니다.”

Vice President, Financial Services Organization



HP Hadoop solutions

Hadoop 솔루션에 적합한 마켓리딩 시스템과 서비스, 그리고 HP만의 차별화된 기술력

AppSystem

HP AppSystem for Apache™ Hadoop™

- Enterprise Ready
- Easy to Deploy & Manage
- Turnkey Solution

Reference Architecture

cloudera

MAPR
TECHNOLOGIES

Hortonworks

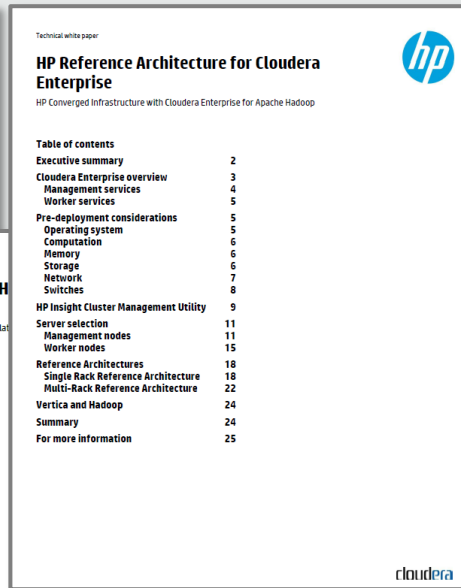

<http://www.hp.com/go/hadoop>



HP Reference Architecture for Hadoop Series

Hadoop 구성을 위한 HP의 표준 아키텍처

- HP 와 alliance 를 맺은 3대 Hadoop 배포판 벤더와 협력하여 검증된 표준 아키텍처
- 기본 개념부터 실제 장비 구성에 대한 파트별 상세 가이드, 그리고 BOM 까지 제공




Source : HP Solutions for Apache™ Hadoop™

<http://h71028.www7.hp.com/enterprise/us/en/partners/hadoop.html>



HP AppSystem for Apache Hadoop

빅데이터 기술의 능력을 극대화시킬 수 있는 엔터프라이즈급 Hadoop 플랫폼



- Enterprise ready Hadoop platform – 배송과 함께 즉시 사용 가능
- HP Cluster Manager 를 통한 보다 간편한 구축 및 확장
- 실시간 성능 데이터를 시각화하여 보다 쉬워진 관리
- 더 빨라진 적재, 분류, 분석 가능 - 10TB at 120GB/min

1

**Push-button
deployment**

800

**Nodes in minutes
not months**

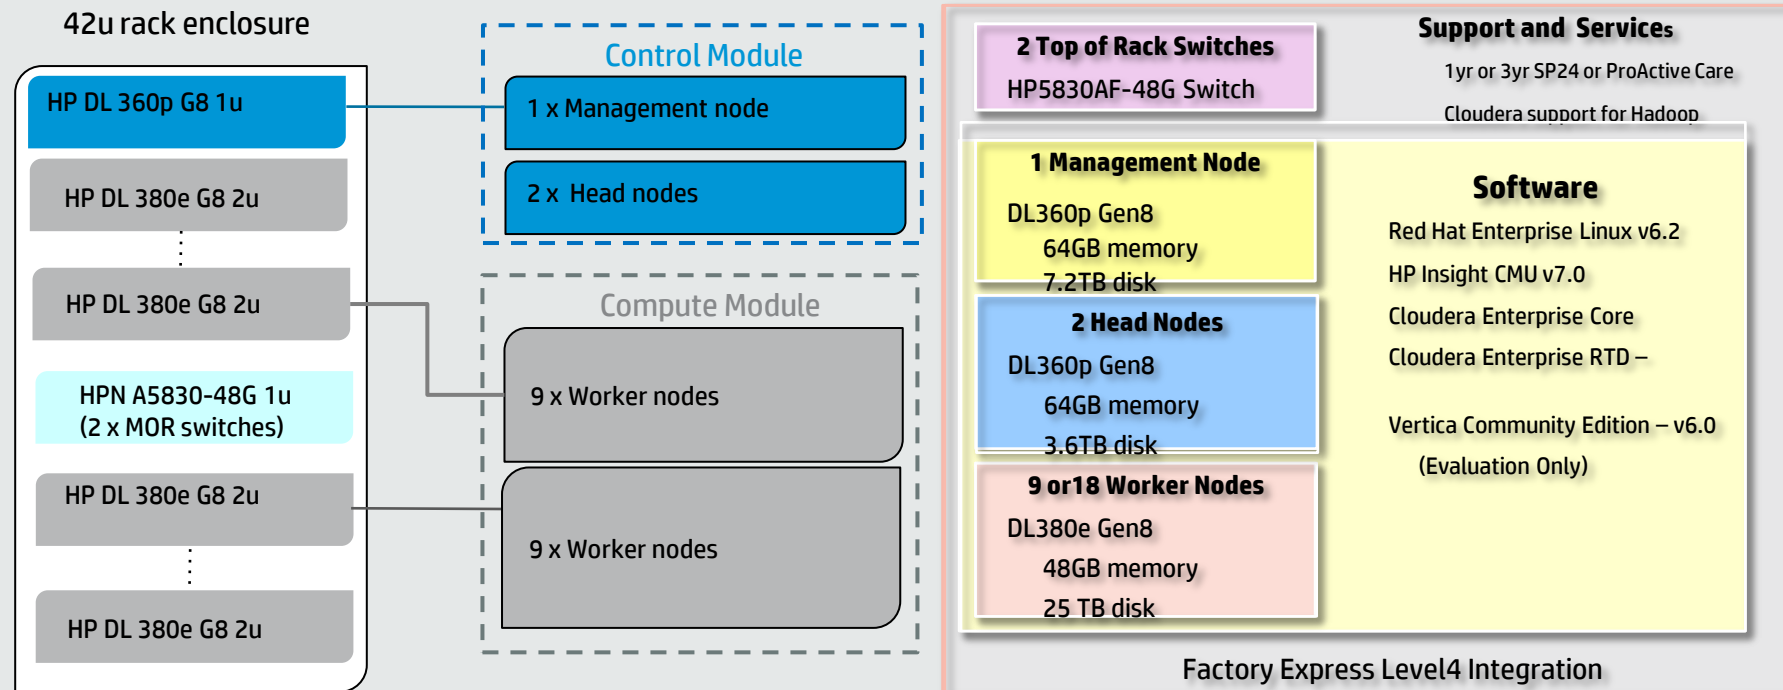
3.8x

Times Faster



HP AppSystem for Apache Hadoop

기본 2가지 구성 제공 – full rack (21노드, HDFS 384TB)와 half rack (12노드, 192TB)



“Hadoop guys recommend commodity hardware for Hadoop cluster,
but they recommend carrier-grade server for Master node.”

“하둡 클러스터는 일반적인 하드웨어로도 충분합니다.
하지만, 마스터 노드를 위해서는 더 높은 가용성의 장비가
필요합니다.”



Component for HP AppSystem for Hadoop

전체 시스템 관리용 서버 및 Hadoop 전용 관리 서버

Management Node (r01mgt) HP ProLiant DL360p Gen8



- CPU : 2 x E5-2667 2.9 GHz (HT enabled)
- HDD & Array controller : 8 x 900GB SFF SAS 10K & P420i
- Memory : 64 GB (8 x HP 8GB DDR3)
- I/O : 4 x 1GbE FlexLOM NICs (Dual-Homed Bonded Pairs)
- iLO4 Management Engine

Head Nodes (r01hn01, r01hn02) HP ProLiant DL360p Gen8



- CPU : 2 x E5-2667 2.9 GHz (HT enabled)
- HDD & Array controller : 4 x 900GB SFF SAS 10K & P420i
- Memory : 64 GB (8 x HP 8GB DDR3)
- I/O : 4 x 1GbE FlexLOM NICs (Dual-Homed Bonded Pairs)
- iLO4 Management Engine

Component for HP AppSystem for Hadoop

실제 데이터 저장 및 처리 서버 및 연결 스위치

Worker Nodes (r01wn01 – r01wn18)

HP ProLiant DL380e Gen8



- CPU : 2 x E5-2440 2.4 GHz (HT enabled)
- Array controller : P420 Smart Array Controller
- HDD : 12 x 2TB 7.2K LFF SATA MDL (Data)
2 x 500GB SATA 7.2K LFF MDL (Mirrored OS and Runtime)
- Memory : 48 GB (6 x HP 8GB DDR3)
- I/O : 4 x 1GbE NICs
- iLO4 Management Engine

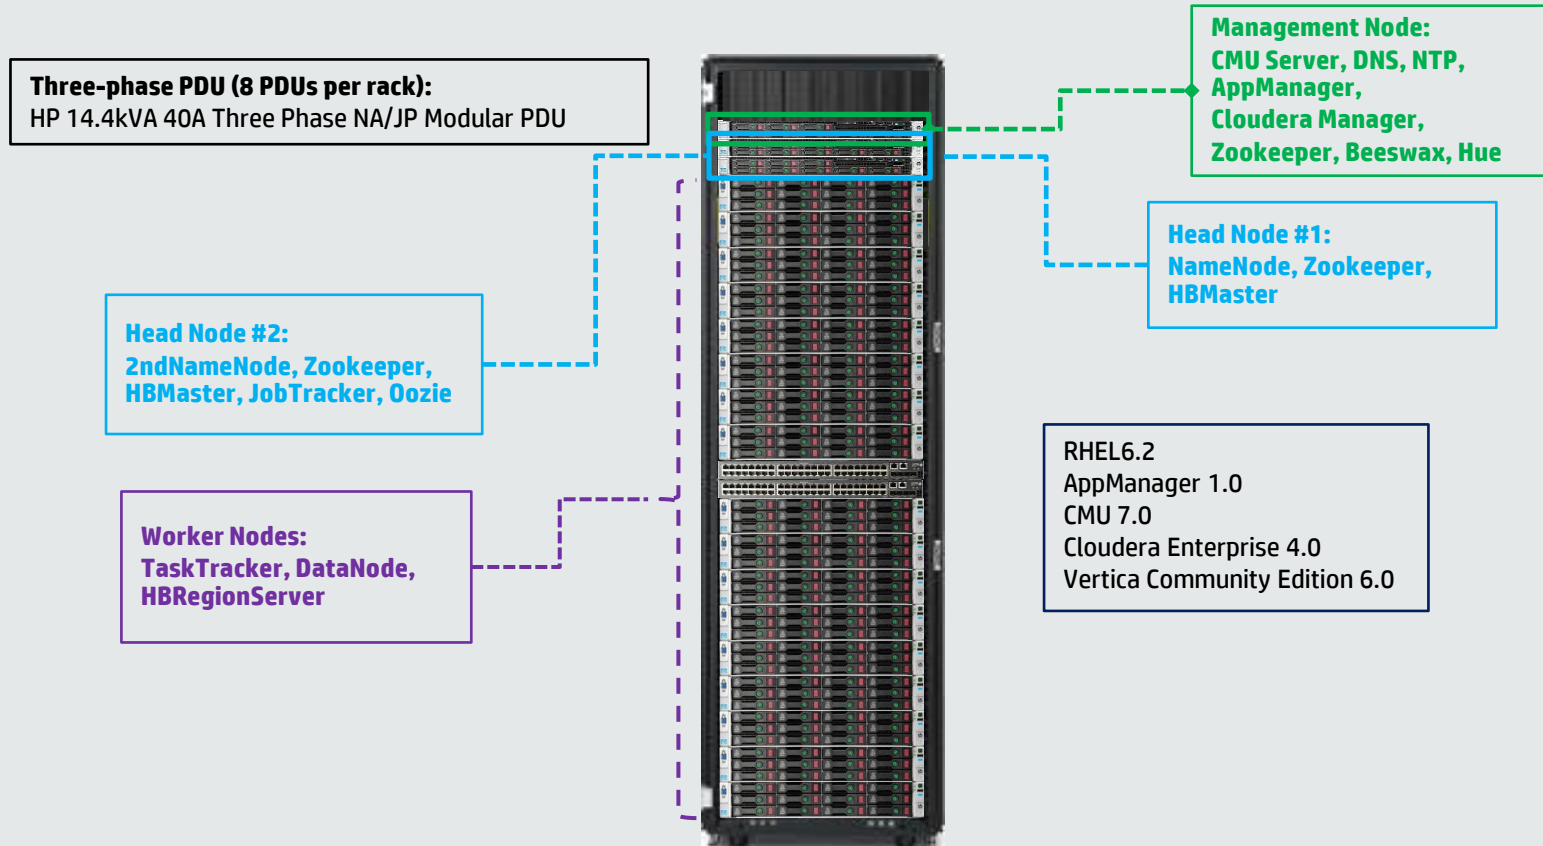
Network Switch

2 x HP 5830AF-48G



- 1GB Packet Buffer
- 48 x1GbE Ports and 4 x 10GbE Ports
- 2 x 10GbE ports used for IRF Bonding, 2 reserved for rack expansion
- Two Vlans:
 - Vlan1: Private network 172.24.0.0/15
 - Vlan2: Enterprise Customer Network (ECN), reconfigurable

Summary: HP AppSystem for Apache Hadoop



Management Solution

Deployment 와 Configuration 을 위한 도구 제공

Cloudera: Hadoop 관리와 모니터링

- Cloudera Manager – Deployment and Monitor

HP: 인프라 관리와 모니터링

- Insight CMU – Provision, Monitor, Control
- AppManager – Compliance Check
- HP CT – Firmware & BIOS level Check



Cloudera Manager 4

The screenshot displays the Cloudera Manager 4 interface. At the top, there is a navigation bar with tabs for Services, Hosts, Activities, Logs, Events, and Reports. A search bar and user profile (admin) are also visible. Below the navigation bar, a timeline shows activity from Oct 18 3:00 AM to Oct 18 4:20 AM. The main content area is titled "Services (Current)" and includes an "Add Cluster..." button. It is divided into two sections: "HP Hadoop - CDH4" and "Cloudera Managed Services".

HP Hadoop - CDH4

[Add a Service](#) Actions

Name	Type	Status	Health	Role Counts	Actions
hbase1	HBase	✓ Started	✓ Good	9 Region Servers, 2 Masters, 1 Gateway	Actions
hdfs1	HDFS	✓ Started	✓ Good	1 SecondaryNameNode, 3 HttpFS, 1 NameNode, 1 Balancer, 9 DataNodes, 1 Gateway	Actions
hue1	Hue	✓ Started	✓ Good	1 Beeswax Server, 1 Hue Server	Actions
mapreduce1	MapReduce	✓ Started	✓ Good	1 JobTracker, 9 TaskTrackers, 1 Gateway	Actions
oozie1	Oozie	✓ Started	✓ Good	1 Oozie Server	Actions
zookeeper1	ZooKeeper	✓ Started	✓ Good	3 Servers	Actions

Cloudera Managed Services

Name	Type	Status	Health	Role Counts	Actions
mgmt1	Cloudera Management Services	✓ Started	✓ Good	1 Event Server, 1 Host Monitor, 1 Activity Monitor, 1 Reports Manager, 1 Alert Publisher, 1 Service Monitor	Actions



Management and Monitoring – Cloudera Manager

The screenshot displays the Cloudera Manager 4 interface for the 'hdfs' service. The top navigation bar includes 'Services', 'Hosts', 'Activities', 'Logs', 'Events', and 'Reports'. The main content area is divided into several sections:

- Service hdfs:** Shows the service status as 'Was Running with Good Health (Oct 5, 04:30:59 PM)'. Navigation tabs include 'Status', 'Instances', 'Commands', 'Configuration', and 'History'.
- HDFS Summary (Oct 5, 4:30:59 PM):** A table showing throughput and latency for read and write operations, and a capacity bar chart.
- Status and Health Summary (Oct 5, 4:30:59 PM):** A table listing the status and health of NameNode, Secondary NameNode, DataNode, and Balancer.
- Health Tests (Oct 5, 4:30:59 PM):** A list of health tests with their results.
- Charts (Oct 5, 1:57:17 PM - 5:05:41 PM):** A line chart titled 'HDFS Canary Duration' showing the duration of canary tests over time.
- Events (Oct 5, 1:57:17 PM - 5:05:41 PM):** A search interface for events, showing 56 results.

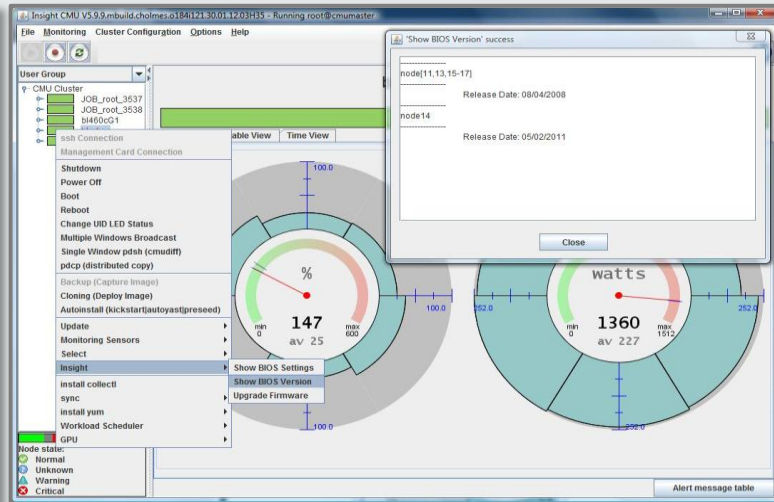
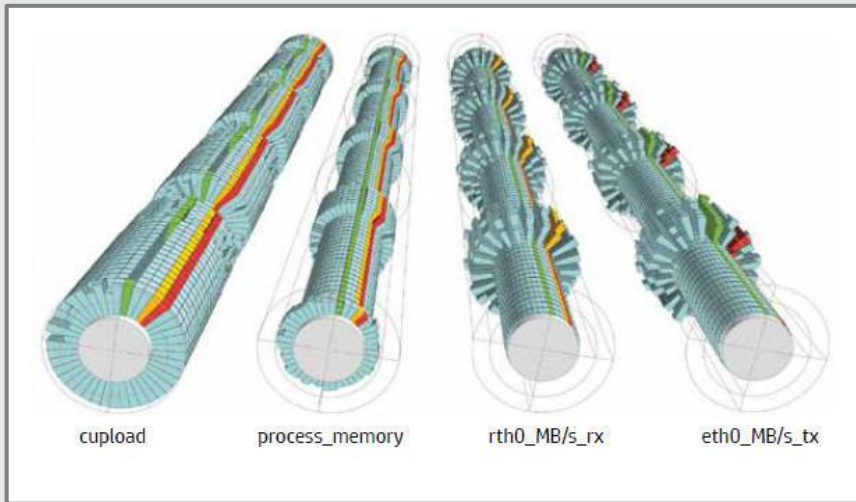
A blue banner is overlaid on the chart area with the text: **Cloudera Manager 4 – View Service Health and Performance**

Event Type	Event Code	Start Time	Host	Service	Role	Message
Log Message	EV_LOC_EVENT	Oct 5, 02:18:48 PM	c0639-hal.cloudera.com	hdfs	st_nameNode	DIR* NameSystem.startFile: failed to create file /user/hudson/benchmarks/TestDFSIO/lo_data/test_io_85 for DFS



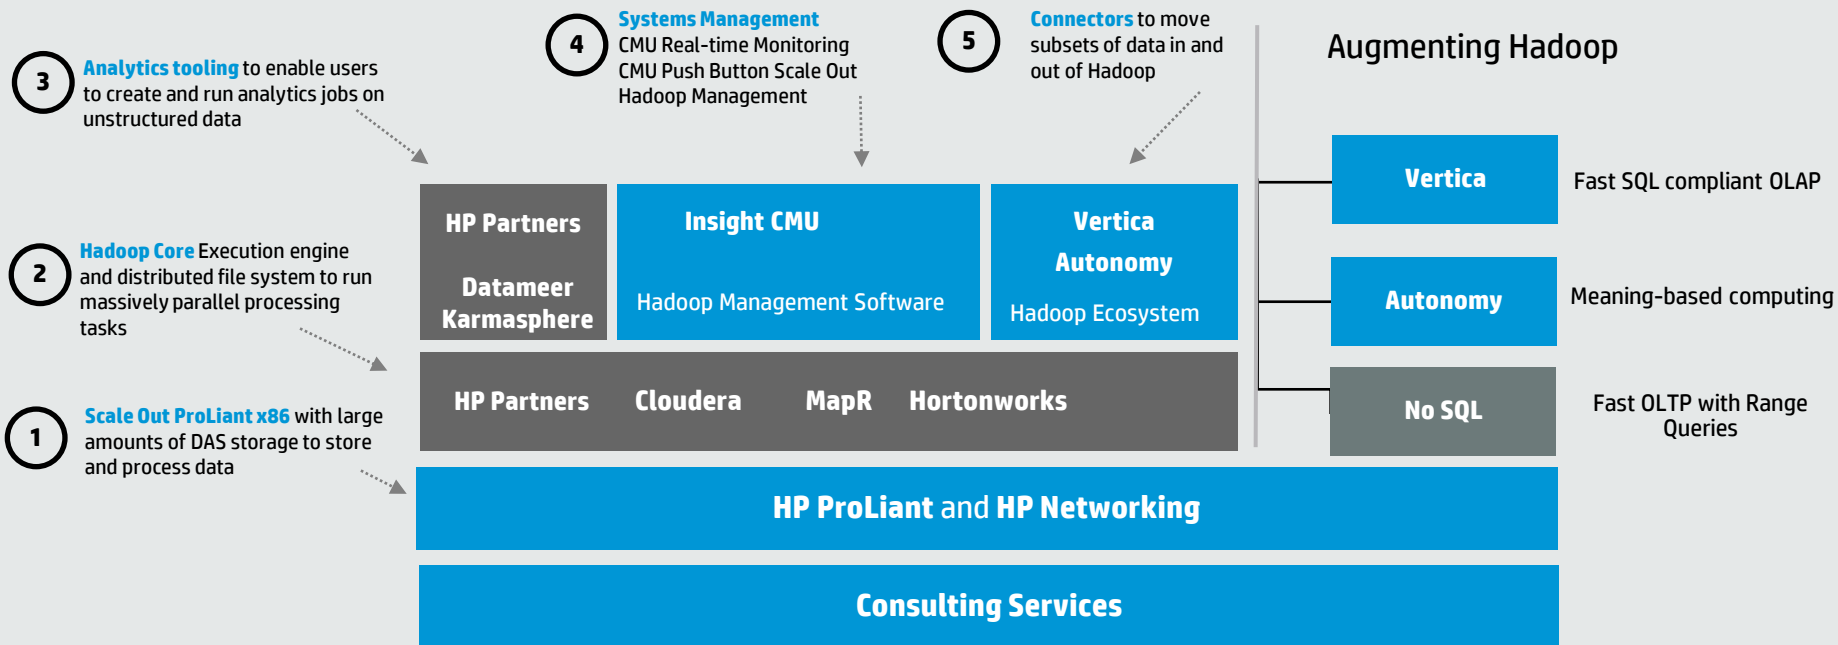
Management and Monitoring – HP CMU

- HP CMU (Cluster Management Utility) 는 push-button 식의 scale-out 관리 지원
- 동시에 1,000 대 이상의 서버 노드 관리 가능
- 빅데이터 환경의 중앙 집중형 콘솔 제공
- 복제를 통한 provisioning (optimized image propagation)을 통해 연속적인 시스템 확장 제공
- 대량의 시스템 노드 설치와 구성에 대해 필요한 단순 반복 작업 해소



HP Information Management Solutions

빅데이터 생태계의 A to Z 를 가능케 하는 확장형 플랫폼 제공




New Platform for Big data (SL454x)

빅데이터를 위해 설계된 첫번째 서버

SL4540 & SL4545
(1X60)

4Q12

Object Storage




**1 node
configuration**

Intel E5-2400 Series CPU를
탑재한 1대의 서버 및
60개의 3.5" HDD 설치 가능

SL4540 & SL4545
(2X25)

4Q12

Big data analytics




**2 node
configuration**

Intel E5-2400 Series CPU를
탑재한 2대의 서버 및 대당
25 개의 3.5" HDD 설치 가능

SL4540 & SL4545
(3X15)

1Q13

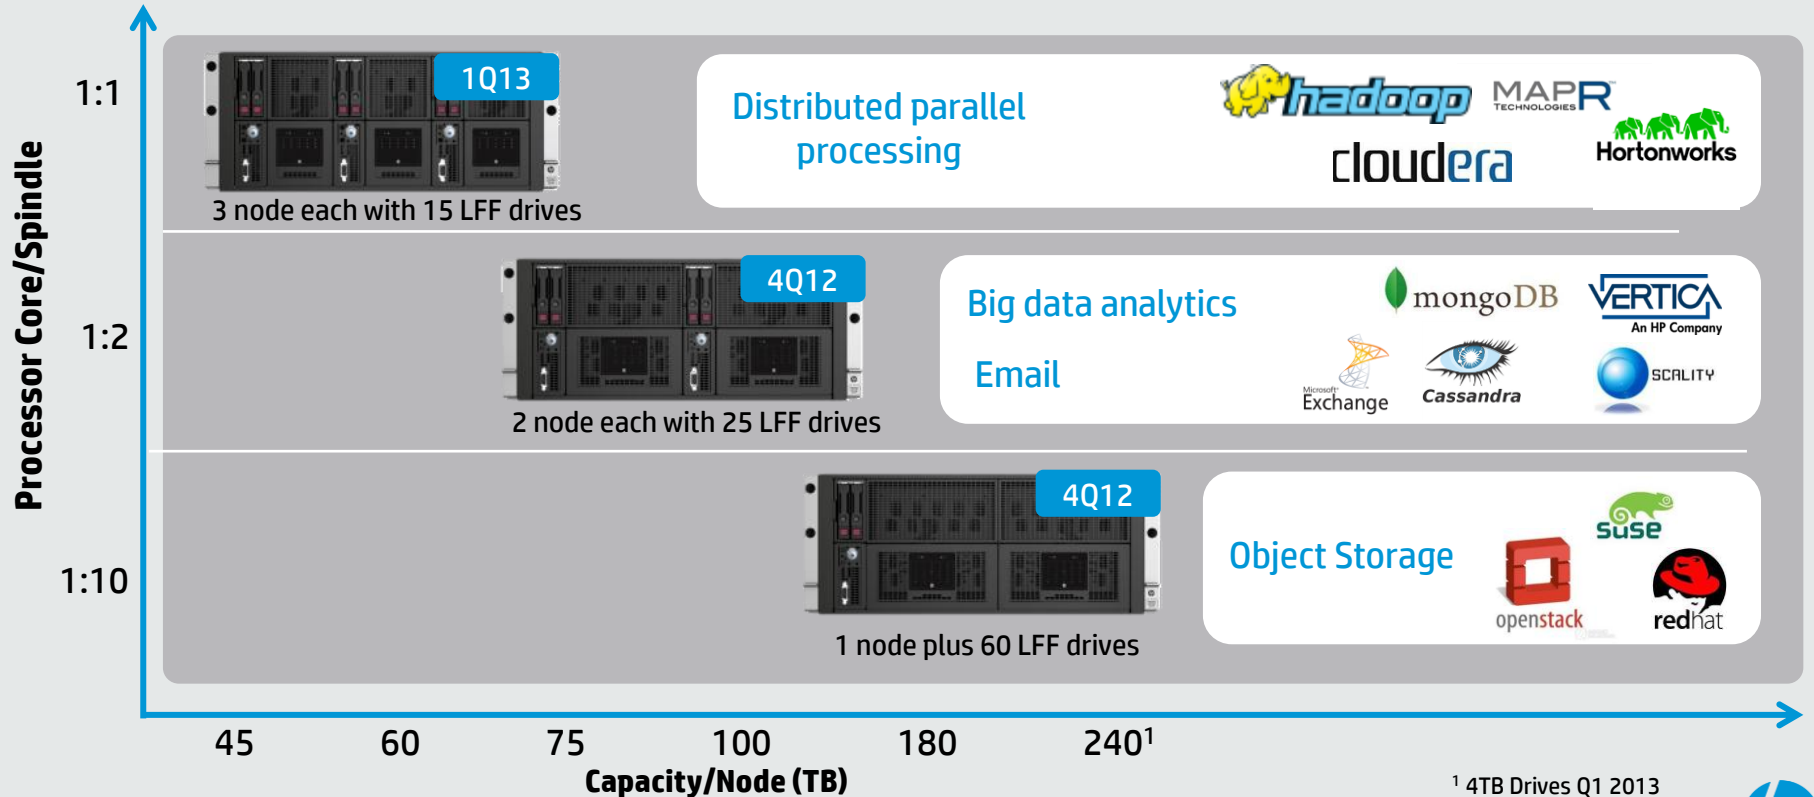
Hadoop



**3 node
configuration**

Intel E5-2400 Series CPU를
탑재한 3대의 서버 및 대당
15 개의 3.5" HDD 설치 가능

Optimized for big data workloads



Thank you

