Moderne Big Data Architektur – auf Hadoop Basis

Hans-Jürgen Fuks

EMEA Presales

Hans-Juergen.Fuks@hp.com

28.4.2015





Your Independent HP Business Technology Community

Agenda

Big Data Introduction

HP Haven Architecture

Hadoop – Basics

HP's Next Generation Big Data Reference Architecture





Your Independent HP Business Technology Community

Big Data Introduction



Every Enterprise action leaves a unique footprint 1% -5 % of the Digital Universe that actually is

being tagged and analyzed





The world is changing and accelerating

Big Data is no longer just a Buzzword – It's EVERYWHERE and growing...



Market Landscape

- Innovate quickly to remain competitive
- Global customer base, global workforce
- Digital data explosion intensifies compliance and security requirements



- 2x global cloud infrastructure growth in 24 months
- 75–80 % of new cloud apps will be Big Data intensive
- 10x explosion of new cloud apps in next 4 years
- 3x expansion of cloud developer community



Why does the business love big data?



Source : Harvard Business Review, October 2012 Source : McKinsey : Big Data – The next frontier for innovation, competition and productivity

Hadoop Market Opportunity

One of the fastest growing market segments



Different Kinds of Data need **DIFFERENT** Infrastructure **Human Information Machine Data** Annual Growth ~100% 90% of Information —> **Business** Data ~10%

10% of Information



Machine Data: Micro-transactions from machines



McKinsey : Big Data – The next frontier for innovation, competition and productivity

- Medical equipment
- Utility networks and meters
- Car and truck fleets
- Security sensors
- Home automation
- Touch-streams from games
- Drones
- Pollution sensors
- Transport sensors



Human Information: Meaning from interaction

- Faces
- Places
- Logos
- License Plates



Images

- Transcripts
- People



- Words
- Meaning





- Sentiments
- Meaning



- Complaints
- Sensitive info



Email

- Numbers
- Words
- Huge volumes



Documents



Big Data needs a different approach

One platform for structured, semi, and unstructured to profit from 100% of data

Yesterday's data warehouse and analytic infrastructure



Migration from disaggregated data sources to a unified system will be KEY to every organization



Big Data use cases by industries

			Ō	Å		ŧ
Retail & Manufacturing	Public Sector	Mobility	Healthcare	Telecom	Finance	Energy & Environment
 Brand monitoring Supply chain optimization Defect tracking RFID Correlation Warranty management 	 Law enforcement Fraud detection Intelligence service Traffic flow optimization 	 Data and voice quality Subscriber location Overage billing 	 Electronic patient record Medical imaging Gene sequencing Pharmaceutical research 	 Call detail records Customer retention Network optimization 	 Algorithmic trading High-frequency trading Risk modeling Anti-money laundering Fraud detection Portfolio analyzis 	 Billing Metering Weather forecasting Seismic logging Natural Resources Exploration

Horizontal Use Cases

- Sentiment analysis
- Social CRM / network analysis
- Churn mitigation
- Brand monitoring
- Cross and Up sell
- Loyalty & promotion analysis
- Web application optimization

- Marketing campaign optimization
- Brand management
- Social media analytics
- Pricing optimization
- Internal risk assessment
- Customer behavior analysis
- Revenue assurance

- Logistics optimization
- Clickstream analysis
- Influencer analysis
- IT infrastructure analysis
- Legal discovery
- Equipment monitoring
- Enterprise search

HP Haven Architecture



Big Data needs a unified approach

One platform for structured, semi, and unstructured to profit from 100% of data.





Having the right PLATFORM at the CORE is KEY HAVEn **Data Lake** PLATFORM Hadoop/HDFS Enterprise n Vertica Autonomy Apps **Unstructured Data Unstructured Data Security** like **Structured Data Structured Data** N number of **Context-Aware ArcSight for** Flexible Schema's **Semi-structured Data Applications** Analytics and **Unknown-structures** Intelligence Analyze at extreme Powering Process and index scale and Speed... Catalogue massive Collect & unify Software all information in real-time volumes of machine data + your apps distributed data 5 **69** 0 Transactional Social media Video Audio Email Texts Mobile IT/OT Search engine Documents Images data

HP's Platform called Haven is about INTEGRATION It is not a single Product or part number

HAVEn provides a platform for:

- Connectors,
- Applications,
- And Engines







semi-structured : XML, JSON, ...



Hadoop – Basics



Intel on Hadoop

"Within a couple of years, Hadoop will be the number one application. It will be running on more servers than any other single app. It will be more common for enterprise IT than their ERP system,"

Diane Bryant, senior VP and general manager of Intels data center group, at the Intel Developer Forum in San Francisco September, 2014



The Hadoop History

Created by Doug Cutting, named after his child's stuffed elephant !

- Pre-History (2002-2004)
 - Nutch sort/merge based web crawling
- Nutch Grows Up (2004-2006)
 - Google publish GFS and MapRed papers
 - Added DFS and MapRed to Nutch
 - Better scaling, easier to run
- Hadoop is Born (2006-2008)



- Doug Cutting hired by Yahoo
 - Dedicated Yahoo developers
- Running on 188 nodes / sort benchmark in 47.9H (04-2006)
- Hadoop splits from Nutch
- First stable release (09-04-2007)
- Web scale in 2008
- Hadoop is Open Source (2008)
 - As an Apache Top Level Project (01-2008)
 - Yahoo! use Hadoop to claim TeraSort benchmark (04-2008 / 209s 910 nodes)
 - 4000 node test cluster (07-2008)

"The name my kid gave a stuffed yellow elephant. Short, relatively easy to spell and pronounce, meaningless, and not used elsewhere: those are my naming criteria. Kids are good at generating such..."







The Beginning - Hadoop 1.x







The Pace of Change





The Pace of Change



And how people are buying Hadoop is changing also....



Strategy of ...integrating with Apache Hadoop YARN Store all data in a single place, interact in multiple ways

Single Use System

Batch Apps

1st Gen of Hadoop

MapReduce

(cluster resource management & data processing)

HDFS (redundant, reliable storage)

Multi Use Data Platform

Batch, Interactive, Online, Streaming, ...





Hadoop gets diverse - Hadoop 2.x



- Container based architecture
- Yarn resource management
- MapReduce is just another container based application
- Other projects adopting the Yarn Container model to run within Hadoop
- Various alterative file structures being adopted



Comprehensive Partner Ecosystem for Hadoop

Engineering level engagement with Hadoop partners enables HP to provide guidance



Open Partner strategy provides choice to customers

Apache Software Foundation – the Hadoop custodian





Hadoop Differences.

Performance and Scalability



Dependability/ Reliability

cloudera



Manageability

Data Access

10070	opensource		, i i opriciary
	Hortonworks	Cloudera	MapR
Performance and S	alability		
Data Ingest	Batch	Batch	Batch and streaming writes
Metadata Architecture	Centralized	Centralized	Distributed
HBase Performance	Latency spikes	Latency spikes	Consistent low latency
NoSQL Applications	Mainly batch applications	Mainly batch applications	Batch and online/real-time applications
Dependability			
High Availability	Single failure recovery	Single failure recovery	Self healing across multiple failures
MapReduce HA	Restart jobs	Restart jobs	Continuous without restart
Upgrading	Planned downtime	Rolling upgrades	Rolling upgrades
Replication	Data	Data	Data + metadata
Snapshots	Consistent only for closed files	Consistent only for closed files	Point-in-time consistency for all files and tables
Disaster Recovery	No	File copy scheduling (BDR)	Mirroring
Manageability			
Management Tools	Ambari	Cloudera Manager	MapR Control System
Volume Support	No	No	Yes
Heat map, Alarms, Alerts	Yes	Yes	Yes
Integration with REST API	Yes	Yes	Yes
Data and Job Placement Control	No	No	Yes
Data Access			
File System Access	HDFS, read-only NFS	HDFS, read-only NFS	HDFS, read/write NFS (POSIX)
File I/O	Append only	Append only	Read/write
Security: ACLs	Yes	Yes	Yes
Wire-level Authentication	Kerberos	Kerberos	Kerberos, Native

100% OpenSource Proprietary Proprietary

Table Source: Robert D. Schneider, Hadoop Buyer's Guide, Ubantu, 2014



Investments



- HP \$50M investment in HWX
- Strategic Joint Engineering between HP EG & HF Software and HWX Engineering
 - HP Servers (YARN Labels)
 - HP Autonomy IDOL
 - HP Vertica (HDFS and Hive)
 - HP Operations Mgmt (Ambari)
- HDP

- HP has invested to provide L1 & L2 support
- HDP is part of our standard offering for "as a service" in HP Helion Private Cloud
- HWX Strategic Partnerships with MSFT & SAP





HP and Hortonworks also included a seat on the Hortonworks board of directors, which will be filled by HP Executive Vice President and Chief Technology Officer Martin Fink.

Given Martin's current duties overseeing HP Labs/R&D and leading HP's cloud strategy, as well as <u>his</u> <u>leadership in HP's overall open-</u> <u>source and Linux strategies</u>, it's clearly an understatement to say that Martin has the right mix of experience that will make him an excellent addition to the Hortonworks board.



HP Reference Architecture(S) for Hadoop Optimized and tested configurations

- Flexible, pre-approved & optimized configurations
 - Scaling from 4 to thousands of HP ProLiant Gen8 Servers
 - Sized to customer's workload and storage needs
 - Impressive Processor and Storage density
 - <u>Reference Architecture</u> for unique HW configuration
- Processor, Drives, Network, 1TB/4TB disk size etc.

Breakthrough economics, density, simplicity

SL4540 3/15



400TB Hadoop usable for a full rack











HP ProLiant DL380 Gen9 - Hadoop Worker node

15 Data drives (12 x LFF for Data + 3 x LFF for Data with H240 + 1 x m.2 for OS)

Cloudera – Hortonworks :

- Worker Node MapR:
- **Control Services**
- Worker Services

Core/Spindle ratio = 1.33



HP ProLiant DL380 Gen9 – 15LFF: up to 90TB raw ~ 22.5TB Hadoop usable (min. without compression)

- 2 x 10-Core Intel Xeon E5-2660v3@2.6 GHz (20 cores 40 Threads)
- 8 x 16GB 128 GB RAM (DDR4 Registered Memory) ٠
- 90 TB raw storage (15 x 6TB LFF 6G SATA 7.2K)
- 120GB for OS (m.2 device)
- 3LFF Rear SAS/SATA Kit
- coming H240ar 12Gb 2-ports Int FIO Smart Host Bus Adapter
- H240 12Gb 2-ports Int Smart Host Bus Adapter ۰
- 120GB M.2 SSD ML-DL Sngl Enablement Kit
- Ethernet 2P 10GbE 561FLR





HP Insight CMU – pushbutton scale-out management

Control 1000s of nodes instantly



- Integrated and serves as central management console for HP Hadoop solutions
- Enables provisioning via cloning (optimized image propagation) for seamless scaling
- Reduces time and pain of system installation & configuration for each node in Hadoop cluster
- Battletested in HP cluster deployments at top 500 sites



2,500 nodes cluster view

HP's Next Generation Big Data Reference Architecture - Minotaur



Some interesting released Hadoop features

Architecture trends

• HDFS Tiering / Heterogeneous Storage Tiers (HDFS-2832)

For example, HBase can request that its data files (Hfiles) be stored on SSD. Then when HBase does writes and reads from HDFS, these requests will hit SSD and provide the latency requirements that HBase needs for supporting near real time applications.

• Phase2:

- HDFS-5682 Application APIs for heterogeneous storage
- HDFS-7228 SSD storage tier
- HDFS-5851 Memory as a storage tier (beta)
- HDFS Archival Storage Design (HDFS-6584)
 - Introduces a new concept of storage policies. For accommodating future storage technology and different cluster characteristics, cluster administrators will be able to modify the predefined storage policies and/or define custom storage policies.
 - Data policy names : Very Hot \rightarrow Hot \rightarrow Warm \rightarrow Luke Warm \rightarrow Cold



Hadoop gets asymmetric

but I thought we were taking the work to the data...



Yarn Labels

Allows applications running in yarn containers to be constrained to designated nodes in the cluster

HDFS Tiering

Allows the creation of pools of storage for SSD, HDD and Archive leveraging difference server configurations



Hadoop Big Picture

Note: N	Business Applications										
Of Ord Graph Processing Streaming Security Streaming Graph Processing Bagel, GraphX Bound Cooperation Workflows Machine Learning Data Mgmt Machine Learning Morkflows Morkflows <td< td=""><td>ne me</td><td>Ad-Hoc SQL Drill, Shark, Presto</td><td colspan="2">BI SQL Hive, Stinger, Impala</td><td colspan="2">Analytic SQL Vertica, Cassandra</td><td colspan="2">ODS SQL Trafodion, MarkLogic, Couchbase</td><td>Schemaless SQL Hadapt, Polybase</td><td colspan="2">HP DSM</td></td<>	ne me	Ad-Hoc SQL Drill, Shark, Presto	BI SQL Hive, Stinger, Impala		Analytic SQL Vertica, Cassandra		ODS SQL Trafodion, MarkLogic, Couchbase		Schemaless SQL Hadapt, Polybase	HP DSM	
Workflows Oozie, Cascading Machine Learning Mahout Data Mgmt HCatalog, Falcon Machine Learning MLib Workflows Oozie, Cascading Machine Learning Distributes Machine Learning MLib Workflows Oozie, Cascading Workflows Oozie, Cascading Workflows Oozie, Cascading Workflows Oozie, Cascading Machine Learning Workflows Oozie, Cascading Machine Learning Machine Learning <td>Log Co Flu</td> <td></td> <td>Graph Processing Pregel, Giraph</td> <td colspan="2">Streaming Storm</td> <td colspan="2">Security Knox, Rhino</td> <td colspan="2">Streaming Spark Streaming</td> <td>Graph Processing Bagel, GraphX</td> <td>oring p</td>	Log Co Flu		Graph Processing Pregel, Giraph	Streaming Storm		Security Knox, Rhino		Streaming Spark Streaming		Graph Processing Bagel, GraphX	oring p
Stript Script SQL NoSQL SQL Script Batch Compute Batch Compute Pig SQL Hive NoSQL SQL Spark Batch Compute Batch Compute Spark Spar	Budge		Machine Learning D Mahout Ho		Data HCatalo	Data Mgmt HCatalog, Falcon		Learning ^{Lib}	Workflows Oozie, Cascading	& Monit	
Image: Power Paragregic Paragregi Paragregic Paragregi Paragregic Paragregic Pa	Access Data Exch webHDFS Sqoop Coordination S Apache Zookeeper, Ap	Batch Compute MapReduce, MapReduce2	Script Pig	SQL Hive	NO: HBase, M MapR Bi	NoSQL HBase, MongoDB, MapR Big Tables		<mark>Script</mark> _{Spark}	Batch Compute Spark	gement	
Resource Management Varn Resource Management Varn, Mesos, Amazon EC2 Distributed File System HDFS2, MapR FS, Google FS, Amazon S3 Distributed File System		Execution Engine Tez				Execution Engine Spark				er Mana	
Distributed File System HDFS2, MapR FS, Google FS, Amazon S3		Resou		Resource Management Yarn, Mesos, Amazon EC2				<mark>u</mark> Cluste			
	FS , Fuse,		Distributed File System HDFS2, MapR FS, Google FS, Amazon S3								HP CM

HP ProLiant server

Hadoop learns SQL

	cloudera [®] IMPALA	HIVE	APACHE	Spark SQL
Primary Sponsor	Cloudera	Hortonworks	MapR	Databricks
SQL Executor	Impala	Hive/Tez	Drillbit	Spark
Associated File Structure	Parquet	Orcfile	Nested data model on numerous fs	In Memory



Trafodion - Introduction



Open source project to develop transactional SQL-on-Hadoop database engine

HBase

– Rides the unstoppable Hadoop wave!

- Transforms how companies store, process, and share big data
- Affordable performance, elastic scalability, availability
- Open source project downloadable for free
- Apache 2.0 software license
- Eliminates vendor lock-in and licensing fees
- Leverages community development resources and speed
- Schema flexibility and multi-structured data
- Capturing and storing all data for all business functions

Transactional SQL

- Full-function ANSI SQL

- Reuses existing SQL skills and improves developer productivity
- Distributed ACID transaction protection
- Guarantees data consistency across multiple rows, tables, and SQL statements
- Targeted for operational workloads!
- Optimized for real-time transaction processing applications, operational reporting, and Operational Data Stores (ODS)
- Leverages 20+ years of HP investments



HP's approach to address Big Data demands IT infrastructures must evolve to handle Big Data demands

Current traditional Big Data approach

- Compute and storage are always collocated
- All servers are identical
- **Data is partitioned across servers** on direct-attached storage (DAS)



Challenge....

- Compute to storage ratio
- Memory requirements
- Node width (narrow vs wide nodes)
- Special Accelerators GPU's, FPGA's, etc.

HP's NEXT Gen Approach for Big Data

- <u>Separate compute and storage</u> tiers connected by Ethernet networking
- <u>Standard Hadoop installed asymmetrically</u> with storage components on the storage servers and yarn applications on the compute servers



Multi Use Data & Compute Platform Batch, Interactive, Online, Streaming,

Building blocks for the HP Big Data Reference Architecture ...vs an Appliance Approach



Store all data in a single place, interact in multiple ways



Benefits of HP Big Data Reference Architecture

HP Moonshot and SL4540 addresses a variety of enterprise big data needs





Cluster consolidation

Multiple big data environments can directly access a shared pool of data

Flexibility to scale Scale compute and storage independently

Maximum elasticity

Rapidly provision compute without affecting storage

Breakthrough economics

Significantly better density, cost and power through workload optimized components

Current DFSIO testing on Moonshot/Argos



Advantages* of HP Big Data Reference Architecture

A new standard for Big Data delivery at scale



Traditional Architecture





HP Big Data Reference Architecture

* Normalized on performance

Independent Scaling of Compute and Storage

Traditional Architecture



B 80 Bo l a -B. 08.8 80 80 HOT COLDER 60% of the compute 4x compute 1.7x compute 60% of the storage capacity 1.5x the storage capacity 2x the storage capacity 72% of the Hadoop IO 2.1x the Hadoop IO 2.9x the Hadoop IO

HP Big Data Reference Architecture

Compared with traditional architecture, full rack



HP Moonshot System + SL4540 for Big Data

HP Big Data Reference Architecture TESTED With....

	Hadoo	p D	istr	ribution	
Hortonworks	Hortonworks Data Platform	ı	Cloudera Enterprise 5		
	Opera	atin	ıg S	ystem	
		Lin	ux		
	ŀ	larc	lwa	ire	
HP ProLiant m710) Server Cartridge			HP ProLiant SL4540 Scalable System	
Intel® Xeon® E3 Processor	480GB storage	Compute W	prker Node	Intel [®] Xeon [®] IPB E5 Processor storage	
32GB Memory	Dual port	/orker Node	Storage Wo	192GB Dual port Memory 10GbE	

Built on Standard Hadoop Distributions

No proprietary software. Leverage the latest versions of Hadoop and consumer plug-ins

Optional HP Consulting Services

Expedite the sizing and configuration of the infrastructure through the Hadoop Reference Architecture implementation service

Optional Factory Build

Hardware racked, wired and tested, delivered to your data center



With this Approach...

- All components are now generally available to deploy our Reference Architecture
 - ✓ Integrated and Tested with Best Practices
 - ✓ Built by HP's Factory Express
 - ✓ And support by HP ...



Value of this Strategy...

- IT Infrastructure...
- Ability to have a federation of different platform compute needs
- Greater flexibility with Compute to Storage ratio...
- SPEED to move DATA (between Compute nodes and Storage) and Offload CPU's with data movement.
- Don't need to build a unique Cluster... (use a common set data)
- Best of Breed Hardware... Price/Performance/Density/Power No Vendor LOCK-in...

Hadoop USER...

• Different workloads...

because each tool has different needs use the right tool for the job

- Focused Hadoop and Open Standards... no proprietory storage Arrays/Appliances)
- **Built on HDFS** as a common Data Layer...With Yarn to Optimize Workloads... on Industry Std. Hardware (density/price perf/power)
- Enhanced... capabilities with classes, data locality, and scheduling for better agility/utilization ...integration with hardware, software and needs.

And a BETTER CHOICE Using the Right Tool, with the Right Person, for the Right JOB...



Big data is a federation of platform needs



Low cost, dense x86 servers with high speed, low cost networking for mainstream work



FPGA assisted for offloading stream processing and CEP





Large, dense, low cost memory with a scale-out optimized CPU and fabric

DSP assisted

cartridges for

extraction

offloading voice

and video feature



Lowest cost, dense storage with maximum IO capability



HPC GPU Servers for acceleration of deep learning algorithms



"Wide Nodes" for special needs such as scale up for single large address space

• Then we need to unify this into a rational architecture that can share data





Minotaur CI for Big Data long term view

Evolve to support multiple compute and storage blocks

Workload Optimized compute nodes to accelerate various big data software



Multi-temperate Storage using HDFS Tiering, NoSQLs and Objectstores



Collaterals

Generell Informations on Big Data

- Big Data Overview: <u>www.hp.com/go/bigdata</u>
- HP and Hadoop: <u>www.hp.com/go/hadoop</u>
- HP's Haven Architecture: <u>www.hp.com/go/haven</u>

Generell reference Architectures:

- <u>http://h17007.www1.hp.com/us/en/converged-infrastructure/converged-systems/bigdata-hadoop.aspx#tab=TAB3</u>
 - HP Big Data Reference Architecture: A Modern Approach:
 - <u>http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA5-6141ENW&cc=us&lc=en</u>
 - HP Big Data Reference Architecture: Cloudera Enterprise reference architecture implementation:
 - <u>http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA5-6137ENW&cc=us&lc=en</u>





