



## BIG DATA IN CAPITAL MARKET

DEJTECH LILLE 1 – JANUARY 2015

# Agenda


1. Finaxys & Finaxys Academy
2. Big Data & Finance Industry
3. Hadoop / HBase
4. Pitfalls and Challenges
5. Product & Trade Design
6. Scaled Risk & Twitter for Trading Decision

Note → Some materials/ideas are from

- Nick Dimiduk – HBase for Architect (HortonWorks)
- Roman Wilhelm (Scaled Risk)
- Lars Hofhansl - HBase Internals (Salesforce)
- Ian Varley – HBase Schema Design (SalesForce)

Some contents are Licensed under a Creative Commons Attribution-ShareAlike 3.0 Unported License.

## Pierre Bittner

- CTO of Finaxys
  - 2000: Lille 1 Master 2 IAGL
  - [pbittner@finaxys.com](mailto:pbittner@finaxys.com)
-  @BittnerPierre

# FINAXYS

## Conseil IT Finance

- **250 collaborateurs** spécialisés en **IT Finance** (création en 2008)
  - CA 2014 : **27 M€**
  - **International**: Londres, Belgique, Luxembourg
  - **Filiales spécialisées**: Mobilité, Médias Sociaux, Cloud Computing & **Big Data (Editeur)**
- Une **large palette d'expertises** : Conseil SI & fonctionnel, Business Analyst, Commando, IT Quant, Chef de Projet, Architecture, intégration, développement Java/C#/C++, production et support applicatifs
- Un **accompagnement global** sur l'ensemble du cycle de vie projet, du cadrage à la recette, associé avec une **démarche qualité intégrée**
- Reconnaissance du **savoir-faire**
  - Spécialisation sur secteur banque et assurances (90% du CA)
  - Ranking fournisseur: Top 5 BNP CIB, Top 10 SGCIB, Top 3 CACIB
  - Maitrise des projets à engagement de résultats (Proof of Concept, Forfait, ...)
- Positionnement **stratégique** de nos offres
  - Conseil et mise en œuvre de solutions **Big Data** sur mesure (Hadoop/HBase) – Scaled Risk
  - Approche pragmatique du **Continuous Delivery / DevOps** et de la transformation Agile
  - **Finaxys Academy** – L'incubateur de jeunes talents de l'IT Finance

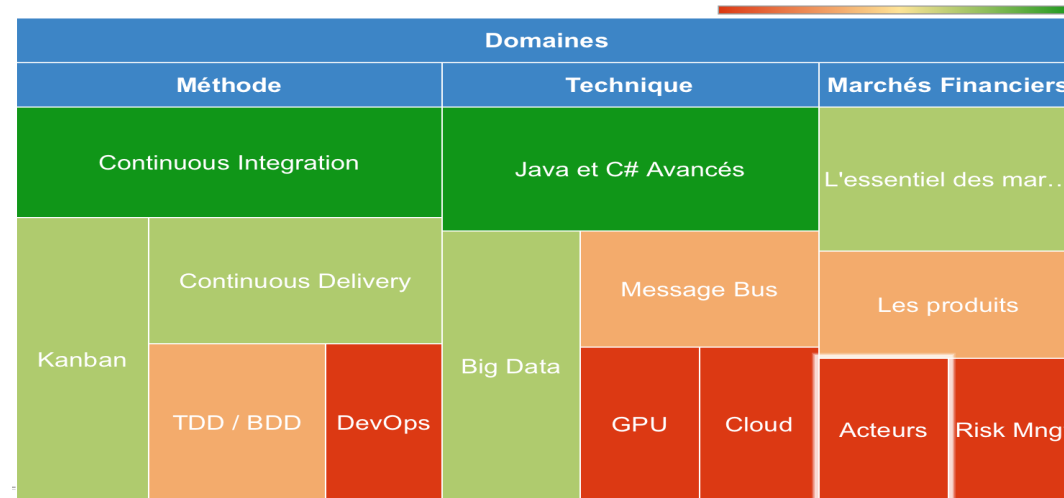
# Finaxys Academy

Incubateur de jeunes talents

La FINAXYS Academy, c'est une université interne incubatrice de jeunes talents, qui forme les consultants pendant leur stage de fin d'étude (6 mois) chez FINAXYS puis chez nos clients, en apportant un savoir-faire technique d'excellence mais également des connaissances fonctionnelles et méthodologiques (Finance de marché, Développement Objet, Big Data, GPU, Scrum, etc).

4 jeunes ingénieurs ont suivi le cursus en 2014. Grâce à la qualité et pertinence du parcours, deux d'entre eux ont intégrés les exigeantes équipes de Front-Office un mois avant la fin de leur parcours.

Début 2015, **8 nouveaux étudiants** intégreront l'académie.



Inscription: [rh@finaxys.com](mailto:rh@finaxys.com)

## Tronc commun

- Scrum et les méthodes agiles
- Le Continuous Delivery
- L'essentiel des marchés
- Programmation avancée

## Les Majeures

- Le Big Data en Finance
- Cloud/DevOps
- Les applications de Trading

## Les Mineures

- Pricing et calcul parallèle
- Infrastructure as Code

# Actors of Financing Industry

## Banking Group

Bank of America, CRÉDIT AGRICOLE, NOMURA, UBS, HSBC, ING, BARCLAYS, Deutsche Bank, SOCIETE GENERALE

## Corporate & Investment Bank

BNP PARIBAS CORPORATE & INVESTMENT BANKING, JPMorgan, NATIXIS, SOCIETE GENERALE Corporate & Investment Banking, CRÉDIT AGRICOLE CORPORATE & INVESTMENT BANK, Goldman Sachs

## Online Bank, Private Banking

BOF BANK, PICTET, INDOSUEZ PRIVATE BANKING, SOCIETE GENERALE

## Data Prov., Clearer, Depository

Bloomberg, markit, LCH.CLEARNET, REUTERS, euroclear, SWIFT

## Insurance

AXA, Allianz Global Investors, Groupama, AIG, GENERALI GROUP, SwissLife

## Asset Manager

Amundi ASSET MANAGEMENT, INVESTMENT MANAGERS, caceis INVESTOR SERVICES, BLACKROCK, AVIVA INVESTORS LYXOR

## Retail Banking

LA BANQUE POSTALE, LCL BANQUE ET ASSURANCE, BANQUE POPULAIRE BRED

## Custody Services

BNP PARIBAS SECURITIES SERVICES, SOCIETE GENERALE Securities Services

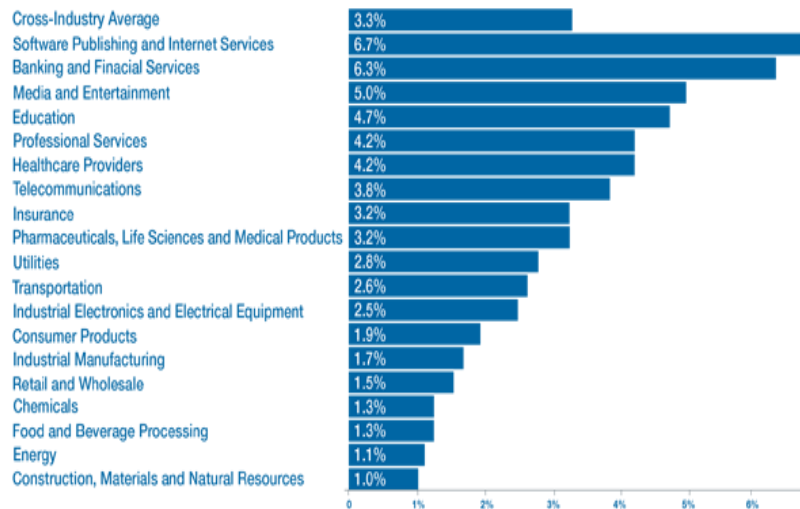
## Stock Exch., Broker, Central Bank

Euronext, NYSE, Newedge, Nasdaq, BANQUE DE FRANCE, EUROSYSTEME, Kepler Cheuvreux, Caisse des Dépôts

# Finance, Banking, Insurance: Biggest investor in IT

6.3% of revenue spent on IT, #2  
3.3% industry average

## IT SPENDING AS A PERCENT OF REVENUE, BY INDUSTRY



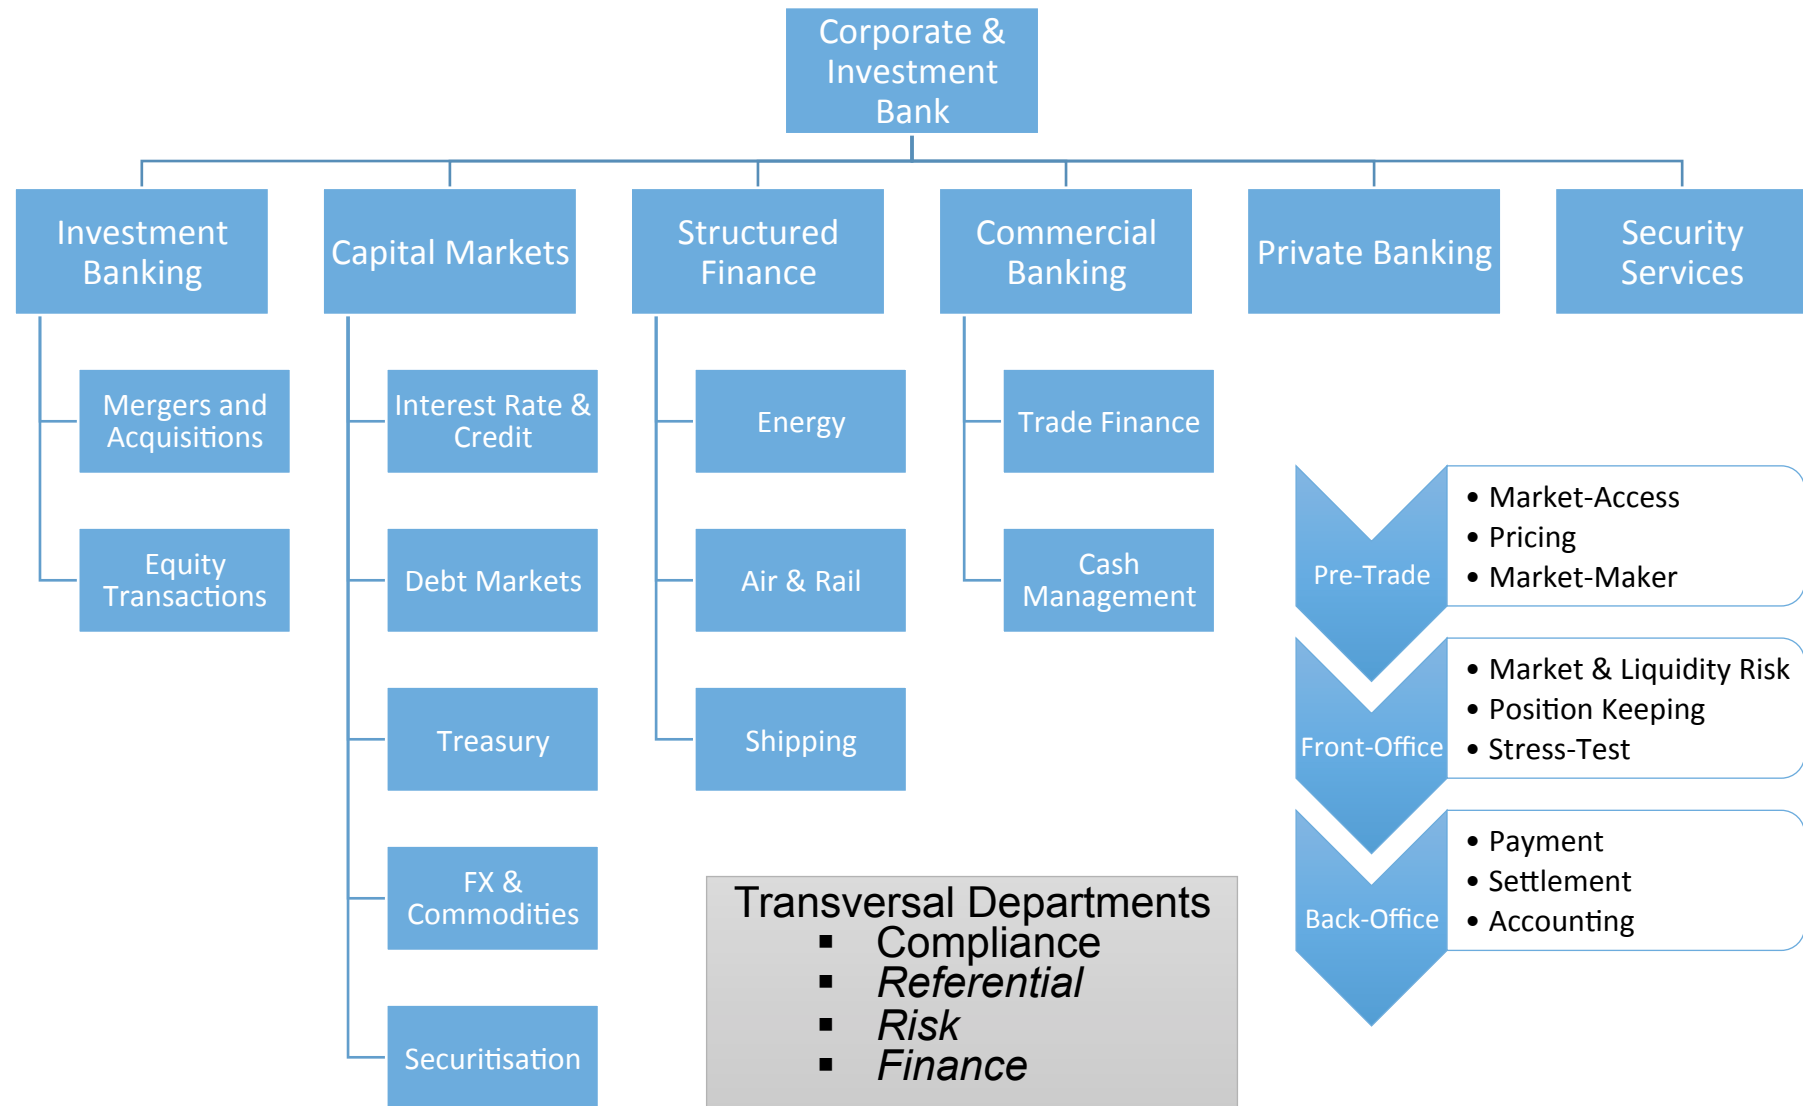
Source: IT Key Metrics Data 2014, Gartner Benchmark Analytics

## Among top Industries Investing in IT

Banking	\$59.4 billion
Communications & Media	\$50.7 billion
Construction	\$34.2 billion
Discrete Manufacturing	\$33.7 billion
Securities and Investment Services	\$31.5 billion
Health Care	\$24.9 billion
Insurance	\$23.5 billion
Process Manufacturing	\$21.5 billion
Resource Industries	\$19.6 billion
Retail	\$19.0 billion
Professional Services	\$15.5 billion
Transportation	\$15.0 billion
Utilities	\$9.6 billion
Wholesale	\$5.7 billion
Consumer & Recreational Services	\$4.1 billion
<b>Total</b>	<b>\$368.5 billion</b>

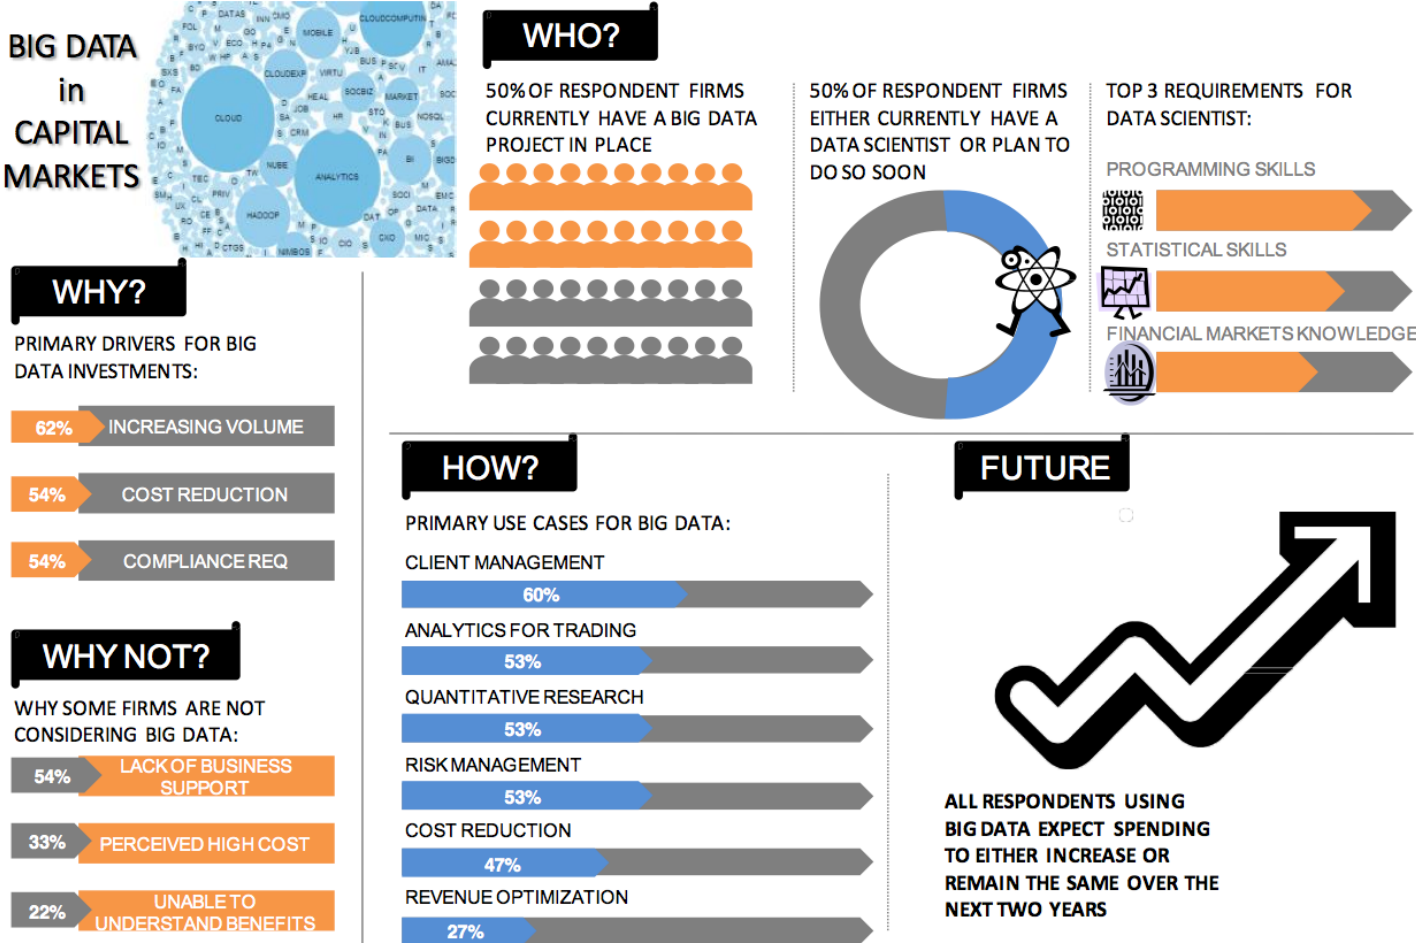
- IT is considered strategic for Business
- Speed, Efficiency in Trading and Risk management with cost control are a key differentiators
- Smaller trends to externalize development

# Corporate & Investments Bank Organization



# Reuters: Highlight of Big Data Survey in Capital Market

Volume, Variety, Velocity – Just the start of the Journey

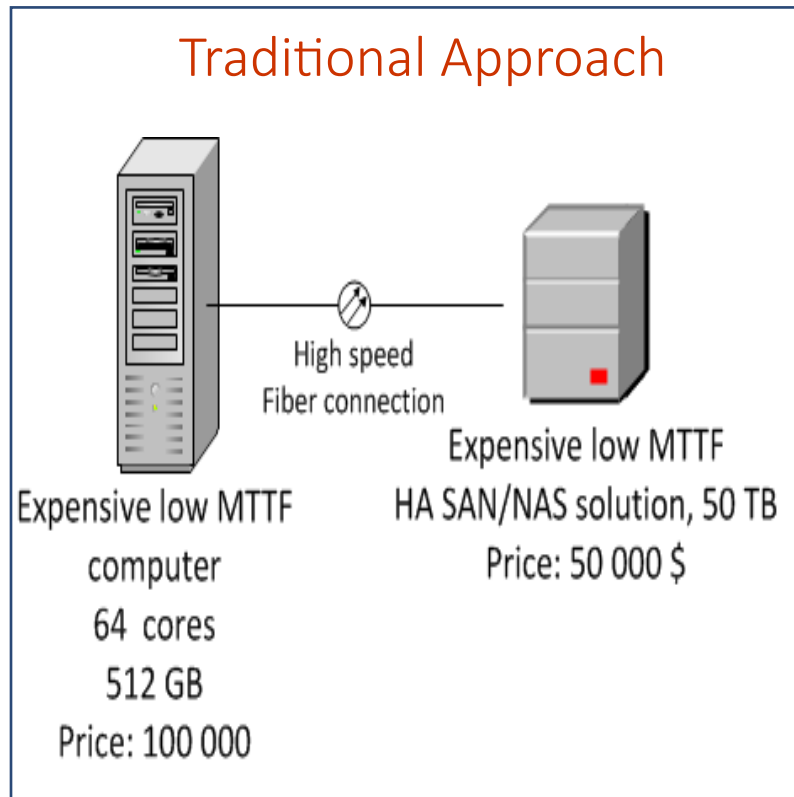


Source: Aite Group survey of 22 capital markets participants, May and June 2014

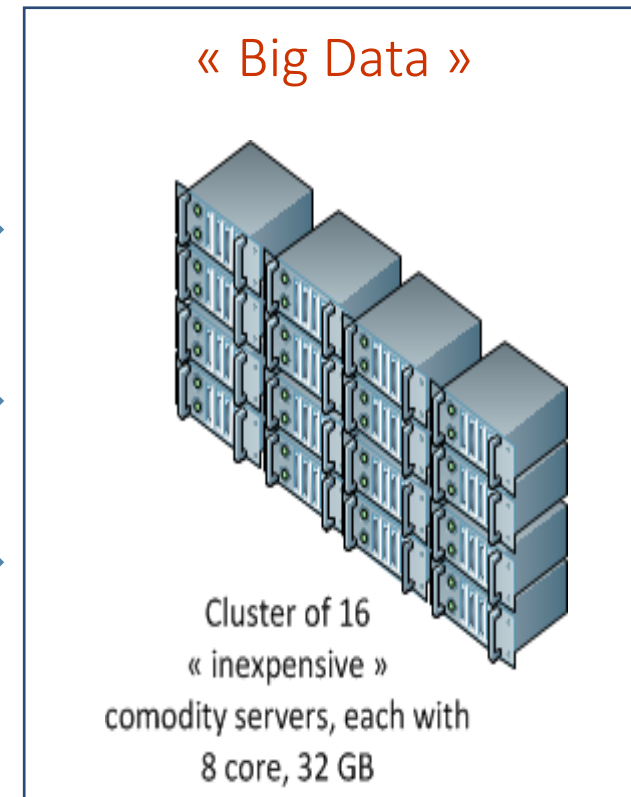



# Big Data: architecture

## How to handle large data set (> 3 TB)



Reliable but expensive hardware

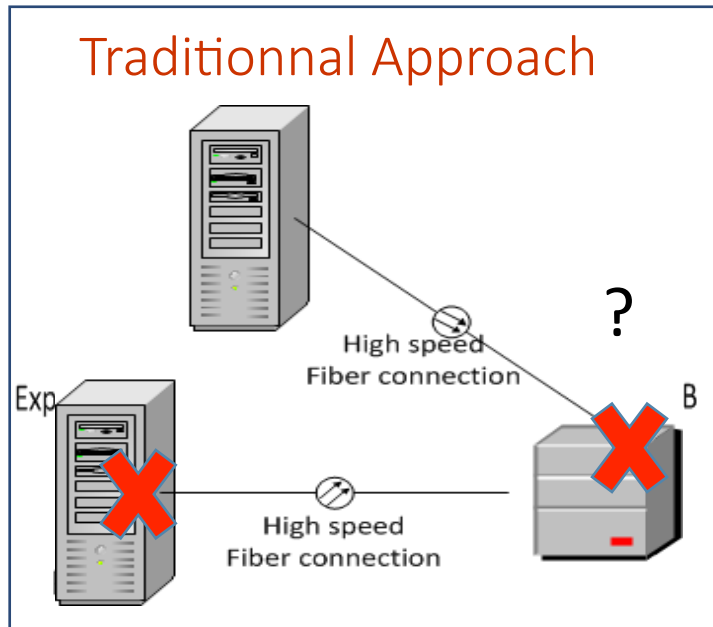


Commodity hardware, availability and distribution handle by software layer

- Performance by collocated CPU and storage
- Split and distribute data across cluster
- Maximizing usage of RAM (no OS disk sync)


# High Availability

« Even the most reliable hardware fails »



Traditional: hardware redundancy

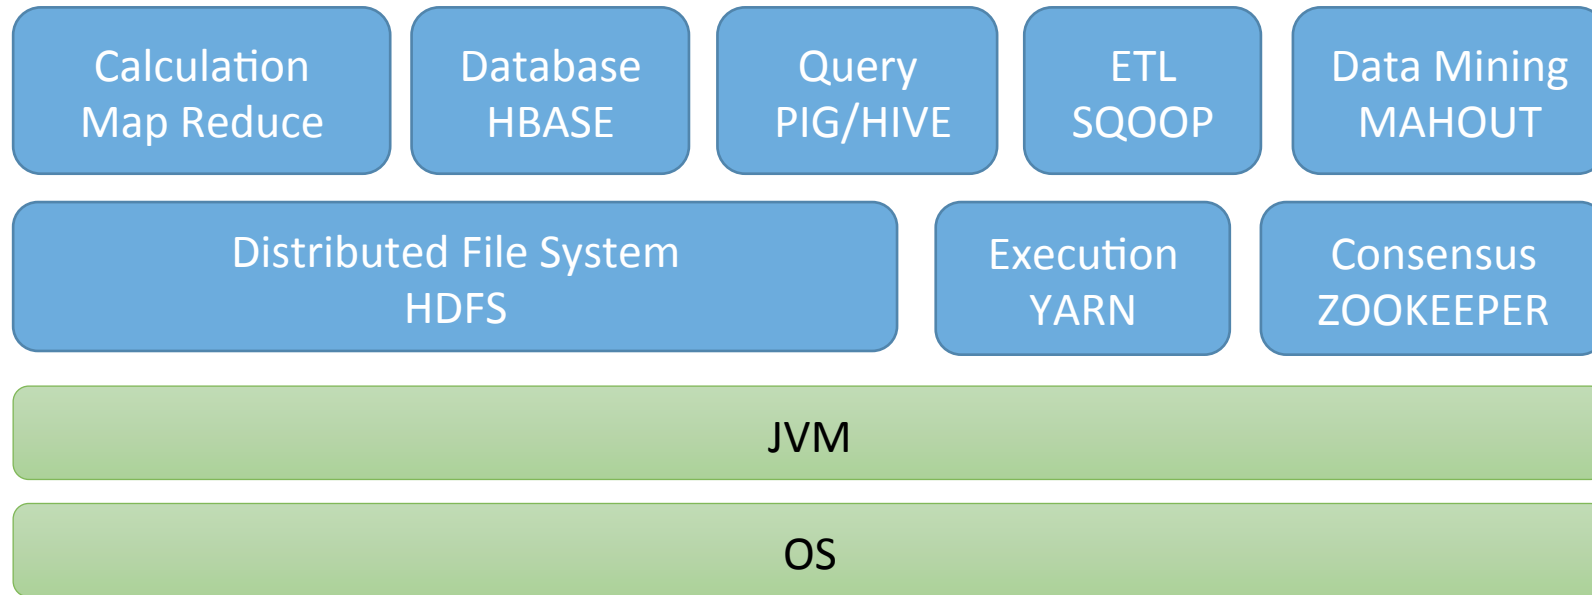
- Expensive hardware that must be replicated
- Risk during failure
- Long recovery time



Big data: data redundancy

- In the core of the architecture
  - Memory and Network faster than Disk
  - Use collocated server as fast backup system
  - Up to 30% of hardware lost
- => Best service continuity

## What (with Hadoop Wording) ?




### Quite independent components:

- You can use ZooKeeper alone
- You can use HBase on top of a different file system
- You can do HIVE queries w/ or w/o HBase
- And so on.

Hive: SQL access, and JDBC

Hadoop, “Linux” of Big Data



# What is Hadoop / Zookeeper / MapReduce / HBase ?

- **Apache Hadoop Distributed Filesystem**
  - Distributed, fault-tolerant, throughput-optimized data storage
  - [The Google File System](#), 2003, Ghemawat et al.
- **Apache Zookeeper**
  - Distributed, available, reliable coordination system / Consensus
  - [The Chubby Lock Service](#)..., 2006, Burrows
- **Apache Hadoop MapReduce**
  - Distributed, fault-tolerant, batch-oriented data processing
  - [MapReduce](#): ..., 2004, Dean and Ghemawat
- **Apache HBase**
  - [BigTable paper from Google](#), 2006, Dean et al.
  - “Bigtable is a sparse, distributed, persistent multi-dimensional sorted map.”
  - **Key Features:**
    - Distributed storage across cluster of machines
    - Fast random, online read and write data access
    - Schemaless data model (“NoSQL”)
    - Self-managed data partitions
    - Strong consistency

## Users

### Early Adopters

- Twitter
- Facebook
- Criteo
- Yahoo!

### 2<sup>nd</sup> Wave in Finance

- Salesforce
- Bloomberg
- American Express
- Goldman Sachs

# HBase Logical Data Model

Table A

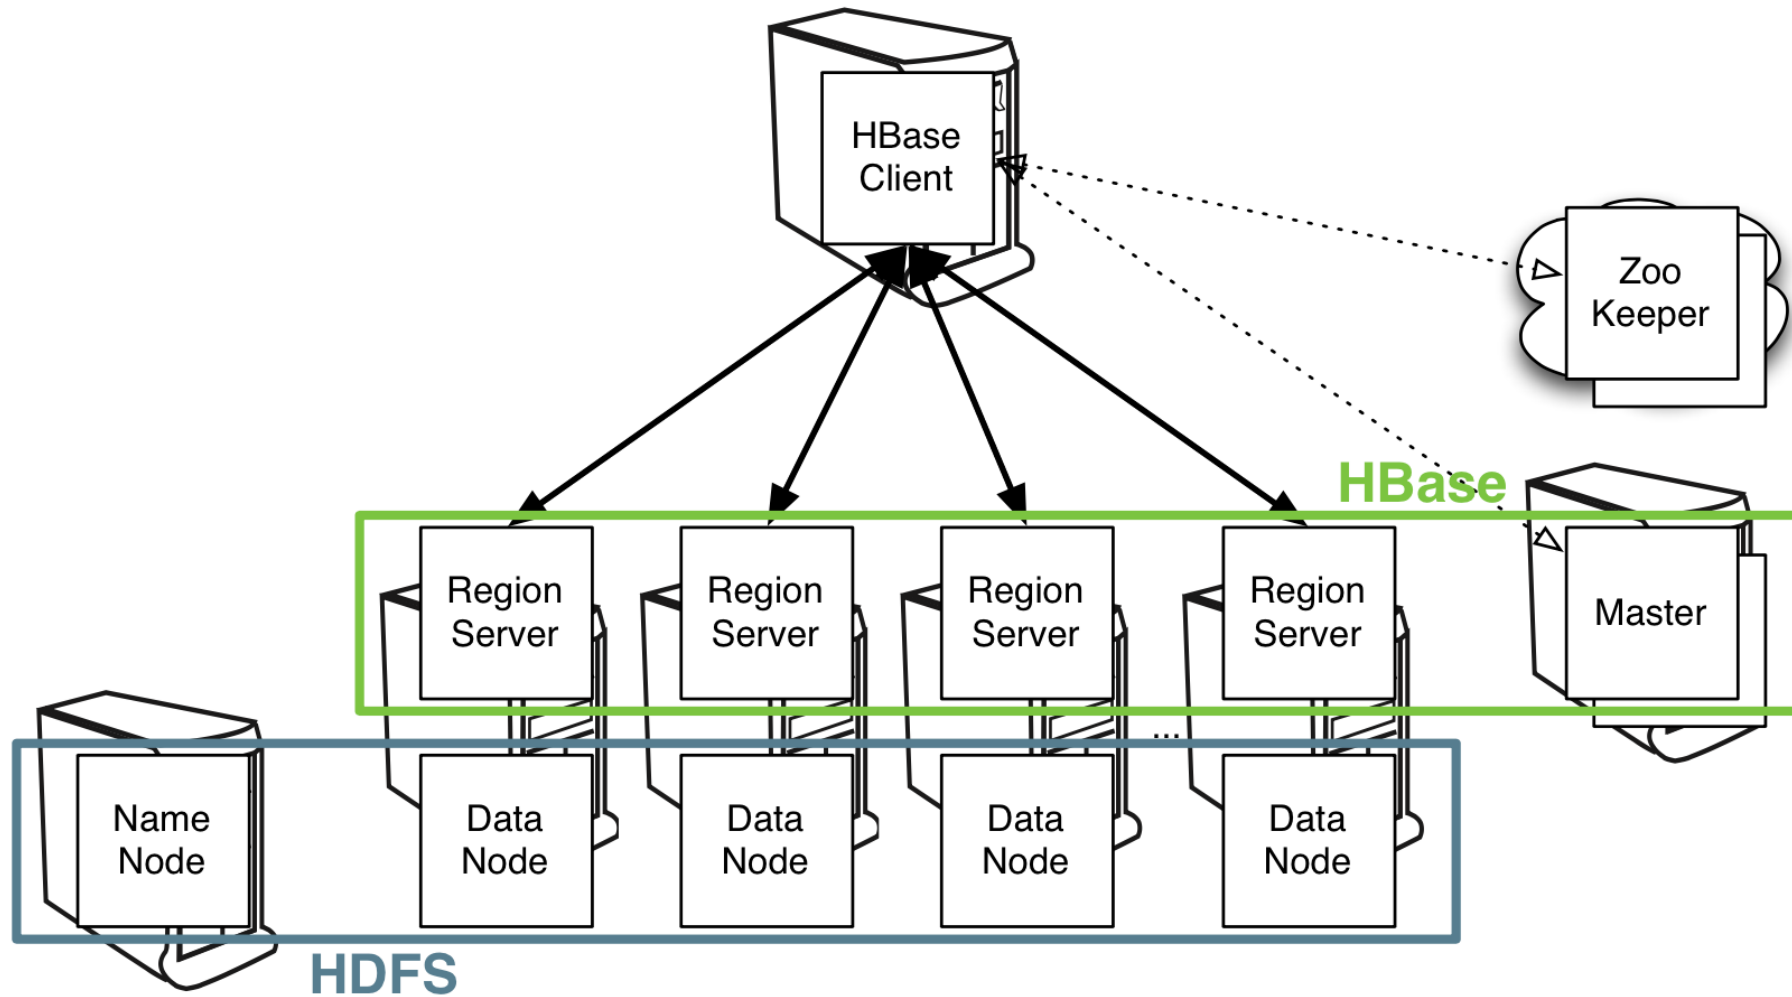
## Column Families

Rows

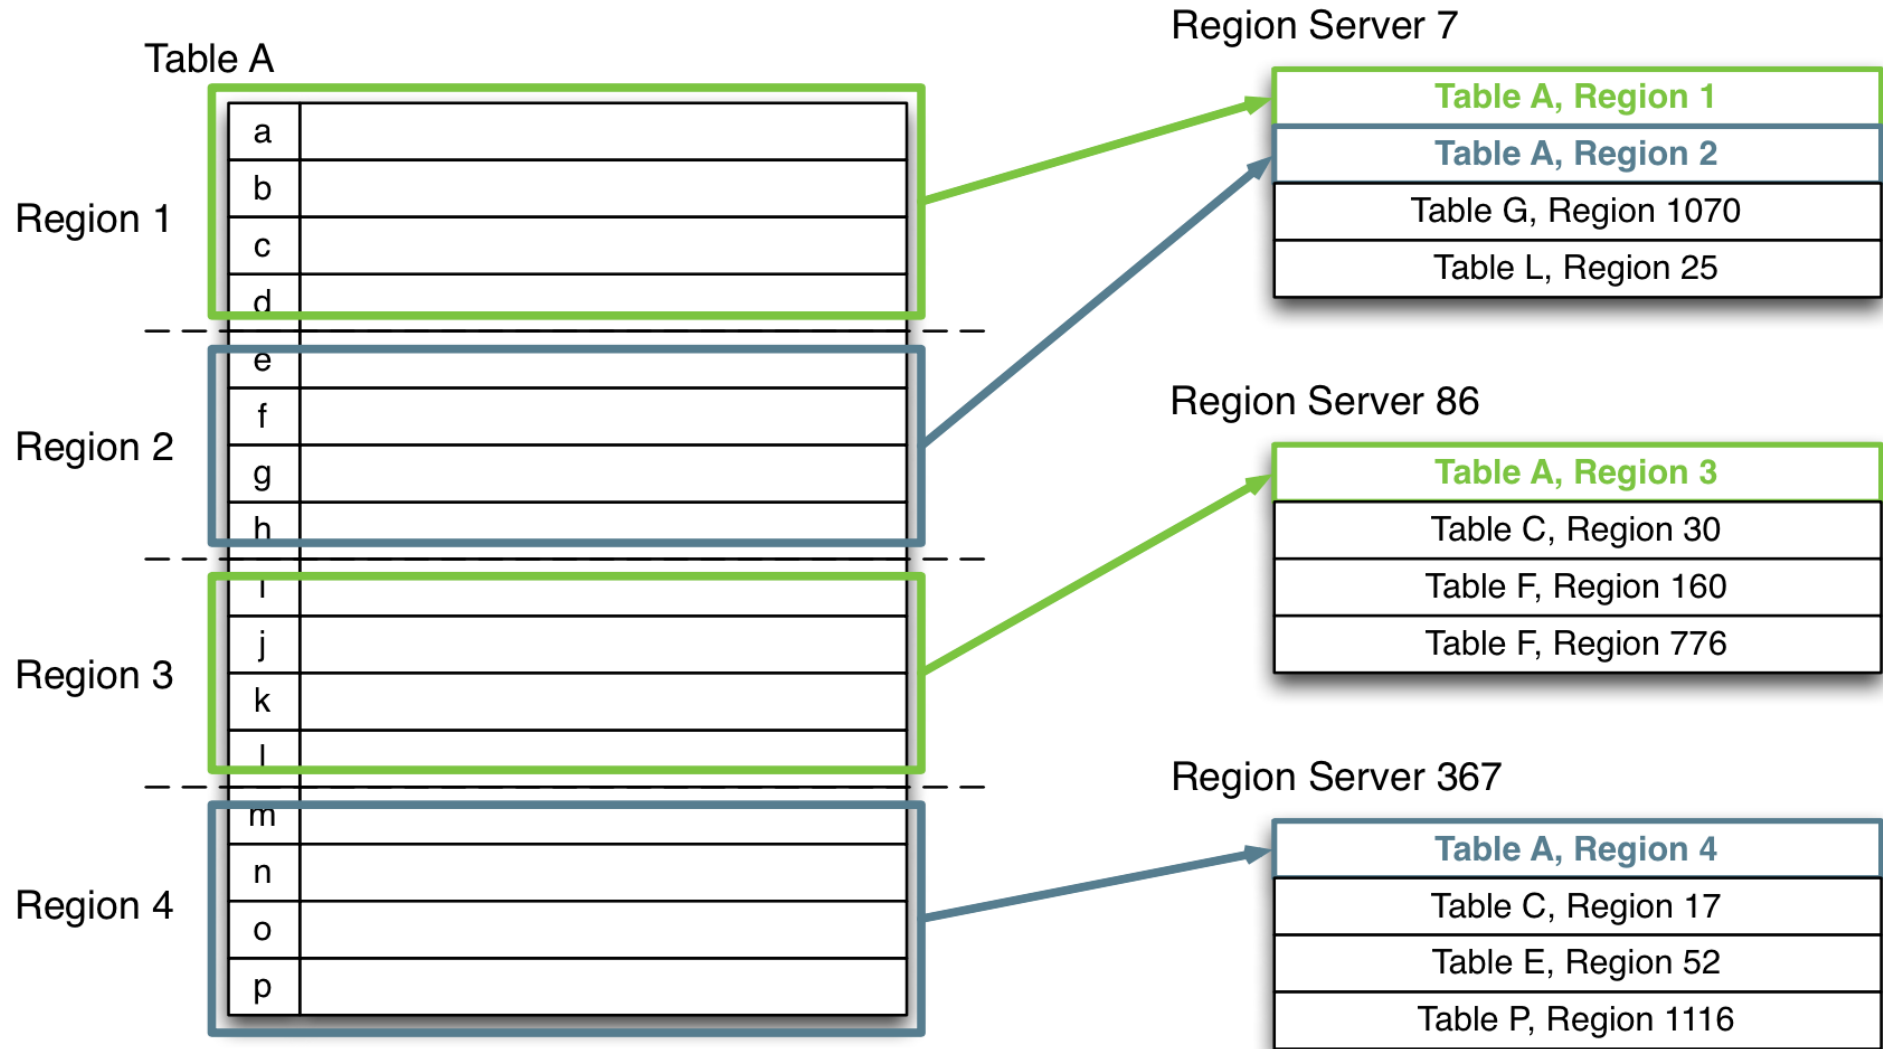
rowkey	column family	column qualifier	timestamp	value
a	cf1	"bar"	1368394583	7
			1368394261	"hello"
		"foo"	1368394583	22
			1368394925	13.6
		1368393847	"world"	
		cf2	"2011-07-04"	1368396302
		1.0001	1368387684	"almost the loneliest number"
b	cf2	"thumb"	1368387247	[3.6 kb png data]

Ow! Important note: rowkey, column qualifiers and values are bytes; timestamp long integer.

# HBase Physical Architecture



# HBase Logical Architecture



# User API

- **{rowkey => {family => {qualifier => {version => value}}}}**
  - Think: nested TreeMap (Java), OrderedDictionary (C#), OrderedDict (Python)
- **Basic data operations: GET, PUT, DELETE**
- **SCAN over range of key-values**
  - benefit of the sorted rowkey business
  - this is how you implement any kind of "complex query" \*
- **GET, SCAN support Filters**
  - Push application logic to RegionServers
- **INCREMENT, APPEND, CheckAnd{Put,Delete}**
  - Server-side, atomic data operations, can be contentious!

\* This is also a foundational component in what we refer to as “schema design” in this “schemaless” database.



# Big Data Pitfalls and Challenges in Capital Market

Big Data is about Volume, Velocity, Variety at cheap price but...

- What about **Verifiability** and **Veracity** ? (Transaction ie. ACID)
- CM mostly have structured data
- We are doing real-time data, not only batch
- What about legacy systems and education ? (our staff speaks mainly SQL)
- We are already doing big data (Data Grid, Compute Grid, Exadata,...)
- What about Disaster Recovery Plan, Infrastructure, online backup (Regulatory)?
- It will be a huge project (multi-year)...
- And we want to deliver quickly new business values!

Let's see some applications of Big Data in Capital Market

# Products & Trade Model

Money Market		
P.Key	Reference	varchar[32]
columns	TradeDate	date
	MType	char[3]
	Counterpart	varchar[32]
	Book	varchar[32]
	Way	char[1]
	ValueDate	date
	MatDate	date
	Notional	numeric
	Ccy	char[3]
	Index	char[5]
	Term	char[2]
	Spread	numeric
	Int. Rate	numeric
	ISIN	char[10]
	Issuer	varchar[32]

Derivative		
P.Key	Reference	varchar[32]
columns	TradeDate	date
	OptionType	char[3]
	Counterpart	varchar[32]
	Book	varchar[32]
	Way	char[1]
	ExerciceDate	date
	ExpiryDate	date
	StrikePrice	numeric
	Ccy	char[3]
	Underlying	numeric
	Premium	numeric
	Quantity	int
	Symbol	char[10]
	LotSize	int

- Loan
- Deposit
- Commercial Paper
- Certificate of Deposit
- Call Account
- ...

Trade		
P.Key	Id	int
columns	Trader	varchar[32]
	UserCre	varchar[32]
	UserMod	varchar[32]
	EventDate	date
	EventType	varchar[5]
	System	varchar[32]
	ProfitCenter	varchar[32]
	Status	varchar[5]

- Call
- Put
- Warrant
- ...

# Consolidation challenge / Too Many Product Models

As many trade schema than systems, product line (Equity, Fixed Income, Foreign Exchange...) and service layer (FO, BO...)

- Real challenge to find a common dictionary... if possible
- Huge problem for consolidation systems (360° Customer view, Risk, Compliance)
- As many “column” groups as system to integrate for a row
- Start to be Big Data (billions of lines, To of Data); many process in //; one hour SLA

Common fields

Specific fields (\*)

## Money Market

ID	Reference	Way	EventType	EventDate	TradeDate	Counterpart	Book	MMType	ValueDate	MaturityDate	Notional	Currency	Index	Term	Spread	InterestRate	ISIN	Issuer
1	AA001	Buy	Create	13/01/15	13/01/15	GOOGLE	EURO	Deposit	15/01/15	15/06/15	10 000 000,00	EUR	FIXED			2,50000		
2	GU002	Sell	Modify	13/01/15	05/01/15	TWITTER	USD	CD/CP	07/01/15	07/01/16	1 000 000,00	USD	EIBOR	3M	0,50		FR32432432	CACIB
3	MM003	Buy	Payoff	16/01/15	15/01/15	APPLE	OTH	Loan	17/01/15	17/05/15	1 000 000,00	JPY	FIXED			3,50000		

## Foreign Exchange

ID	Reference	Way	Trader	EventType	TradeDate	Counterpart	Book	FXType	SpotDate	ValueDate	BoughtAmount	BoughtCcy	SoldAmount	SoldCcy	SpotRate	FwdBoughtAmt	FwdSoldAmt	FwdRate
1	FXS005	Buy	HCE	Create	14/01/15	AIRBUS	SPTEUR	FXSpot	15/01/15		100 000 000,00	EUR	122 105 000,00	USD	1,22105			
2	FXS006	Sell	HCE	Create	15/01/15	ARCELOR	FWDEUR	FXForward		15/02/15	100 000 000,00	EUR	122 105 000,00	USD	1,21605			
3	FXS007	Buy	HCE	Delete	16/01/15	JPMORGAN	SWPUSD	FXSwap	19/01/15	18/06/15	100 000 000,00	USD	81 896 729,44	EUR	0,81897	100 000 000,00	80 512 056,68	0,80512

## Derivative

ID	Reference	Way	Trader	EventType	TradeDate	Counterpart	Book	OptionType	ExerciseDate	ExpiryDate	StrikePrice	Currency	UnderlyingPrice	Premium	Quantity	Symbol	LotSize
1	DER010	Sell	LRI	Exercise	14/01/15	HUawei	NASDAQ	Call	15/01/15	30/01/15	98,00	USD	114,45	15,40	10	AAPL150102C0	100,00
2	FXS007	Sell	LRI	Create	14/01/15	HSBC	SHE	Put		02/02/15	315,00	USD	312,00	2,14	25	AMZN150102C0	50,00

While this massive amount of data needs common enrichment (data normalization, valuation), aggregation or deep analyses, transactions and schema are fundamental issues!

\*: Reality: hundred of columns

# Versioning and audit of trade transactions

## In RDBMS

- Keep all versions in same table and flag the last version
  - Or use an history table and keep only last version of deals
- Difficulty to identify modified fields
- Must be managed by application (RDBMS doesn't help here)
- What-If data must be accessed by different systems?

Last Version Flag, Previous field

Modified Fields

ID	Ref	Previous	LastVersion	EventType	EventDate	UserCre	UserMod	TradeDate	Counterpart	Way	MMType	ValueDate	MaturityDate	Notional	Ccy	Index	Rate	Calendar
1127	AA001		No	Create	13/1/15 10:10	PBI	PBI	13/01/15	GOOGLE	Buy	Deposit	15/01/15	15/06/15	10 000 000	EUR	FIXED	2,5	TGT
2356	AA001	1127	No	Modify	13/1/15 16:50	PBI	RST	13/01/15	GOOGLE	Buy	Deposit	15/01/15	15/06/15	10 000 000	JPY	FIXED	2,5	TGT
3628	AA001	2356	Yes	Modify	14/1/15 11:25	PBI	LTO	13/01/15	GOOGLE	Buy	Deposit	15/01/15	15/06/15	10 000 000	JPY	FIXED	2,5	JPY+TGT

## Cell Level access

- Now consider storage not as a static array of data but as a set of cells
- Cell can be accessed and modified independently
- Cells are versioned (timestamp). Each revision can be retrieved easily by API


Qualifier	Reference	Event Type	Event Date	UserCre	UserMod	TradeDate	Ctpy	Way	MMType	ValueDate	MaturityDate	Notional	Ccy	Index	Rate	Calendar
Timestamp	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375	42007969375
Value	AA001	Create	13/1/15 10:00	PBI	PBI	13/01/15	GOOGLE	Buy	Deposit	15/01/15	15/06/15	10 000 000	EUR	FIXED	2,5	TGT
Timestamp	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972	42017690972
Value		Modify	13/01/15 16:35		RST								JPY			
Timestamp	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694	42018475694
Value			14/1/15 11:25		LTO											JPY+TGT

Modified Fields

# Trade “Columnar” Representation

RowKey	Column Family	Column Qualifier	Timestamp	Value
Reference:AA001	info	EventType	42007969375	Create
			<b>42017690972</b>	<b>Modify</b>
		EventDate	42007969375	13/1/15 10:00
			<b>42017690972</b>	13/1/15 16:35
			<b>42018475694</b>	14/1/15 11:25
		UserCre	42007969375	PBI
		UserMod	42007969375	PBI
			<b>42017690972</b>	<b>RST</b>
			<b>42018475694</b>	<b>LTO</b>
		TradeDate	42007969375	13/01/15
		Counterpart	42007969375	GOOGLE
		Way	42007969375	Buy
		MMType	42007969375	Deposit
		ValueDate	42007969375	15/01/15
		MaturityDate	42007969375	15/06/15
		Notional	42007969375	10000000,00
		Ccy	42007969375	EUR
			<b>42017690972</b>	<b>JPY</b>
		Index	42007969375	FIXED
		Rate	42007969375	2,5
		Calendar	42007969375	TGT
			<b>42018475694</b>	<b>TGT+JPY</b>

Column version:  
By default, only last version is return




## Row atomicity

- All KVs for a “row” are co-located in a region
- Locks are per row
- Stored in-memory at region server

# Schema Less Model

RowKey	Column Family	Column Qualifier	Timestamp	Value
Reference:AA001	info	EventType	42007969375	Create
		EventDate	42007969375	13/01/15
		TradeDate	42007969375	13/01/15
		Counterpart	42007969375	GOOGLE
		Way	42007969375	Buy
		MMType	42007969375	Deposit
		ValueDate	42007969375	15/01/15
		MaturityDate	42007969375	15/06/15
		Notional	42007969375	10000000
		Currency	42007969375	EUR
		Index	42007969375	FIXED
		Rate	42007969375	2,5
		Reference:DER010	info	Way
Trader	42139416667			LRI
EventType	42139416667			Exercise
TradeDate	42139416667			"2015-01-13"
Counterpart	42139416667			HUAWEI
OptionType	42139416667			Call
ExerciseDate	42139416667			15/01/15
ExpiryDate	42139416667			"2015-01-30"
StrikePrice	42139416667			98
Currency	42139416667			USD
UnderlyingPrice	42139416667			114,45
Premium	42139416667			15,4
Quantity	42139416667			10

No constraint on column qualifier or value per row



- Data is ordered by Row Key
- Sharding is managed by Hbase
- Proper Row Key design is a critical for system performance (latency/throughput on Read/Write usage)

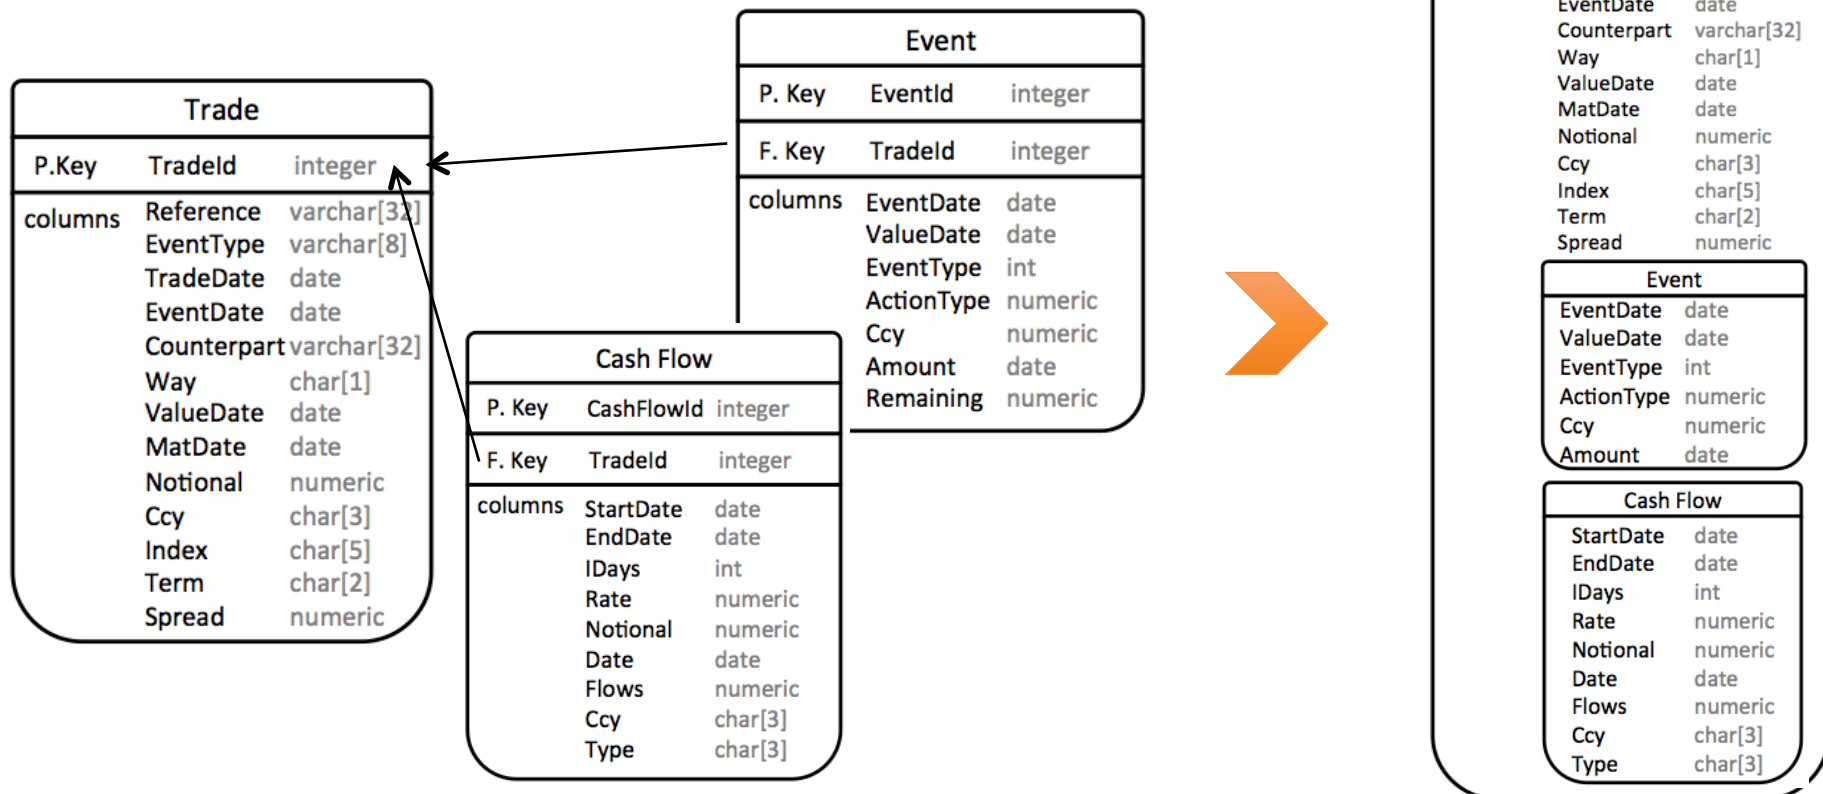
# How to manage relationship between entities? (1/3)

HBase has no foreign keys, or joins, or cross-table transactions.

- This can make representing relationships between entities ... tricky.
- HBase columns can be defined at runtime. A set of dynamically named columns can represent another entity!
- If you put data into the column name, and expect many such columns in the same row, then logically, you've created a nested entity.

Classic use case in finance: Trade, Trade Events and Cash Flows

- Most of the time you have 1-n relationship so nested entities could work just fine.
- But you don't have aggregate functions (avg, sum,...) any more.





# How to manage relationship between entities? (2/3)

## RDBMS representation of entities association

- Update operations trigger transaction on multiple tables

Trade Table

TradeID	Reference	EventType	TradeDate	EventDate	Counterpart	Way	ValueDate	MaturityDate	Notional	Ccy	Index	Term	Spread
1	GU002	Create	22/10/14	22/10/14 11:50	HSBC	Lend	25/10/14	25/01/15	5 333 581,33	EUR	EIBOR	3M	10
2	GU002	Partial Payoff	23/12/14	23/12/14 15:47	HSBC	Lend	25/10/14	25/01/15	5 333 581,33	EUR	EIBOR	3M	10

Cash Flow Table

CashFlowId	TradeId	StartDate	EndDate	IDays	Rate	Notional	Date	Flows	Type	Ccy
1	1			0	0	0	25/10/14	-5333581,33	XNL	EUR
2	1	25/10/14	25/01/15	92	2,145	5333581,33	25/01/15	29236,91	INT	EUR
3	1			0	0	0	25/01/15	5 333 581,33	XNL	EUR

Event Table

EventId	TradeId	Event Date	Value Date	Event Type	Action Type	Currency	Amount	Remaining Amount
1	2	23/12/14	25/12/14	XNL	Partial Payoff	EUR	211157,00	4635542,07
2	2	23/12/14	25/12/14	FEE	Partial Payoff	EUR	22,86	



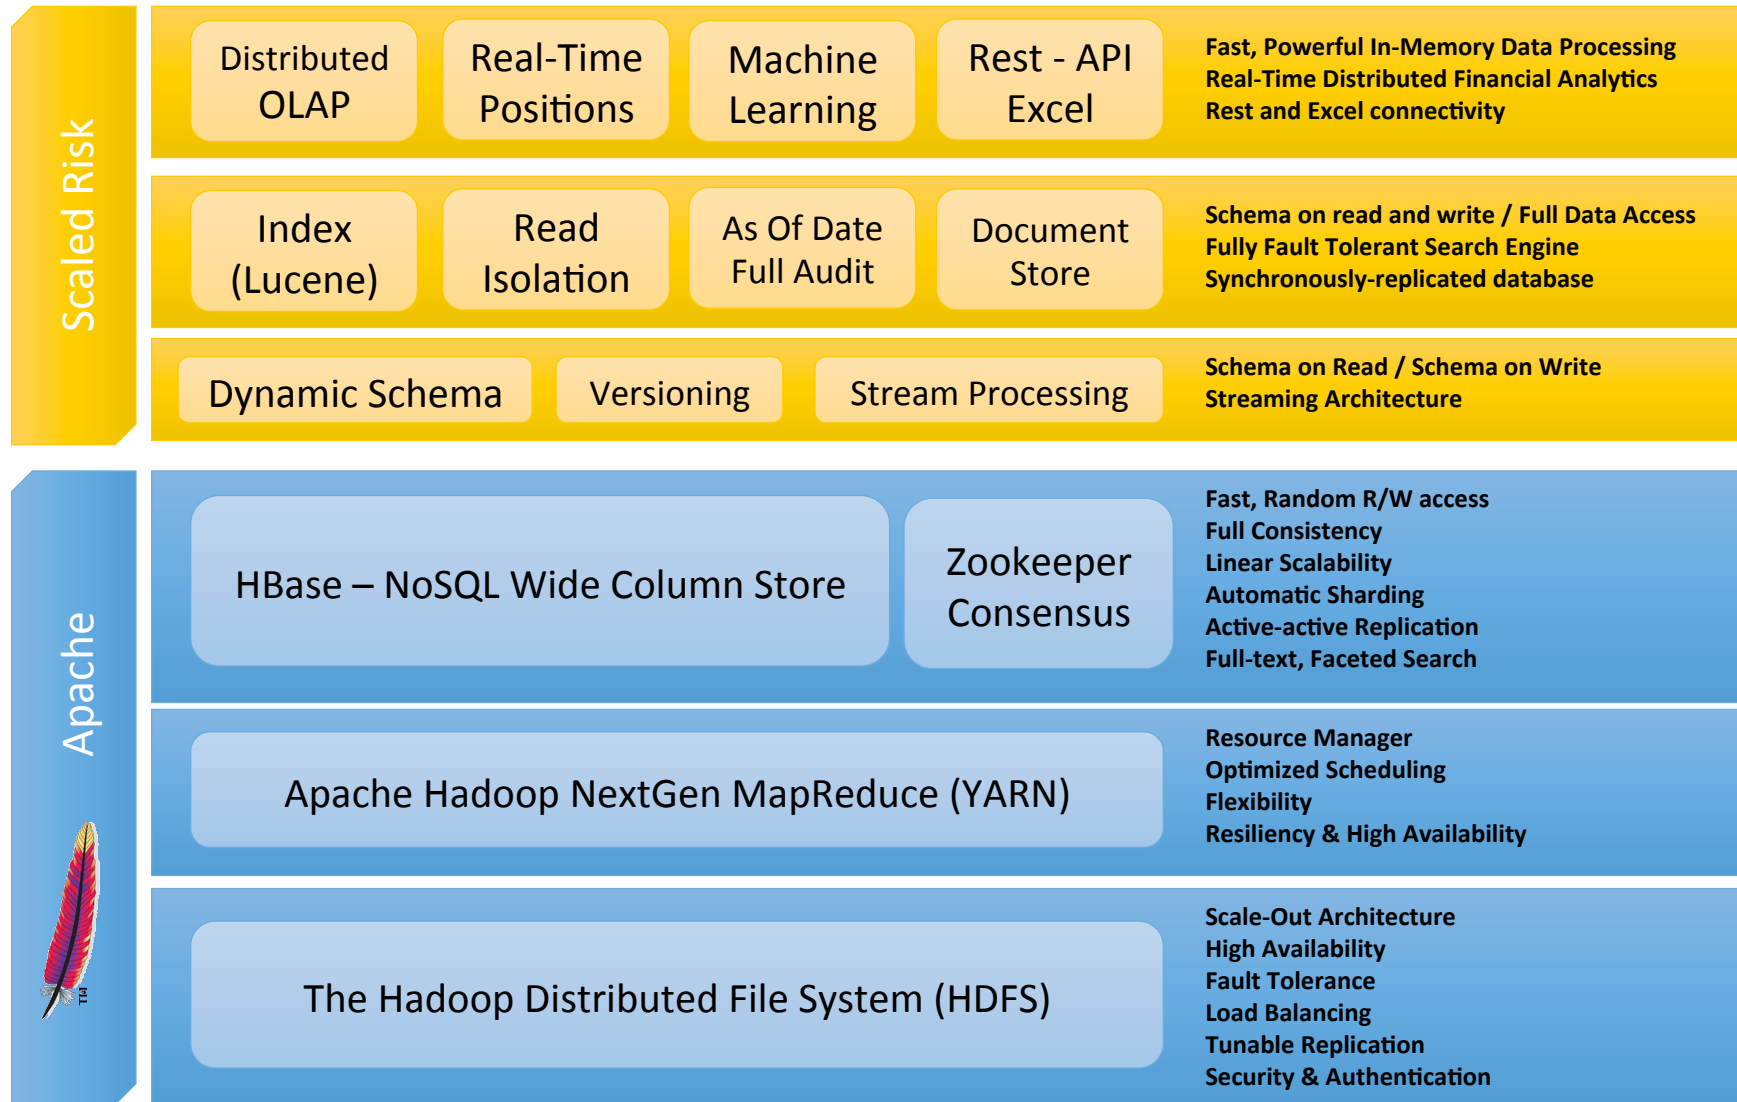
# How to manage relationship between entities? (3/3)

Row Key	Column Family	Column Qualifier	Timestamp	Value
Reference:GU002	info	EventType	42139416667	Create
		TradeDate	42139416667	22/10/14
		EventDate	42139416667	22/10/14
		Counterpart	42139416667	HSBC
		Way	42139416667	Lend
		ValueDate	42139416667	25/10/14
		MaturityDate	42139416667	25/01/15
		Notional	42139416667	5333581,33
		Ccy	42139416667	EUR
		Index	42139416667	EIBOR
		Term	42139416667	3M
		Spread	42139416667	10
		cashflow	CashFlow_1_IDays	42139416667
	CashFlow_1_Rate		42139416667	0
	CashFlow_1_Notional		42139416667	0
	CashFlow_1_Date		42139416667	25/10/14
	CashFlow_1_Flows		42139416667	-5333581,33
	CashFlow_1_Type		42139416667	XNL
	CashFlow_1_Ccy		42139416667	EUR
	CashFlow_2_StartDate		42139416667	25/10/14
	CashFlow_2_EndDate		42139416667	25/01/15
	CashFlow_2_IDays		42139416667	92
	CashFlow_2_Rate		42139416667	2,145
	CashFlow_2_Notional		42139416667	5333581,33
	CashFlow_2_Date		42139416667	25/01/15
	CashFlow_2_Flows		42139416667	29236,91
	CashFlow_2_Type		42139416667	INT
	CashFlow_2_Ccy		42139416667	EUR
	CashFlow_3_...		42139416667	...
	info/event	Event_1_EventDate	42018475694	23/12/14
		Event_1_ValueDate	42018475694	25/12/14
		Event_1_EventType	42018475694	XNL
		Event_1_ActionType	42018475694	Partial Payoff
		Event_1_Currency	42018475694	EUR
		Event_1_Amount	42018475694	211157
		Event_1_RemainingAmount	42018475694	4635542,07
		Event_2_...	42018475694	...

Column Families can be used to separate storage and semantics

Column Qualifiers are used to identify nested entities.  
[Type]\_[id]\_[Attribute]

# Scaled Risk - Product Features



# What's in for me ?

## Data Scientists / IT Quant - Compliance


- Strong Mathematics & statistics skills
- Modeling, Synthesis, Visualization (Story Telling)
- Language: R, Python, Scala
- Query: Hive / SparkSQL
- Artificial Intelligence, Machine Learning

## Big Data Developer / FO – Risk

- Strong programming skills & Optimization
- Java, Scala & others Functional Programming
- Yarn, Spark, HBase, Map Reduce
- Reactive Programming
- Stream Processing
- JSON, XML, CSV

## Hadoop Cluster administrator

- Infra As Code, Puppet, Chef, Automation
- Cloud Computing
- Cloudera, Hortonworks
- Java / JVM Optimization
- Linux
- Kerberos





**Thank You / Questions ? / Contacts  
And Yes, we are hiring!!**

Pierre Bittner – [pbittner@finaxys.com](mailto:pbittner@finaxys.com) /  @BittnerPierre

RH – [rh@finaxys.com](mailto:rh@finaxys.com) /  @Finaxys

