

Jan Valdman

14.5.2013




Big data

Přednáška KIV/SI č.11



Úvod do big data

▶ Exponenciální růst



Source: IDC 2011 Digital Universe Study (<http://www.emc.com/collateral/demos/microsites/emc-digital-universe-2011/index.htm>)

▶ Příklady

- Twitter (12TB tweetů)
- Monitoring telefonních hovorů v Číně
- Fraud analýza během autorizace transakce platební kartou
- Data ze „smart“ online senzorů – odběry energií, doprava, telemetrika...
- Audio-video data, rozpoznání osob na CCTV
- Prediktivní (kyber)bezpečnostní analýza

► Výzvy, které přináší big data

- Dat je moc
- Nejde udělat globální dotaz
- Nejde analyzovat všechna data
- Nejde dělat analýzu v reálném nebo rozumném čase

Proč?

- Big data
 - Big users
 - Cloud computing
-
- Příklad: úspěšná virální aplikace na Appstore nebo Google Play

► Big data – what does it mean?

Volume


Big data comes in one size: large. Enterprises are awash with data, easily amassing terabytes and even petabytes of information.

TB, Records, Transactions, Tables,, Files

Velocity

Often time-sensitive, big data must be used as it is streaming in to the enterprise in order to maximize its value to the business.

Batch, Near time, Real time, Streams



Variety

Big data extends beyond structured data, including semi-structured and unstructured data of all varieties: text, audio, video, click streams, log files and more.

Structured, Unstructured, Semistructured

Veracity

Quality and provenance of received data

Good, Bad, Undefined, Inconsistency, Incompleteness, Ambiguity

► Charakteristika big data – 4 × „V“

- Volume
 - Dat je mnoho a jejich zpracování trvá příliš dlouho
 - Twitter, uzávěrky obchodování ...
- Velocity
 - Obrovské množství dat je potřeba analyzovat v krátkém čase
 - Detekce podvodů, optimalizační úlohy
- Variety
 - Data mají různou strukturu a jsou nestructurovaná: video, senzorická data, text, ...
 - Rozpoznání požadované situace (události) v živém videostramu
- Veracity
 - Rostoucí množství zdrojů dat
 - Problematická důvěra v data

Příležitost

- Business insight
- Netušené možnosti, nové souvislosti
- Online podklady k rozhodování
- Odpovědi na otázky, které dříve nešlo odpovědět
- Konkurenční výhoda

Využití bigdata

- **Digital marketing optimization** (e.g., web analytics, attribution, golden path analysis)
- **Data exploration and discovery** (e.g., data scientists, identifying new data-driven products, new markets)
- **Fraud detection and prevention** (e.g., revenue protection, site integrity & uptime)
- **Social network and relationship analysis** (e.g., influencer marketing, outsourcing, attrition prediction)
- **Machine-generated data analytics** (e.g., Remote device insight, remote sensing, location-based intelligence)
- **Data retention** (e.g. long term conservation, Data archiving)

► Proč big data

Finance: better and deeper understanding of risk to avoid credit crisis – Bale III

Telecommunications: more reliable network where we can predict and prevent failure

Media: more content that is lined up with your personal preferences

Life Science: better targeted medicines with fewer complications and side effects


Retail: a personal experience with products and offers that are just what you need

Government: government services that are based on hard data, not just gut

Architektura big data

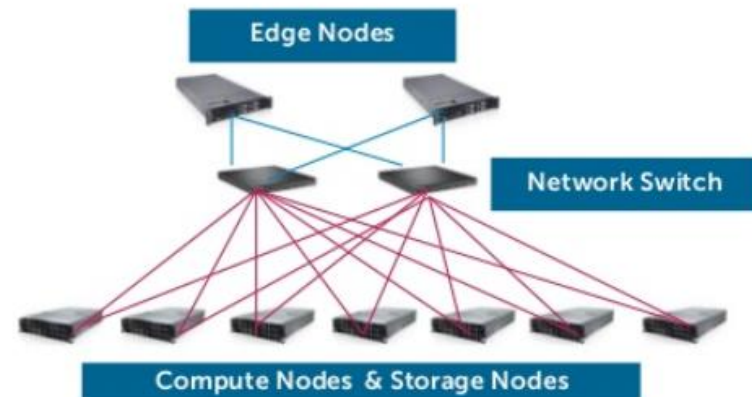
Traditional Architecture

- High availability
- Enterprise database
- Right design for structured data
- Current storage hardware (SAN, NAS, DAS)



Big Data Architecture

- High scalability, availability and flexibility
- Compute and storage in the same box for reducing the network latency
- Right design for structured, semistructured and unstructured data




▶ IBM Watson



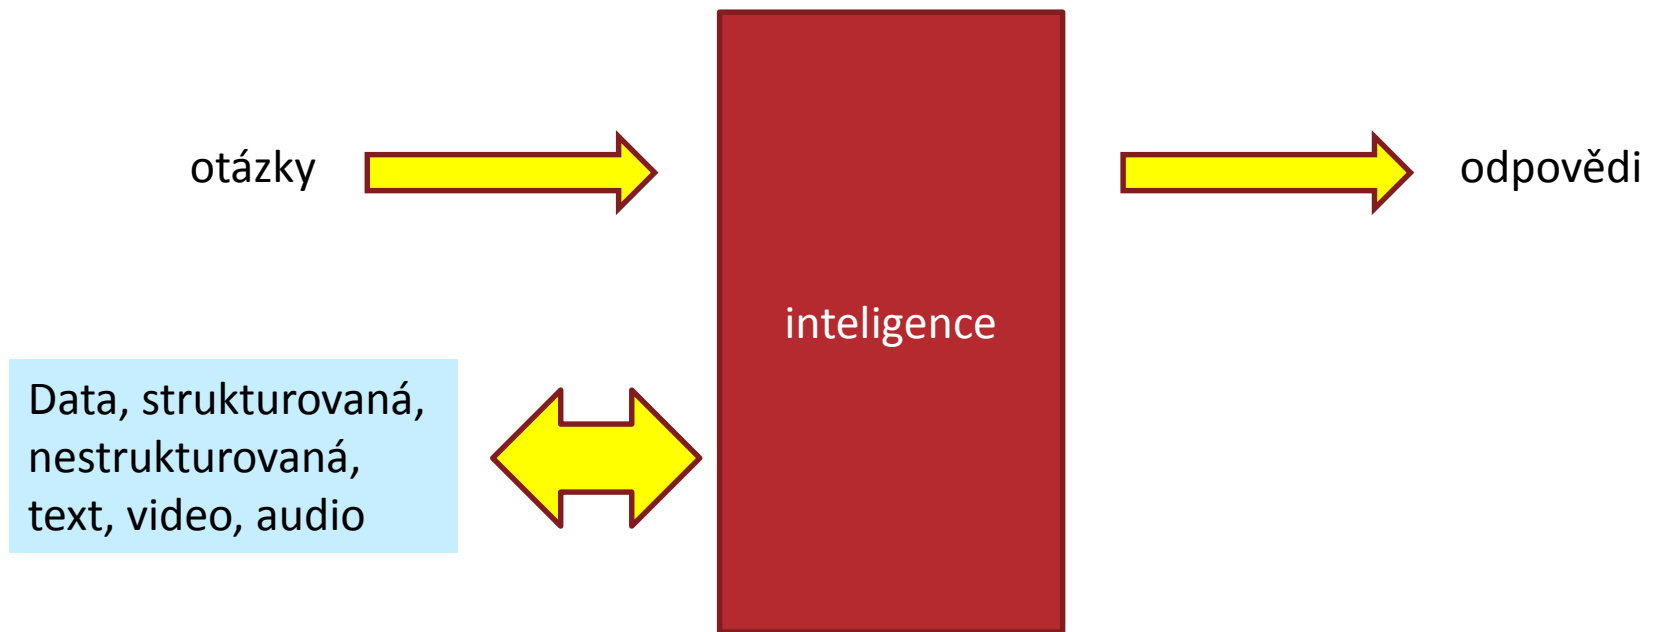
▶ IBM Watson



- Hra „Jeopardy“ (Riskuj)
- V roce 2011 stroj (Watson) vyhrál nad nejlepšími lidmi
- Otázky i odpovědi v přirozeném jazyce, porozumění, fráze, dvojsmysly...
- Bez přístupu na Internet
- Masivně paralelní systém, 90x Power 750 (4 octacore):
2880 POWER7 cores, 16TB RAM, 4TB storage



www.ibm.com/watson



IBM big data platform

The key platform capabilities

- **Hadoop-based analytics:** Processes and analyzes any data type across commodity server clusters.
- **Stream Computing:** Drives continuous analysis of massive volumes of streaming data with sub-millisecond response times.
- **Data Warehousing:** Delivers deep operational insight with advanced in-database analytics.
- **Information Integration and Governance:** Allows you to understand, cleanse, transform, govern and deliver trusted information to your critical business initiatives.

Supporting platform services

- **Accelerators:** Faster time to value with pre-packaged analytical and industry-specific content.
- **Application Development:** Streamline the process of developing big data applications.
- **Information Integration and Governance:** Integrate, protect, cleanse, govern, and deliver your trusted information
- **Systems Management:** Monitor and manage your big data system for secure and optimized performance.
- **Reference Architectures:** Hardware, networking and system software blueprints to accelerate time to value.

▶ IBM Produkty pro big data

- **InfoSphere Streams** Enables continuous analysis of massive volumes of streaming data with sub-millisecond response times.
- **InfoSphere BigInsights** An enterprise-ready, Apache Hadoop-based solution for managing and analyzing massive volumes of structured and unstructured data.
- **InfoSphere Data Explorer** Discovery and navigation software that provides real-time access and fusion of big data with rich and varied data from enterprise applications.
- **IBM PureData powered by Netezza technology** Simplifies and optimizes analytical applications, enabling very complex algorithms to run in minutes not days.
- **DB2 with BLU Acceleration** Advanced, innovative capabilities to accelerate analytic workloads for databases and data warehouses.
- **IBM Smart Analytics System** Data management, hardware, software, & services capabilities that modularly delivers a wide assortment of business changing analytics.
- **InfoSphere Master Data Management** Creates trusted views of your master data for improving your applications and business processes.
- **InfoSphere Information Server** Understand, cleanse, transform and deliver trusted information to your critical business initiatives, integrating big data into the rest of your IT systems.

A large, solid blue triangle pointing to the right, positioned on the left side of the slide.

NoSQL databáze

NoSQL

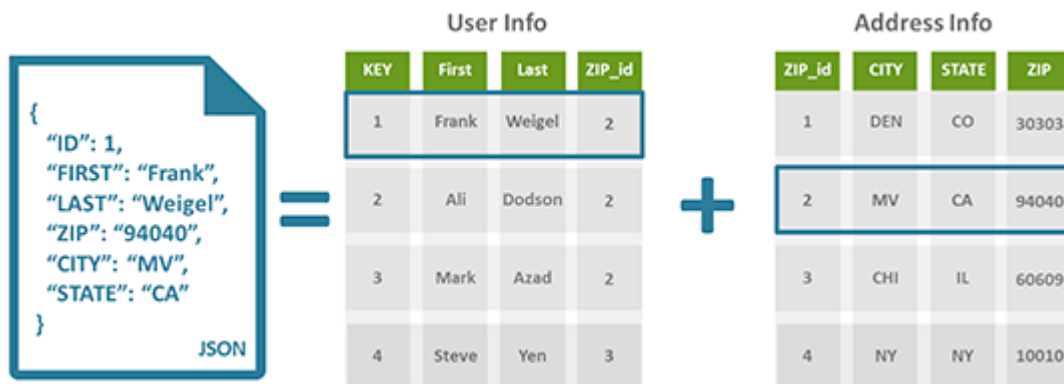
- Obecný název pro řadu databázových aplikací zaměřených na vysoký výkon a škálovatelnost pro zpracování velkých objemů dat
- Jednoduché použití
- Horizontální škálování

Hlavní kategorie

- Key-value – Redis, Riak
- Columns – Cassandra, Hbase
- Graph – InfiniteGraph, Neo4J, FlockDB
- Document – MongoDB, CouchDB

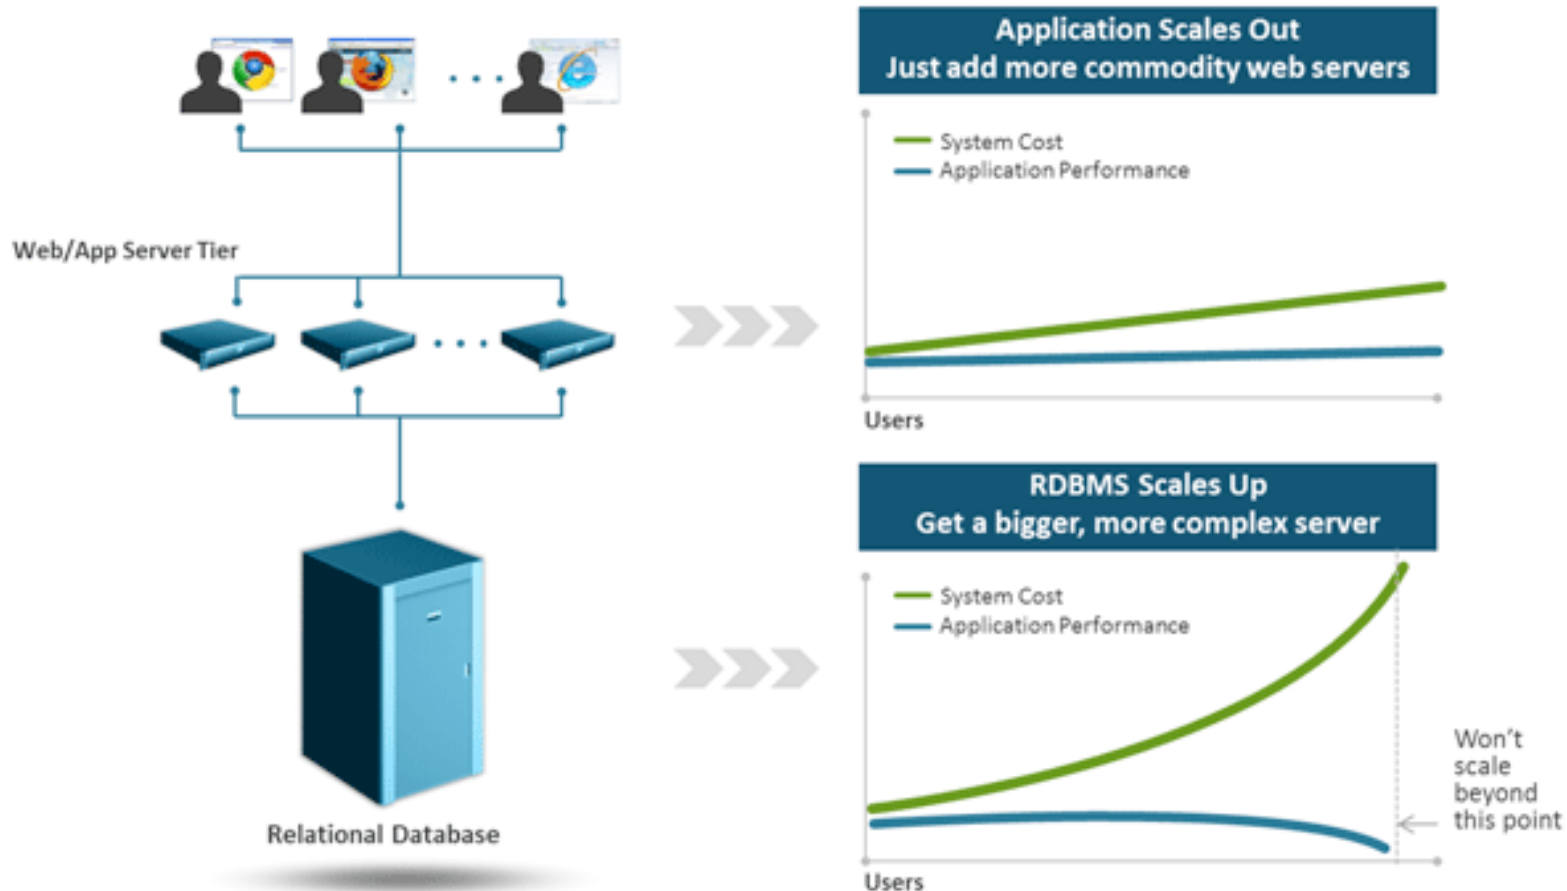
▶ NoSQL databáze pracují na odlišném principu než RDBMS

- NoSQL databases have a very different model. For example, a document-oriented NoSQL database takes the data you want to store and aggregates it into documents using the JSON format.



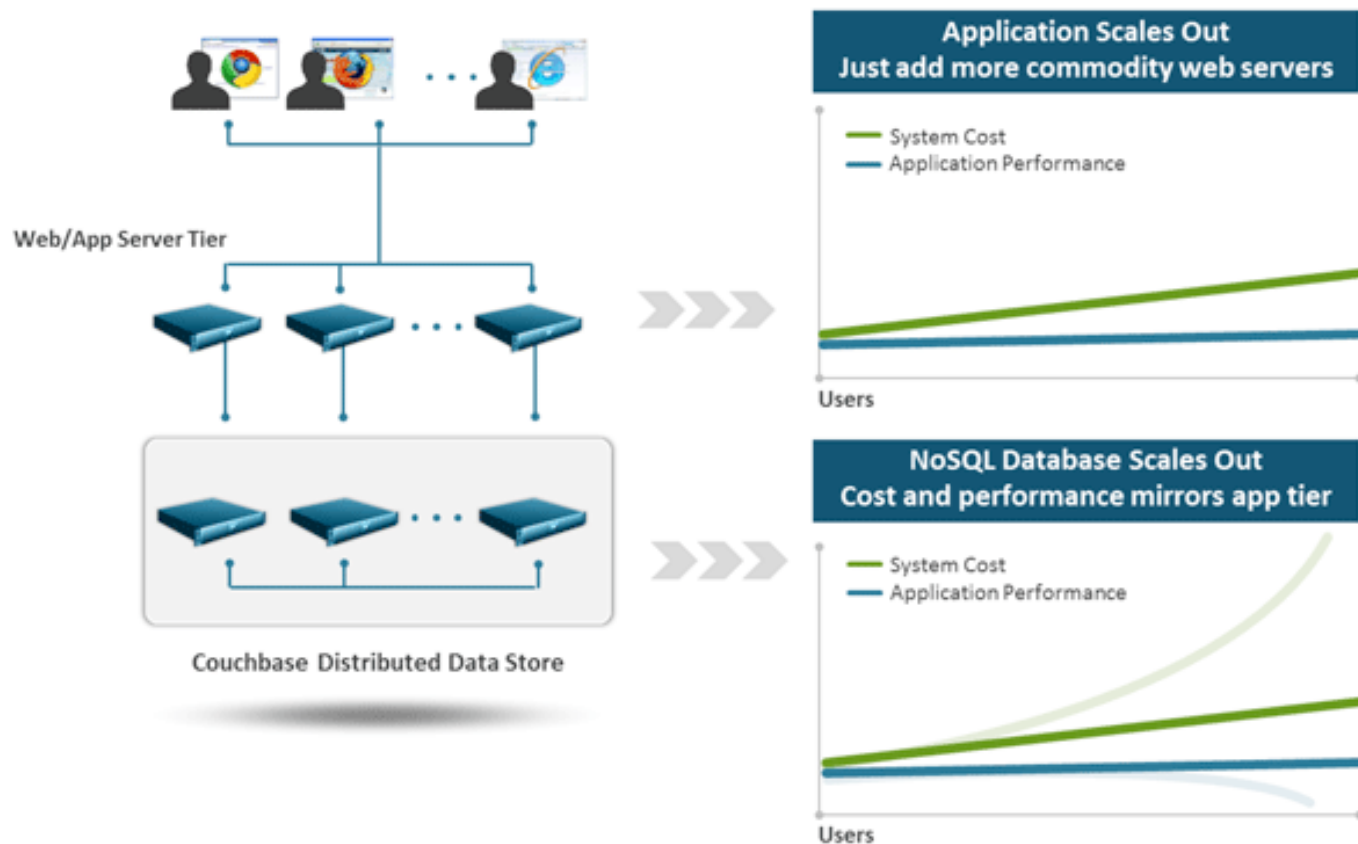
► Škálování SQL databáze

- Škálování relační databázové vrstvy má své limity:



► Šálování NoSQL databáze

- NoSQL databáze se škálují jako distribuovaný systém (Couchbase)



► Existuje asi 150 (!) NoSQL databází

<http://nosql-database.org/>

- Wide column store Hadoop
- Document store MongoDB Crouchbase CrouchDB
- Key-value store (n-tuple) Berkeley DB, OpenLDAP
- Graph database
- Multimodel databases
- Object databases
- Grid databases
- XML databases EMC xDB, IBM DB2
- Multivalue databases
- Event sourcing
- Other IBM Domino,

▶ ACID

- Atomicity
- Consistency
- Isolation
- Durability

▶ HBase

- Clone of Big Table (Google)
- Implemented in Java (Clients : Java, C++, Ruby...)
- Data is stored “Column-oriented”
- Distributed over many servers
- Tolerant of machine failure
- Layered over HDFS
- Strong consistency
- It's not a relational database (No joins)
- Sparse data – nulls are stored for free
- Semi-structured or unstructured data
- Data changes through time
- Versioned data
- Scalable – Goal of billions of rows x millions of columns

Table

Row	Timestamp	Animal		Repair
		Type	Size	Cost
Enclosure1	12	Zebra	Medium	1000€
	11	Lion	Big	
Enclosure2	13	Monkey	Small	1500€

Region

Key

Column



Family

Cell

(Table, Row_Key, Family, Column, Timestamp) = Cell (Value)

▶ Cassandra

- Every node play the same role
- Highly Available
- Really fast reads, really fast writes
- Flexible schemas
- Distributed, Replicated
- No Master, no Slaves
- No Single Point of Failure
- Client can talk to any node
- Written in Java



▶ Cassandra

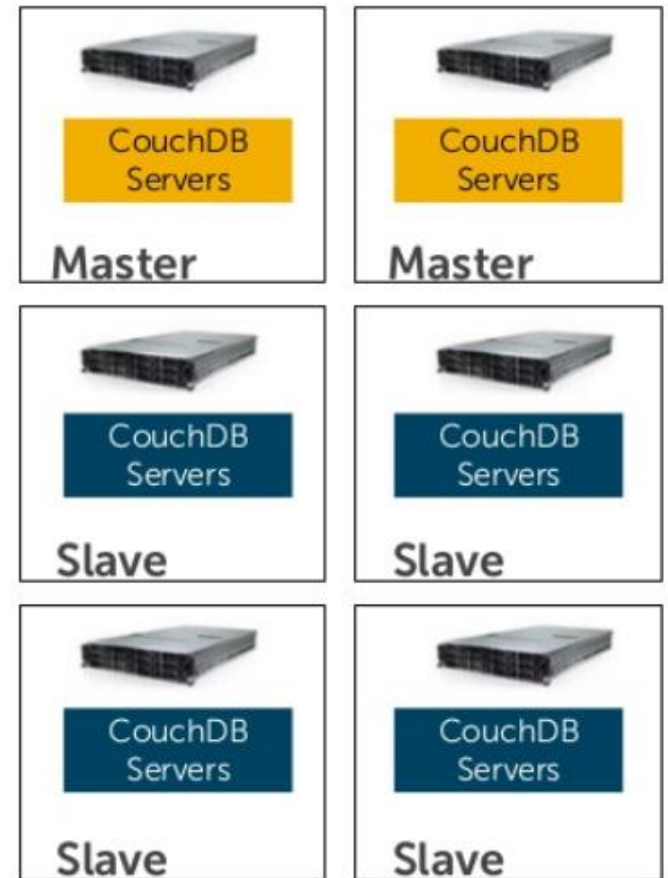


Column
+Name
+Value
+Timestamp

- Column Family
 - Think of it as a DB table
- Column
 - Key-Value Pair (not just a value, like a DB column)
 - Timestamp
- SuperColumn
 - Columns inside a column
 - The value are columns
 - No timestamp
- Keyspace – like a namespace, generally 1 per app
- Indexes
- Queries

▶ CouchDB

- Open Source Distributed Database
- RESTful API
- Schema-less document store (document in JSON format)
- Multi-Version-Concurrency-Control model
- User-defined query structured as map/reduce
- Incremental Index Update mechanism
- Multi-Master Replication model
- Written in Erlang
- Support MapReduce
- Easy to use data storage
- Easy to integrate with web applications : JavaScript, JSON
- Scalability for large web applications : Incremental Replication, bi-directional conflict detection and management
- Query-able and index-able
- Offline by default



- Master → Slave replication
- Master ↔ Master replication
- Filtered Replication
- Incremental and bi-directional replication
- Conflict management

CouchDB

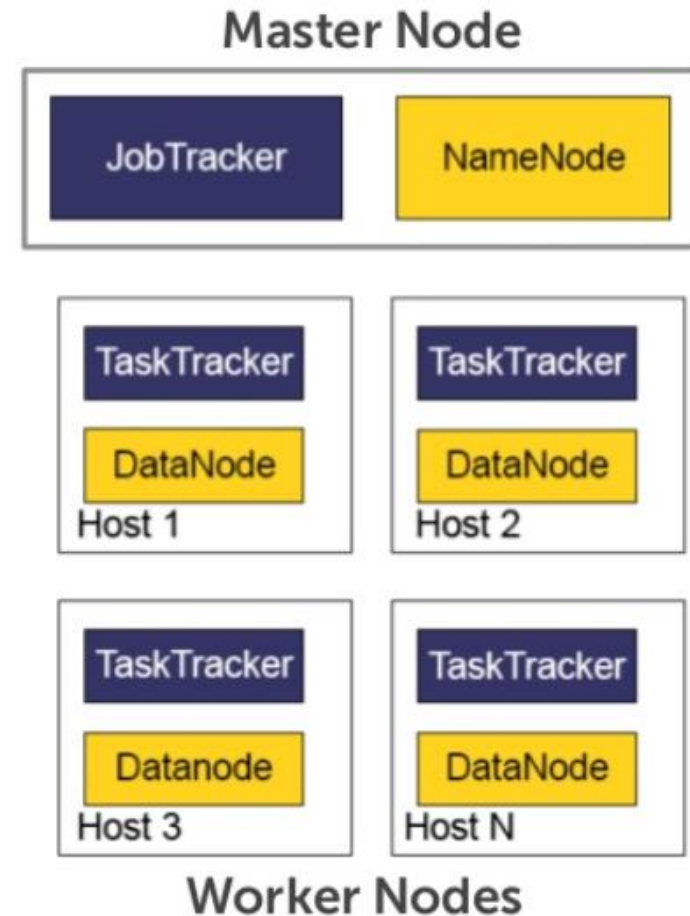
- Document storage
 - CouchDB server hosts named databases, which store documents
- ACID Properties
 - CouchDB never overwrites committed data or associated structures, ensuring the database file is always in a consistent state
- Compaction
 - On schedule, or when the database file exceeds a certain amount of wasted space, the compaction process clones all the active data to a new file and then discards the old file
- Views (Model, Function, Index)
 - View model is the method of aggregating and reporting on the documents in a database, and are built on-demand to aggregate, join and report on database documents
 - View function takes a CouchDB document as an argument and then does whatever computation it needs to do to determine the data that is to be made available through the view, if any. It can add multiple rows to the view based on a single document, or it can add no rows at all
 - View index is a dynamic representation of the actual document contents of a database, and CouchDB makes it easy to create useful views of data. But generating a view of a database with hundreds of thousands or millions of documents is time and resource consuming, it's not something the system should do from scratch each time
- Security
 - To protect who can read and update documents, CouchDB has a simple reader access and update validation model that can be extended to implement custom security models
- Distributed update and replication
 - CouchDB is a peer-based distributed database system, it allows for users and servers to access and update the same shared data while disconnected and then bi-directionally replicate those changes later



Hadoop

Hadoop

- **Hadoop Distributed File System**
 - A scalable, Fault tolerant, High performance distributed file system
 - Asynchronous replication
 - Write-once and read-many (WORM)
 - Hadoop cluster with 3 DataNodes minimum
 - Data divided into blocks, each block replicated 3 times (default)
 - No RAID required for DataNode
 - Interfaces: Java, Thrift, C Library, FUSE, WebDAV, HTTP, FTP
 - **NameNode** holds filesystem metadata
 - Files are broken up and spread over the **DataNodes**
- **Hadoop MapReduce**
 - Software framework for distributed computation
 - Input | Map() | Copy/Sort | Reduce() | Output
 - **JobTracker** schedules and manages jobs
 - **TaskTracker** executes individual map() and reduce() tasks on each cluster node



▶ IBM Netezza

- Přináší zákazníkům rychlé, jednoduché a efektivní řešení v případě potřeby vysokorychlostních analytických dotazů
- Silný nástroj pro analytické dotazy, konkurence pro Oracle a Exadata
- Appliance SW+HW bundle

