



Red Hat


Transakce ...tam a zase zpátky

Ondřej Chaloupka

jOpenSpace 2019

Jsou transakce zlo?

Databáze



- Principles of Transaction-Oriented Database Recovery, Andreas Reuter and Theo Härde, 1983

Transakce jsou...

**abstrakce pro jednoduchou práci
s daty**

Atomicity

Isolation

A C I D

Consistency

Durability

Atomicity == vrácení dat do původního stavu

ACID

Consistency == data jsou konzistentní

Durability == data jsou bezpečně uloženy

A **C** **I** **D**

Isolation == data je možné zpracovávat paralelně

ACID

Serializability!

Lock-based



ID	Jméno	Příjmení
1	Zdeněk	Henek
2	Jan	Novotný

Lock-based for serializability


ID	Jméno	Příjmení
1	Zdeněk	Hánek
2	Jan	Němčický



Lock-based isolation levels



Isolation level	Dirty reads	Lost updates	Non-repeatable reads	Phantoms
Read Uncommitted	may occur	may occur	may occur	may occur
Read Committed	don't occur	may occur	may occur	may occur
Repeatable Read	don't occur	don't occur	don't occur	may occur
Serializable	don't occur	don't occur	don't occur	don't occur

Databáze a transakce




- First version of ANSI SQL was in 1987
- A Critique of ANSI SQL Isolation Levels, Hal Berenson, Phil Bernstein 1995

Lock-based vs. MVCC




CAP a distribuované databáze



- Brewer's CAP Theorem, 2000
- [The Part-Time Parliament](#), Leslie Lamport, 1998
- [Paxos Made Simple](#), Leslie Lamport, 2001


CAP theorem



CAP Consistency != ACID consistency

Availability over Consistency


NoSQL era




- [Amazon's dynamo paper](#), Amazon, 2007

Consistency strikes back

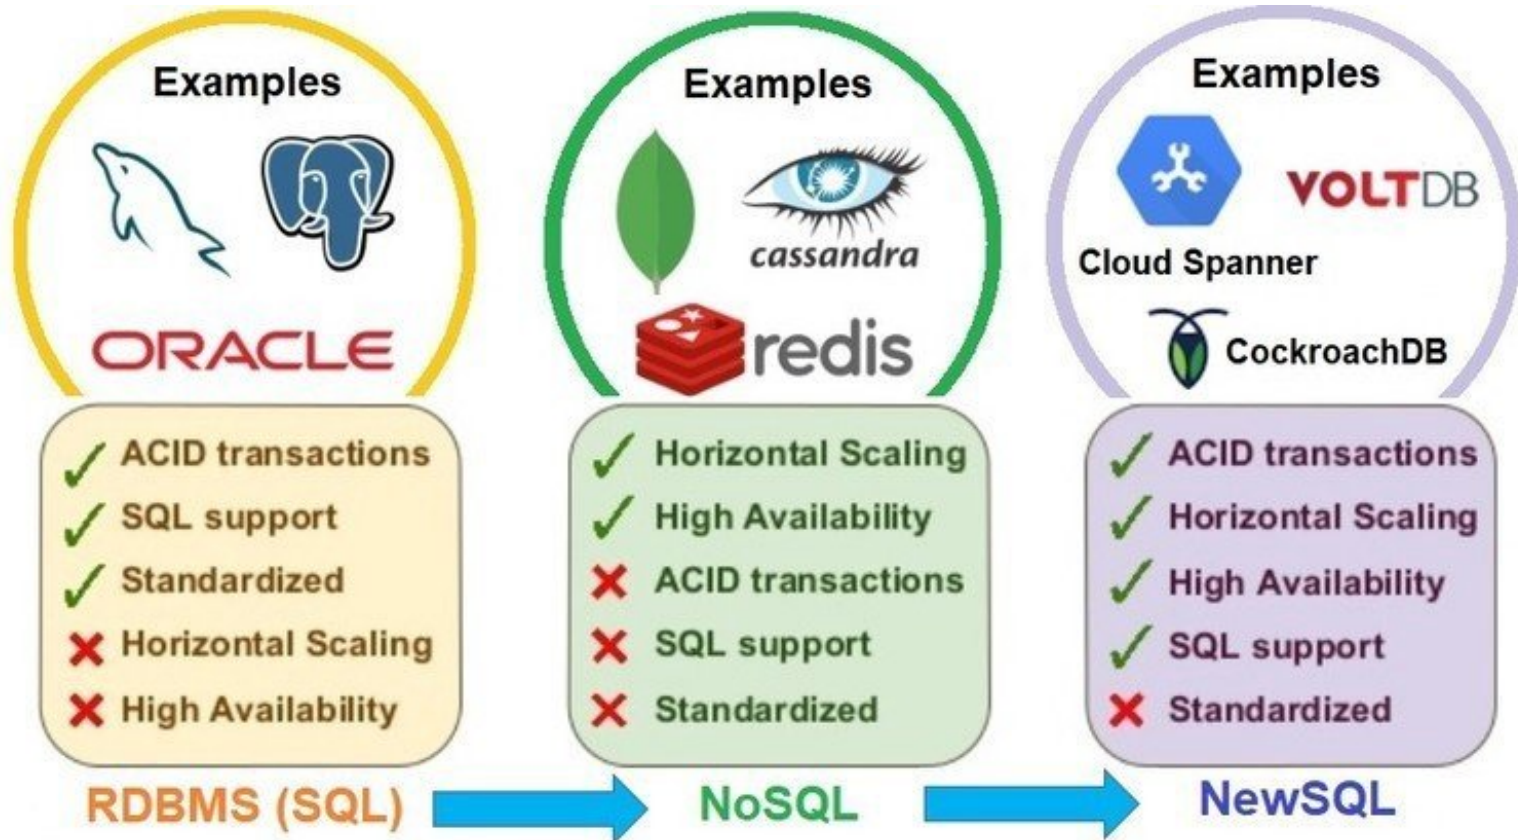
NewSQL era




- [CRDT: Conflict-free replicated data types](#), Shapiro, 2011
- [Spanner](#), Google, 2012
- [Calvin](#), A. Thomson, D. Abadi 2012
- [Jepsen.io analyses](#), Kyle Kingsbury



NoSQL to NewSQL




A co aplikace?



- [XA specification, 1991](#)
- [BASE, an ACID alternative, Ebay, 2008](#)
- [Why banks are BASE, 2013](#)

*"Dva největší problémy při programování jsou **invalidace cache**, pojmenování věcí a chyby o jedničku."*



Every sufficiently large deployment of
microservices

contains an ad-hoc, informally-
specified, bug-ridden, slow
implementation of half of


transactions



Sept 25-26, 2015

thestrangeloop.com

?



- [Saga](#) pattern/comensation transactions
- [Protocols for "client-side" distributed transactions](#)
- Event-sourcing message based approach ([Kafka](#) and friends)


DÍKY ZA POZORNOST



Change photo



CAL EXERCISE


Chaloupka-Juraneek almost-consensus algorithm with (usually) positive social side-effects:

- 1 Go to <https://sched.co/Jcj3>
- 2 Scroll down to Feedback form and click on this icon 
- 3 Talk to 5 randomly chosen people on the corridor during the conference
- 4 Introduce yourselves, explain them the algorithm and ask them to execute points 1.-4. (and continue chatting with the person)


Brno Distributed Systems Meetup Group

 Brno, Czech Republic

 41 members · Public group 

 Organized by **chalda** and 2 others

Share:   

[About](#)[Events](#)[Members](#)[Photos](#)[Discussions](#)[More](#)[Manage group](#) [Create event](#) 