

# A Class of CAP-like Theorems

You can't quote me on any of this.

# Talk outline

- Overview of the CAP theorem
- Aspects of the CAP theorem
- Alternative CAP Theorems:
  - relaxing consistency
  - changing the failure model
  - focus on transactions

# What is CAP?

- Consistency
  - System property
- Availability
  - system property
- Partition-Tolerance
  - failure model (?)
- Pick at most two.

# Which Consistency?

- linearizability - from the shared memory world.
- traditional “strong” consistency model
- all operations “appear” to happen at exactly one point
- wall-clock ordering is respected



# Which availability and Partition tolerance

- infallible per-node internet link
  - cannot use for system traffic
- potentially-eliminable host-pair links
  - must use for system traffic
- availability: all nodes must respond to all internet-link traffic in finite time.
- partitions: any host-pair link can be removed at any time.

# The variants

Many of these variants are real.

Several others are not my invention.

Try to pick out which are which!

# The Not-A theorem

- strengthen failure model: all links are fallible
- upon receipt of user request, all system links destroyed.
- availability failure! NOT A!

# The perfect-world theorem

- weaken failure model - infallible links!
- (this corresponds to taking CA from CAP)

# The SHOWER CAP theorem

- causal consistency
- partition-tolerance support

# The SHOWER CAP theorem

- causal consistency
- partition-tolerance  
support



# The SHOWER CAP theorem

- causal consistency
- partition-tolerance  
support





# Automata: the emissions CAP theorem

- define distributed system as Final Coalgebra
- failure model applies to transition matrix
- consistency model applies to emission Matrix
- CAP holds depending on emission Matrix.



# Trying again! HAT Theorem

- highly-available transactions
- common isolation models work!
- kinda!

Database	Default	Maximum
Action Ingres 10.0/10S [1]	S	S
Aerospike [2]	RC	RC
Akiban Persistit [3]	SI	SI
Clustrix CLX 4100 [4]	RR	RR
Greenplum 4.1 [8]	RC	S
IBM DB2 10 for z/OS [5]	CS	S
IBM Informix 11.50 [9]	Depends	S
MySQL 5.6 [12]	RR	S
MemSQL 1b [10]	RC	RC
MS SQL Server 2012 [11]	RC	S
NuoDB [13]	CR	CR
Oracle 11g [14]	RC	SI
Oracle Berkeley DB [7]	S	S
Oracle Berkeley DB JE [6]	RR	S
Postgres 9.2.2 [15]	RC	S
SAP HANA [16]	RC	SI
ScaleDB 1.02 [17]	RC	RC
VoltDB [18]	S	S

RC: read committed, RR: repeatable read, SI: snapshot isolation, S: serializability, CS: cursor stability, CR: consistent read

# Adapting: COWBOY HAT

- Let's adopt incompatible models anyway!
- what could go wrong?



# Adapting: COWBOY HAT

- Let's adopt incompatible models anyway!
- what could go wrong?



## In summary:

- CAP
- Not-A
- SHOWER CAP
- HAT
- COWBOY HAT

There are more!

- baseball CAP
- newsCAP
- winter CAP
- (winter soldier)  
CAP

# Evaluation

- Implemented in Fedora
- Results inconclusive
- one graph...

