

Scale Out and Conquer: Architectural Decisions Behind Distributed In-Memory Systems

Akmal Chaudhri
Technology Evangelist
GridGain Systems

© 2017 GridGain Systems, Inc.



Agenda

- Partitioning – Pitfalls of Even Distribution
- Affinity Collocation – Laid Out in Numbers
- Multi-threading of Distributed Systems – Things to Know

Partitioning – Pitfalls of Even Distribution

Where request goes?

PUT(K, V)

?



Node 1




Node 2

Affinity Function

Key



Partition



Naïve Function: What's Wrong?


0	2	4
6	8	10

Node 1

1	3	5
7	9	11

Node 2

Naïve Affinity: Redundant Partitions Re-shuffling!




Productized Affinity Functions

- Consistent hashing [1]
- Rendezvous hashing (HRW) [2]

[1] https://en.wikipedia.org/wiki/Consistent_hashing

[2] https://en.wikipedia.org/wiki/Rendezvous_hashing


Rendezvous Affinity: Distribute Evenly at Best




Affinity Collocation – Laid Out in Numbers

Running Transaction: No Collocation

UPDATE(id=1, city=MSK)
UPDATE(id=2, city=MSK)



Non Collocated Data: Number of Operations



2 (2 nodes)

2 (primary + backup)

2 (two-phase commit)

2 (request-response)

--

16 network operations

Running Transaction: Collocated Data!

UPDATE(id=1, city=MSK)
UPDATE(id=2, city=MSK)



0	3
6	9

Node 1


1	4
7	10

Node 2

2	5
8	11

Node 3

Collocated Data: Number of Operations



1 (1 node)

2 (primary + backup)

1 (one-phase commit)

2 (request-response)

--

4 network operations

Upshot

Upshot

- Partitioning – it's not about even distribution only
- Affinity Collocation – a golden concept of distributed systems
- Business Data Model should be adopted/reconsidered
- Thread-per-partition – speeds up most but not all usage scenarios

Q/A