Optimize Your Data with Temperature Awareness

Uma Maheswara Rao G

Intel Big Data Technologies

Contributors: Kai Zheng , Sammi Chen, Wei Zho, Malini Bhandaru

Flash Memory Summit

August 8–10, 2017
Santa Clara Convention Center

Uma Maheswara Rao G

- Senior Software Engineer at Intel Corporation
- Apache Software Foundation Member
- Committer and PMC for Apache Hadoop, Apache BookKeeper, Apache Mnemonic, and Apache Incubator
- Mentor for several Apache Incubator projects

@UmamaheswaraG

Outline

- Background
- HDFS Goodness & Problems
- HDFS Smart Storage Manager
- Wrap up

Data Richness

- Volume: Growing
- Useful Life: both long and short
- Unit Size: large and small
  (OLTP & OLAP)
- Formats : Many
  (Parquet, Apache ORC, Apache CarbonData)

BY 2020

AVG. INTERNET USER **1.5 GB** OF TRAFFIC / DAY

AUTONOMOUS VEHICLES **4 TB** OF DATA / DAY

CONNECTED AIRPLANE **5 TB** OF DATA / DAY

SMART FACTORY **1 PB** OF DATA / DAY

Storage Media and Software Options

Super Hot
Intel Optane, 3D CrossPoint
(ultra fast access but expensive)

HOT
PCIe-SSD

WARM
SATA-SSD, HDD

COLD
S3-Storage/Optical Disks
(slow but ultra cheap)

Intel SPDK!

Fast computing! Intel ISA-L

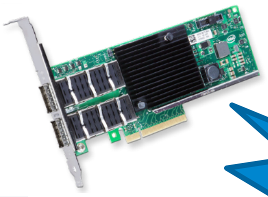

Faster NIC ! (20/50/100 Gb)

Emerging Compute Workloads — Deep Learning

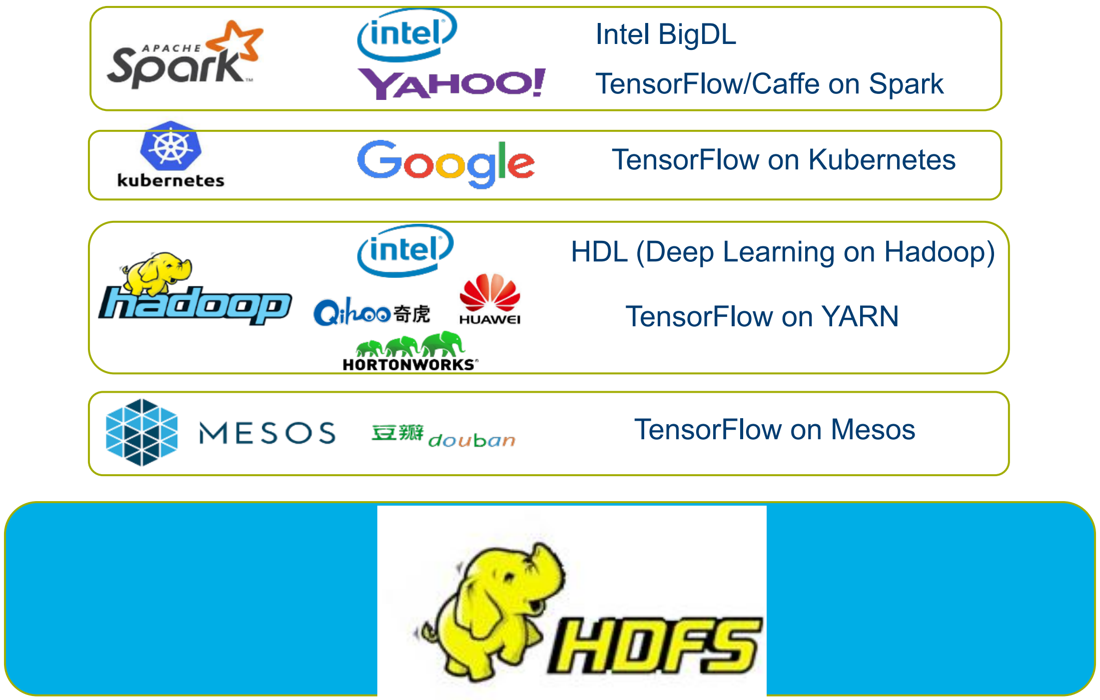

Outline

- Background
- HDFS Goodness & Problems
- HDFS Smart Storage Manager
- Wrap up

HDFS: Storage Overhead
Triple Replication versus Erasure Coding

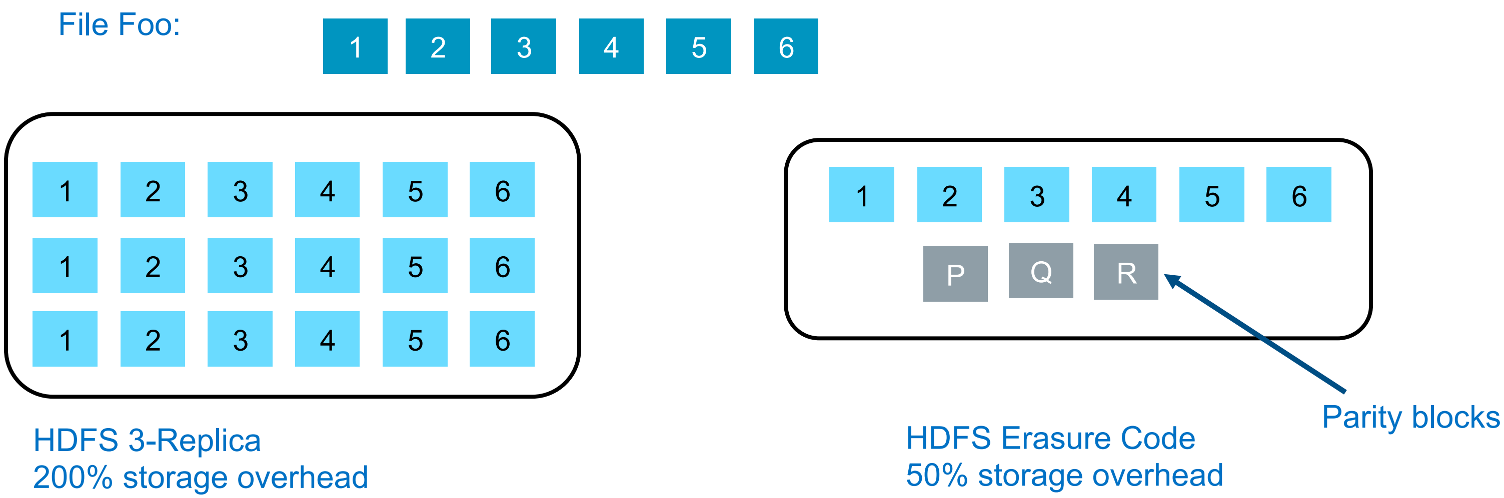

Flexible policy to have different kinds of data layouts.
Could they be dynamically applied??

HDFS: Storage Performance, Cost, Size Variation

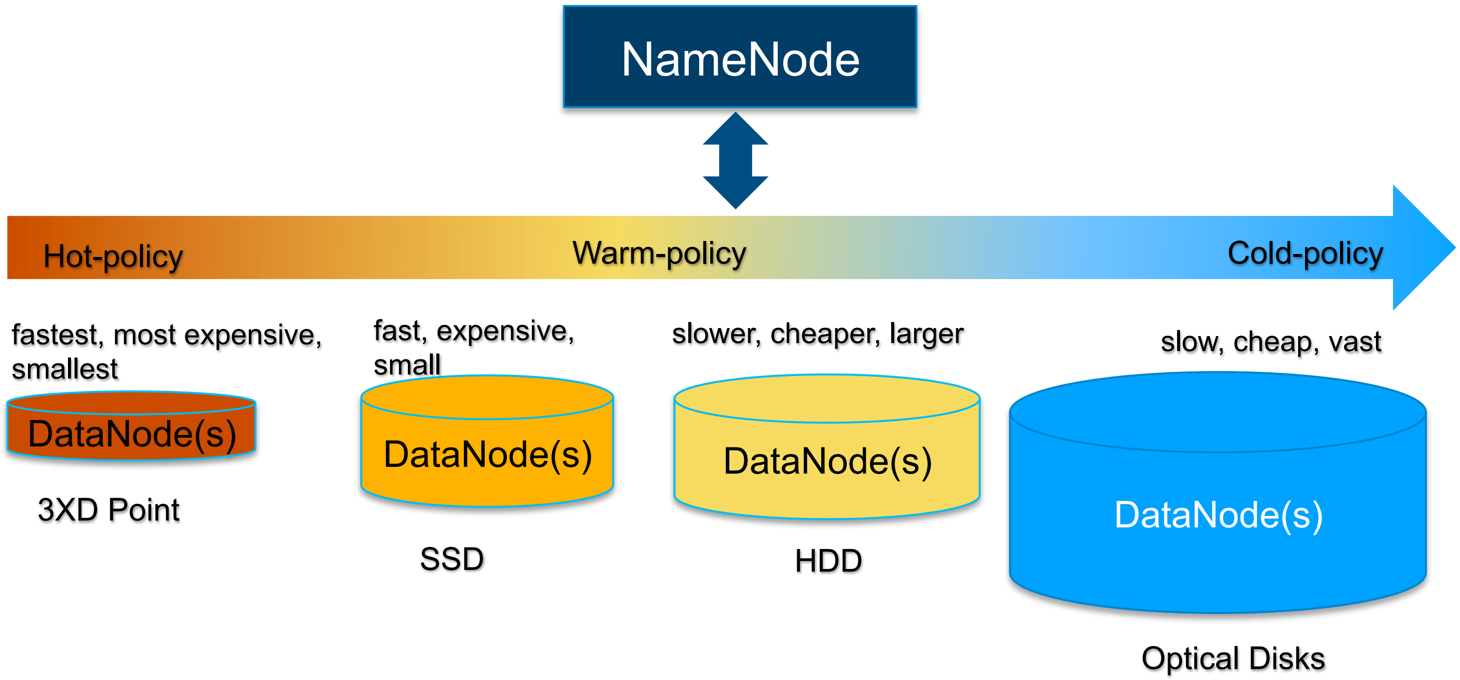

How to migrate when desired data policy changes?

HDFS: Caching -- Pin Hot Blocks in Memory

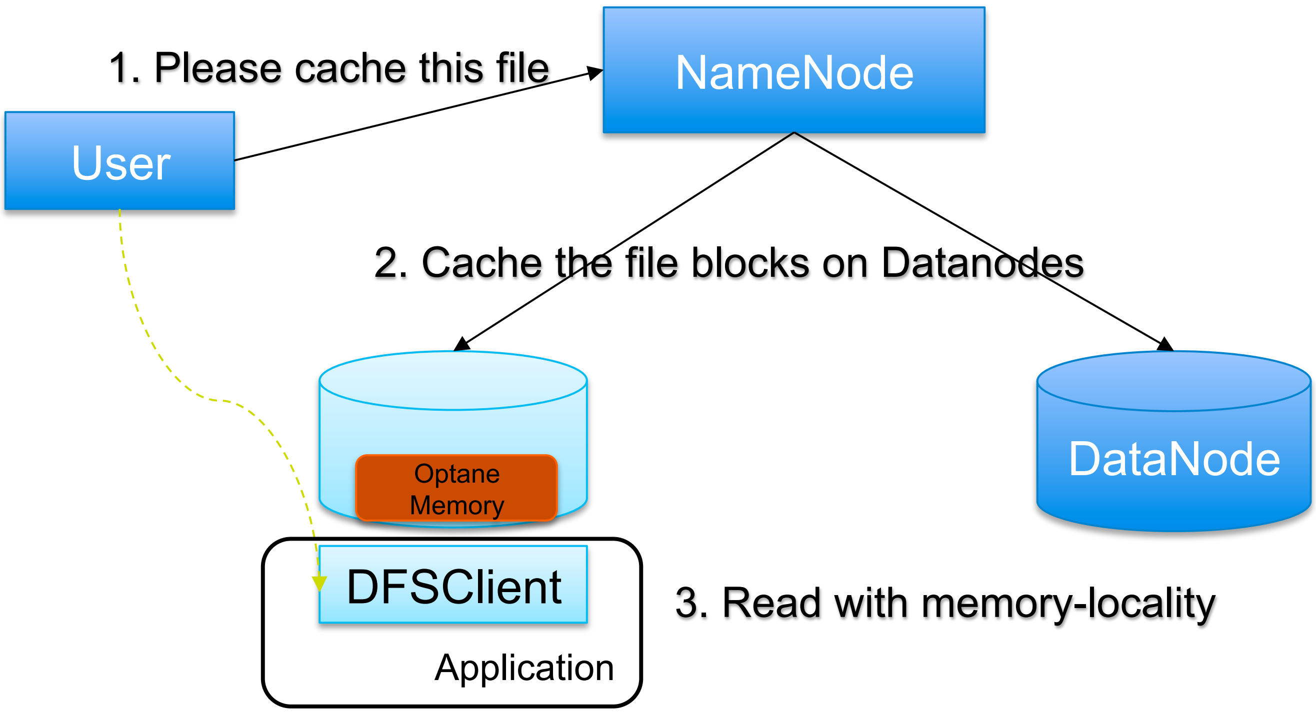

When to cache? Need to do it manually. Don’t forget to clear it after done

Outline

- Background
- HDFS Goodness & Problems
- HDFS Smart Storage Manager
- Wrap up

Smart Storage Manager (SSM)

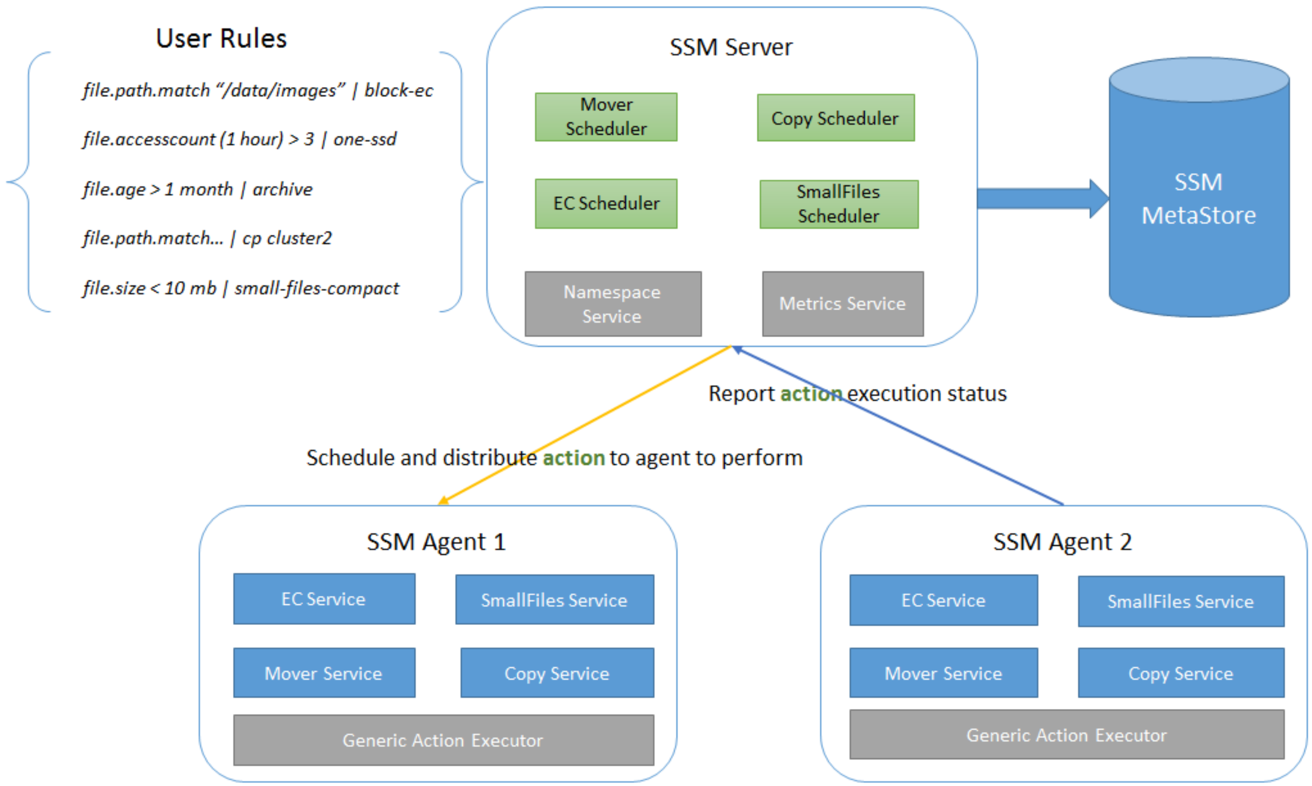

User Defined Rules
based on Usage Patterns & Data Type

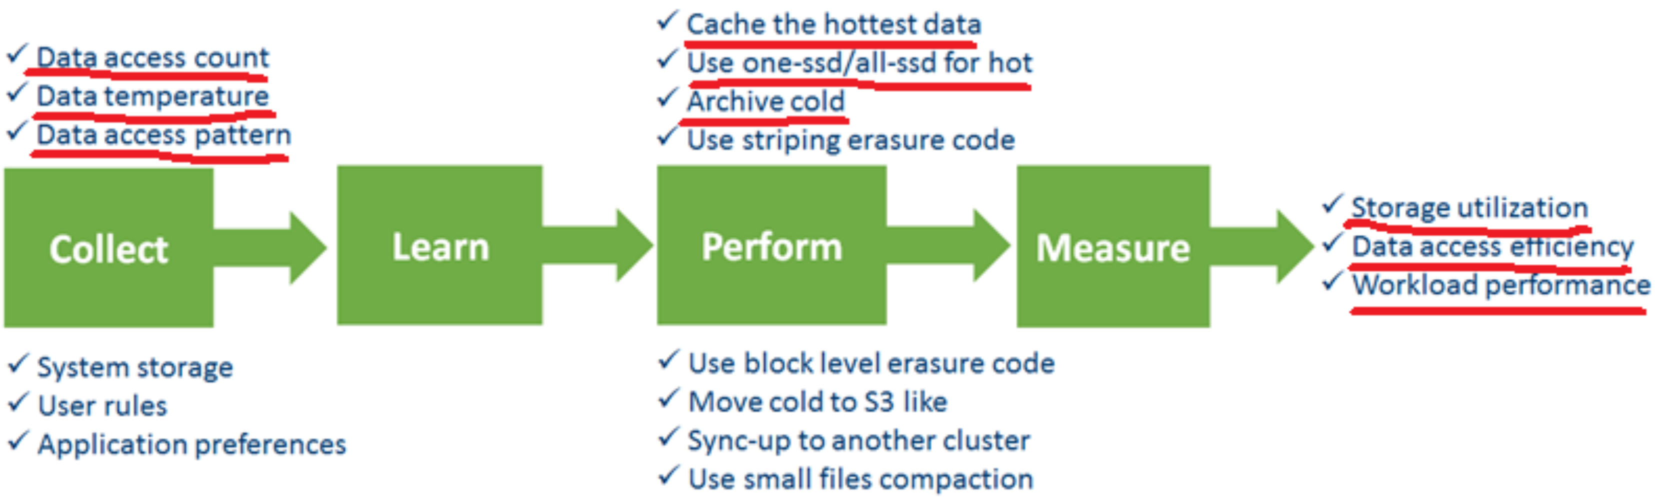

Architectural Highlights

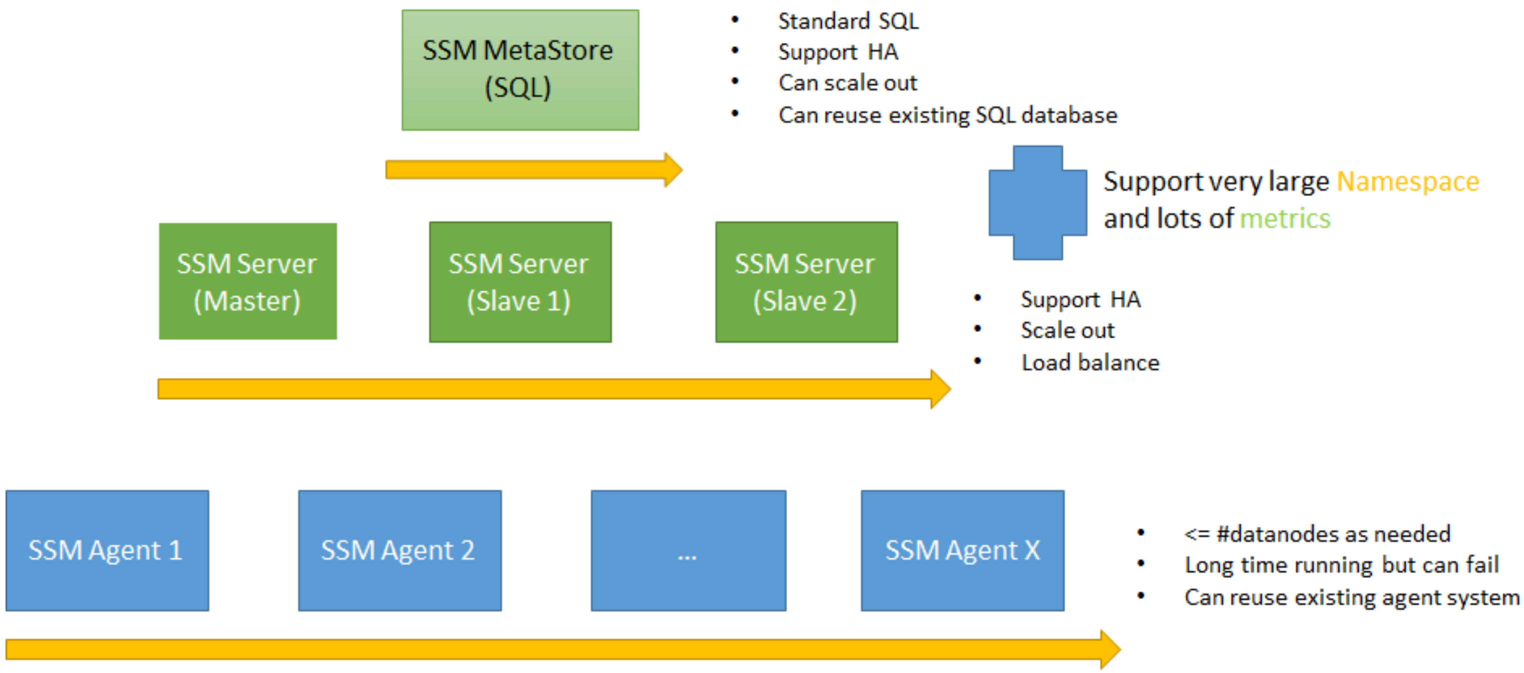

Case 1: Move **HOT** data to fast storage

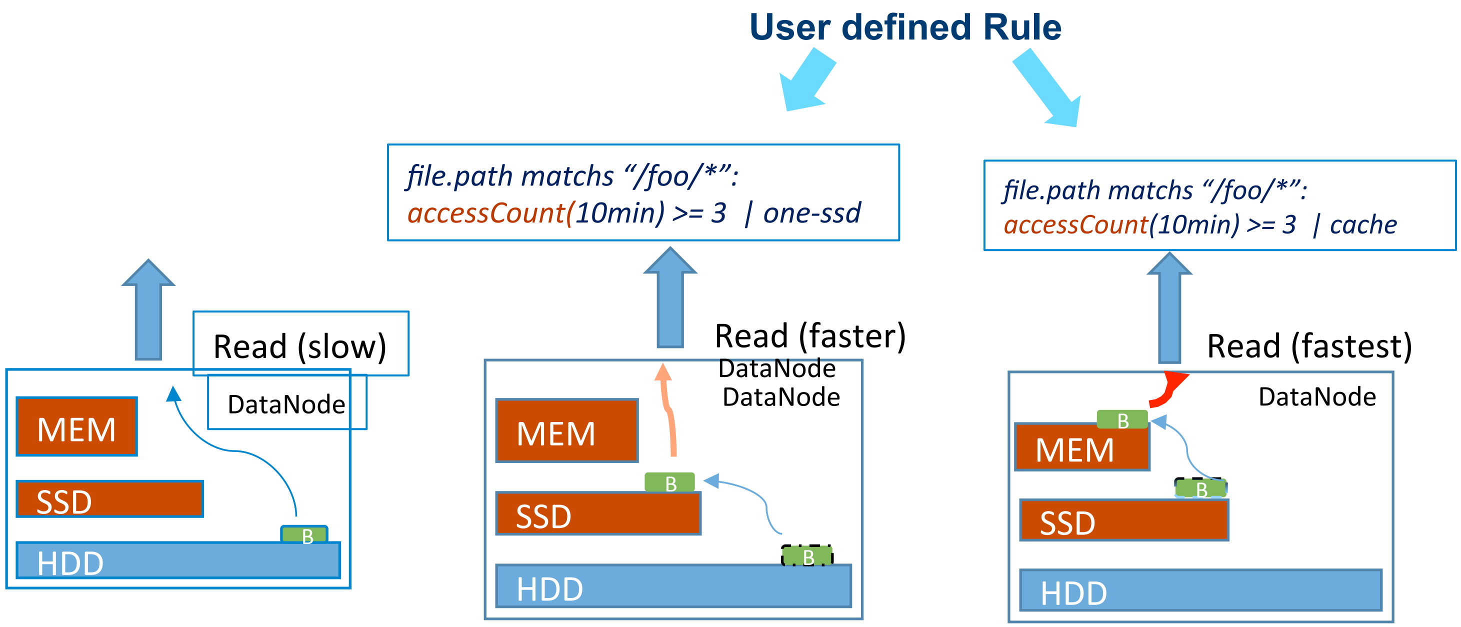

Case 2: Move COLD data to slow storage

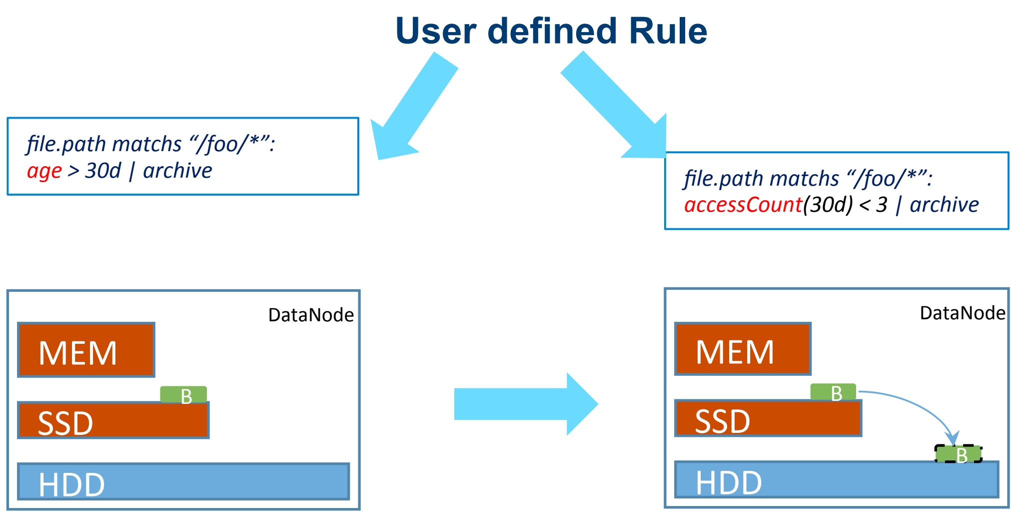

Case 3: ErasureCode COLD data

User defined Rule

file.path matchs "/foo/":*
age > 30d | erasure_code

file.path matchs "/foo/":*
accessCount(30d) < 3 | erasure_code

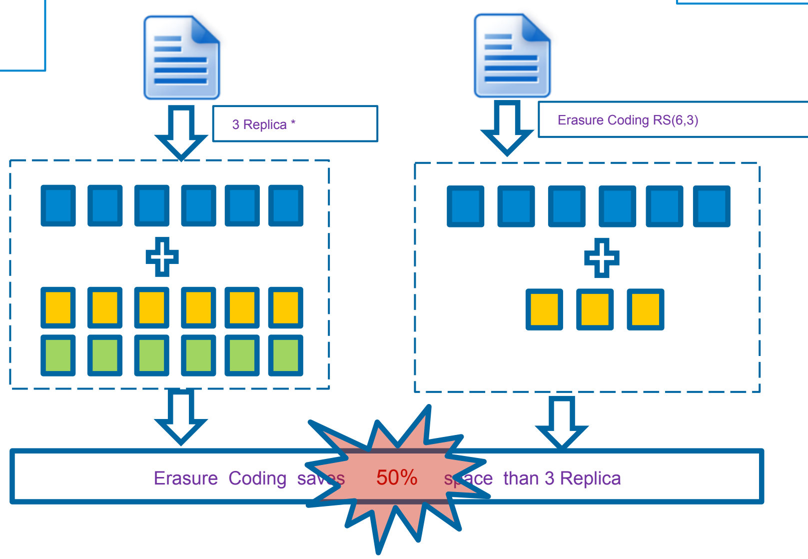

Outline

- Background
- HDFS Goodness & Problems
- HDFS Smart Storage Manager
- Wrap up

SSM Status

Release

– Version 1.0 to be released by Sept 2017

Github

– https://github.com/Intel-bigdata/SSM

Roadmap

- Rule expansion – covering more use cases
- Small File Compaction
- Disaster Recovery via cross-Data Replication

Acknowledgements

Kai Zheng - Architect, Manager

Sammi Chen, Senior Software Engineer

Wei Zho, Senior Software Engineer

Malini Bhandaru – Architect, Principal Engineer

Try It! Get Involved!

Github:

- https://github.com/Intel-bigdata/SSM

Contacts:

- Kai, Zheng (drankye@apache.org | kai.zheng@intel.com)
- Chen, Sammi (sammi.chen@intel.com)
- Zhou, Wei (wei.zhou@intel.com)
- Uma M R Gangumalla (umamahesh@apache.org | uma.gangumalla@intel.com)

Thank you!!!