
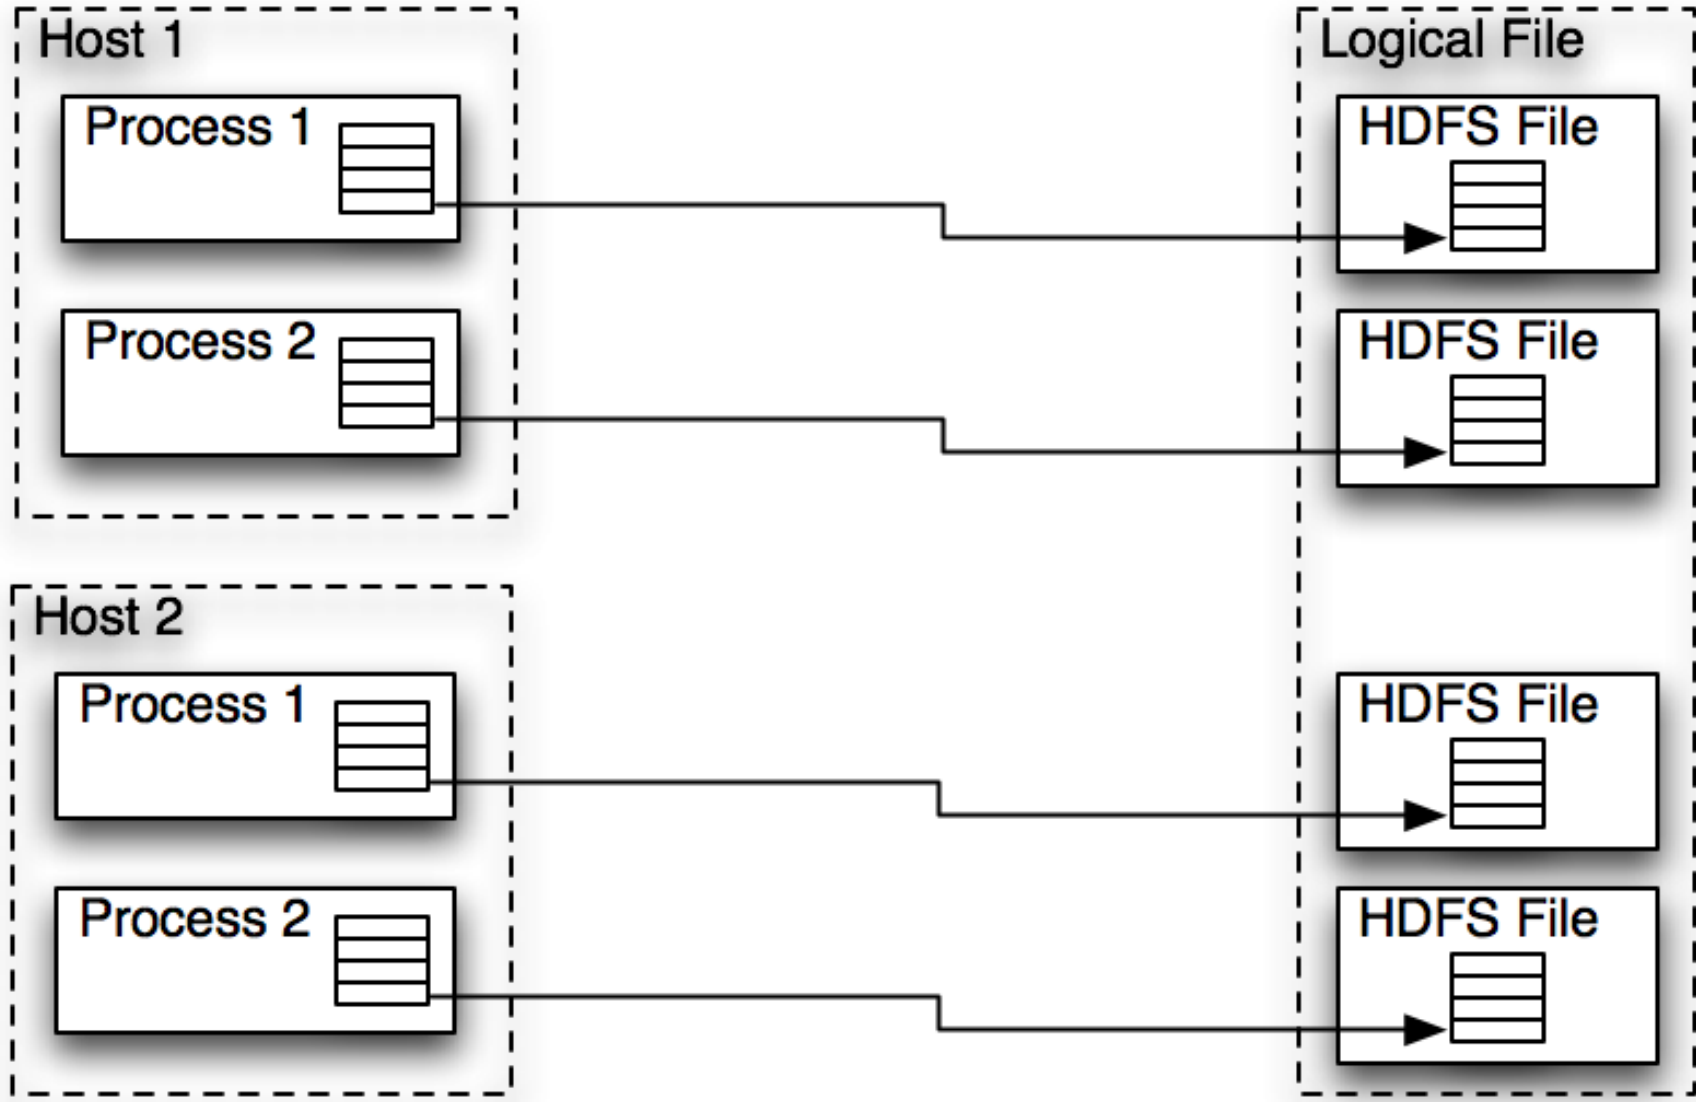


The Collector




Ideal Setup



HDFS Reality



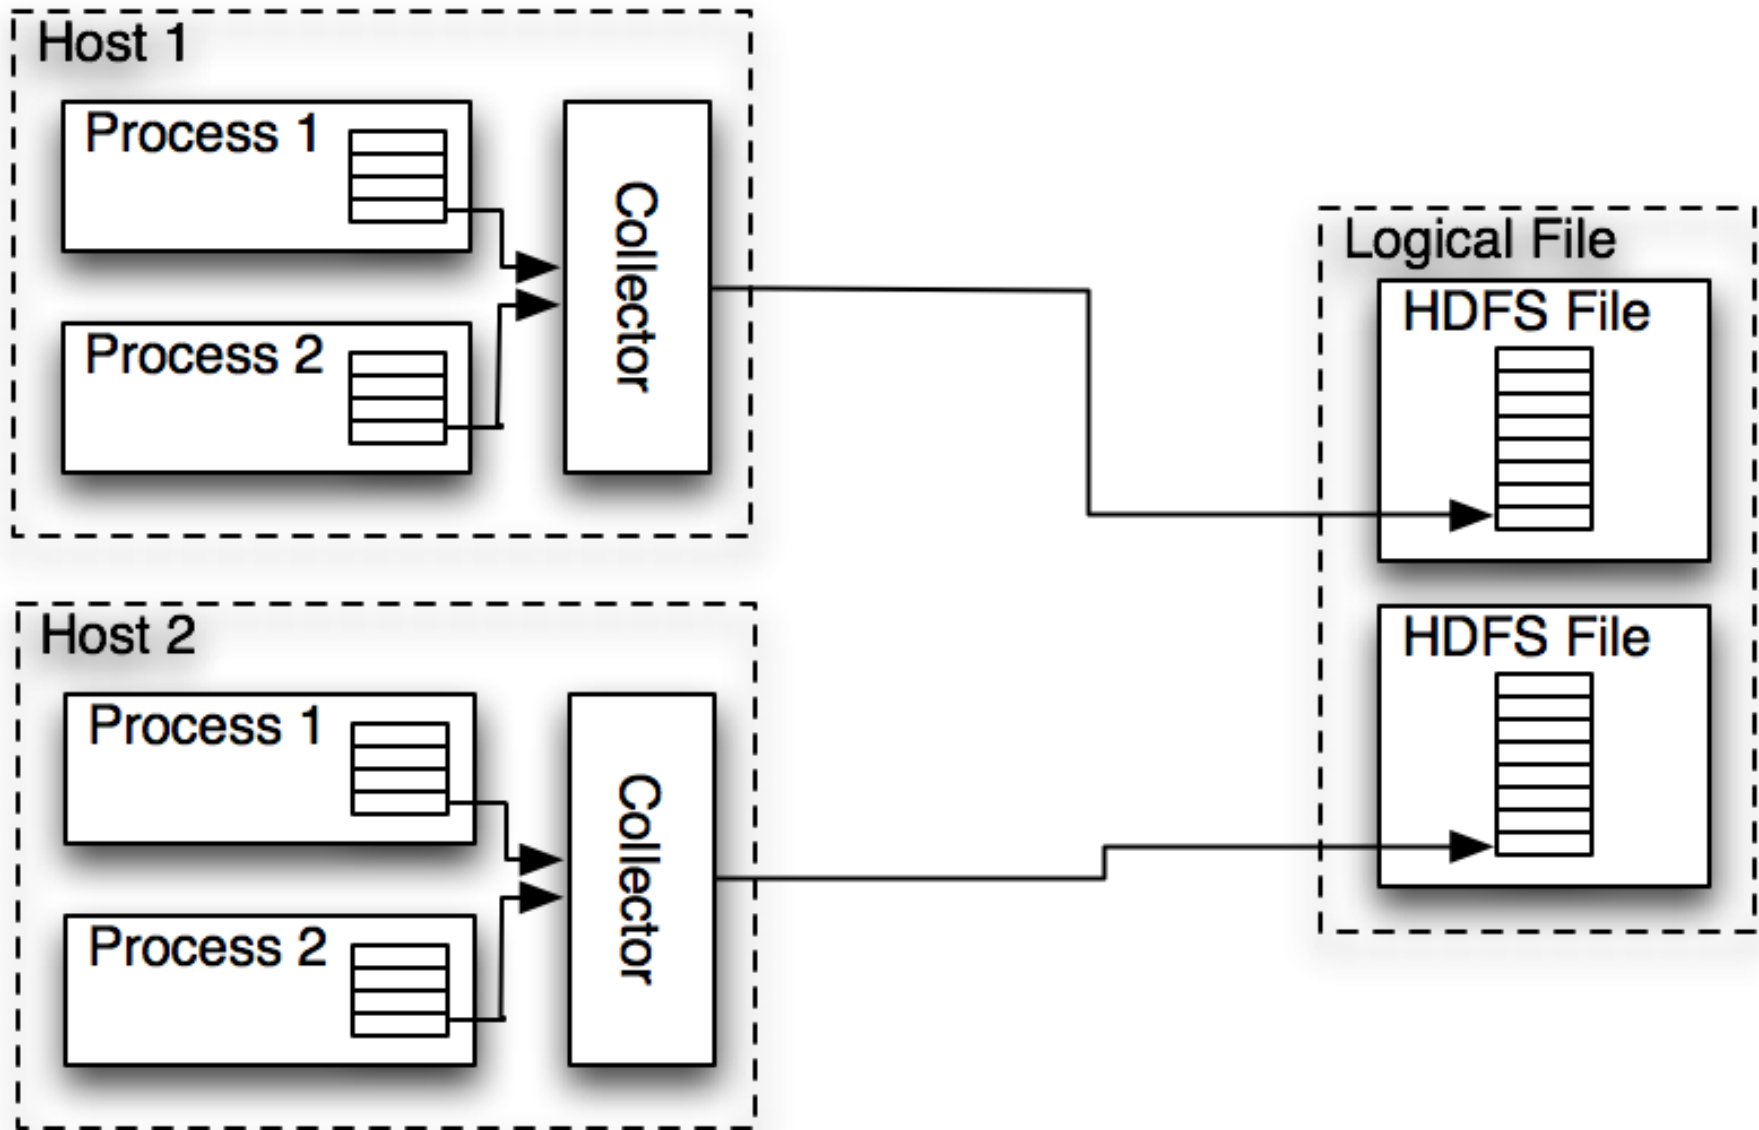
HDFS Reality



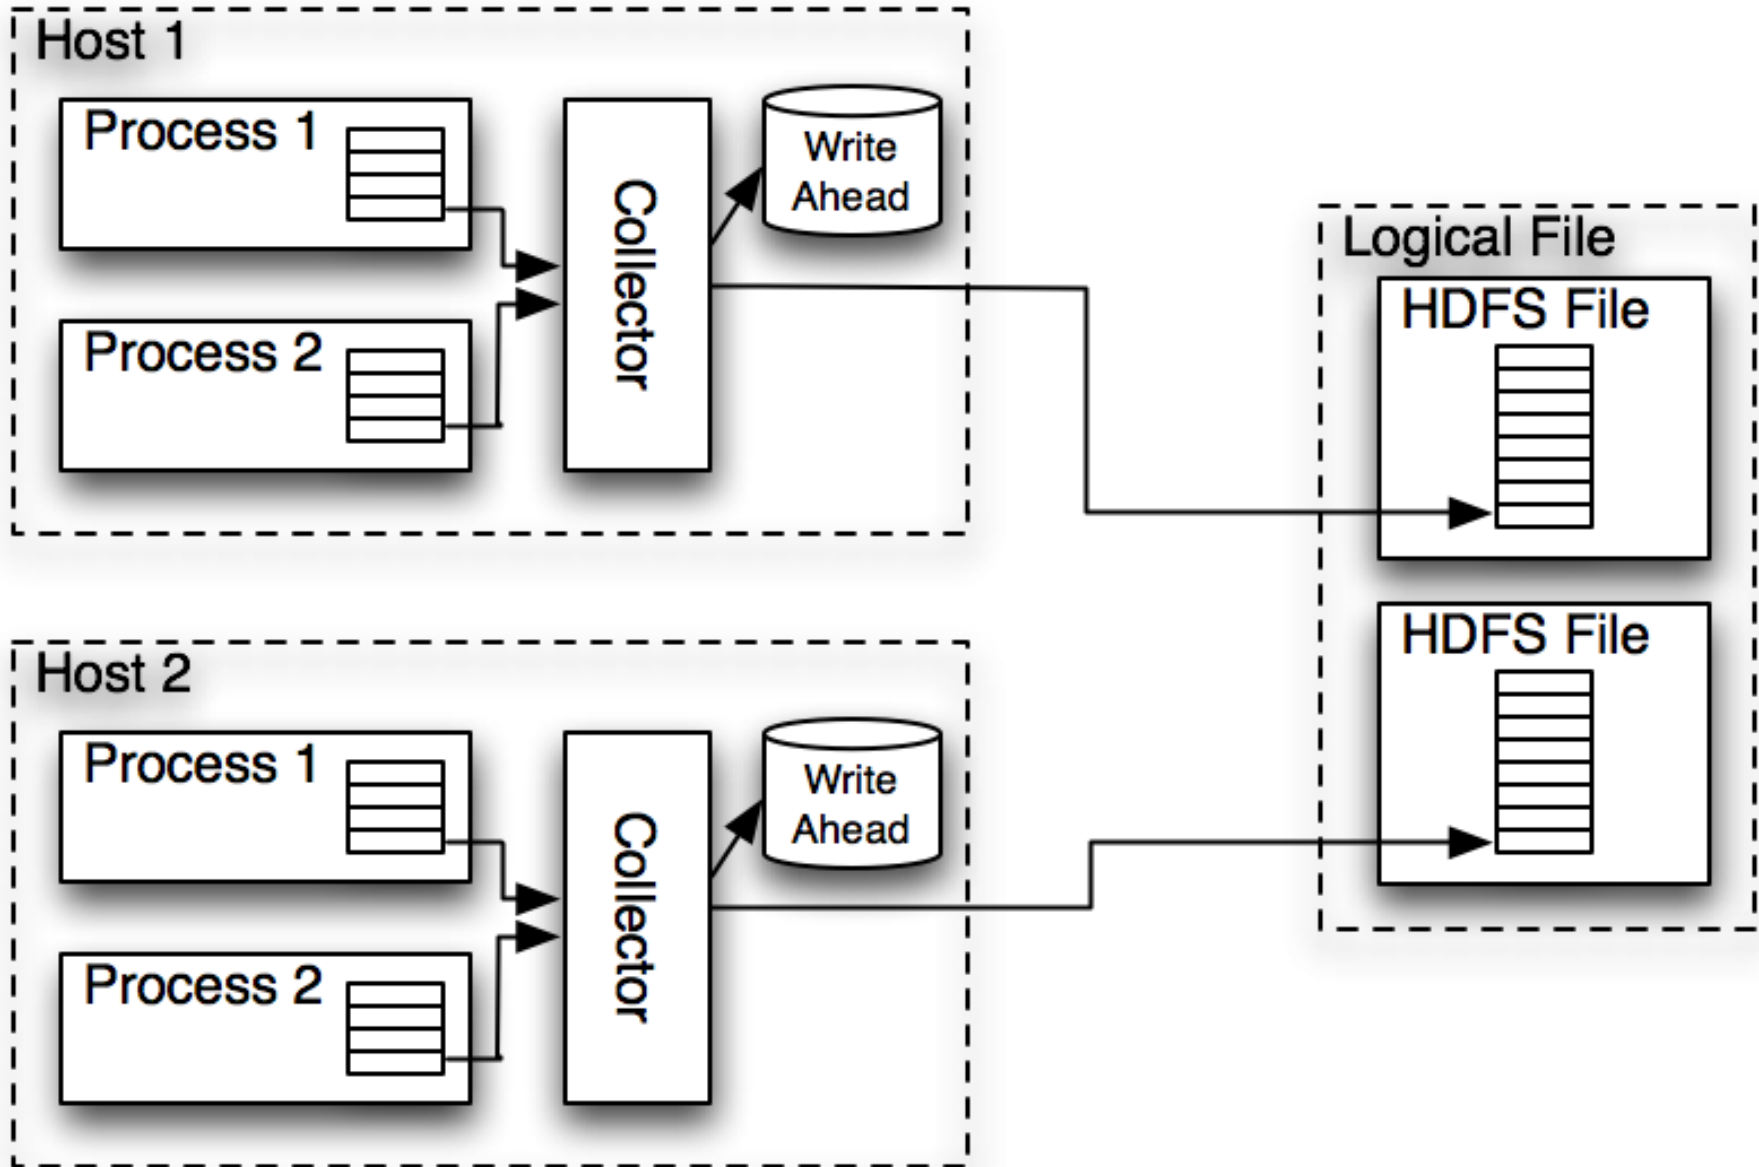
Goals

- Language Neutral
- Fast
- Scalable
- Easy to use and manage
- Durable

The Collector



The Collector



Bonus Round: The Consolidator

