HFetch

Hierarchical data prefetching in multi-tiered storage environments

@ SC'19 Poster Presentation
Hariharan Devarajan, hdevarajan@hawk.iit.edu

SCALABLE COMPUTING SOFTWARE LABORATORY

ILLINOIS INSTITUTE OF TECHNOLOGY

Explosion of data

Increases the cost to access data

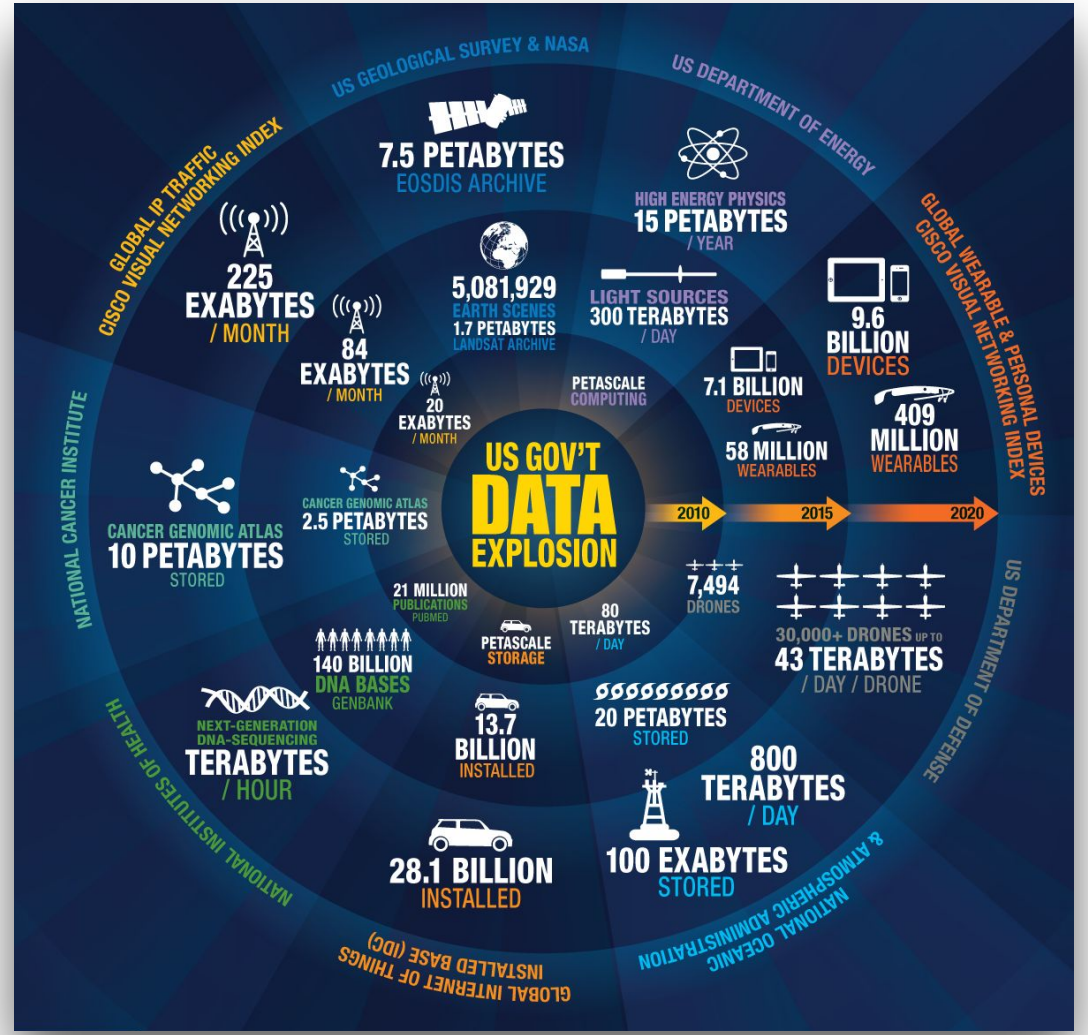

Data Access Optimizations

Can we combine the two approaches?

Goals of HFetch

Server-Push

The server pushes the data to the application in a timely manner instead of applications pulling it.

+

Data Centric

A new prefetching scheme by looking at how data are accessed instead of how applications access it.

+

Hierarchical

The prefetching engine is aware of the multi-tiered hardware and globally places data in it.

HFetch utilizes these three principles to efficiently perform hierarchical prefetching

Design of HFetch

- **Server-Push**
  - Event are captured through kernel's inotify utility
  - Prefetched data is push to the hierarchy
- **Data Centric**
  - Score Incorporates
    - Recency
    - Frequency
    - Sequencing

$$Score_s = \sum_{i=1}^{k} \left(\frac{1}{p}\right)^{\frac{1}{n} * (t - t_i)}$$

- **Hierarchical Placement**
  - The engine calculates placement of prefetch data based on multi-tiered storage and data characteristics.

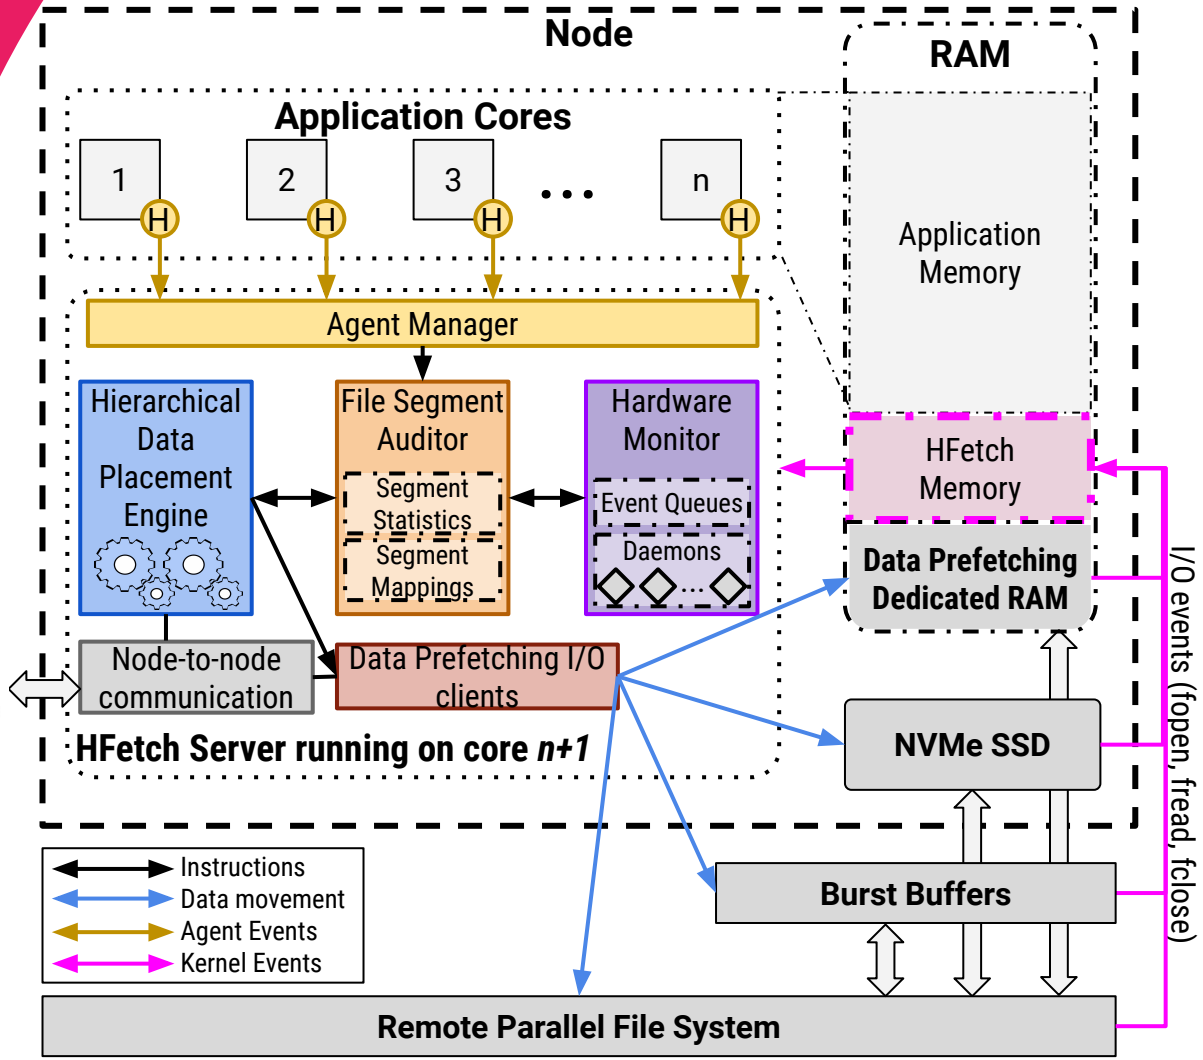

HFetch with Montage and WRF

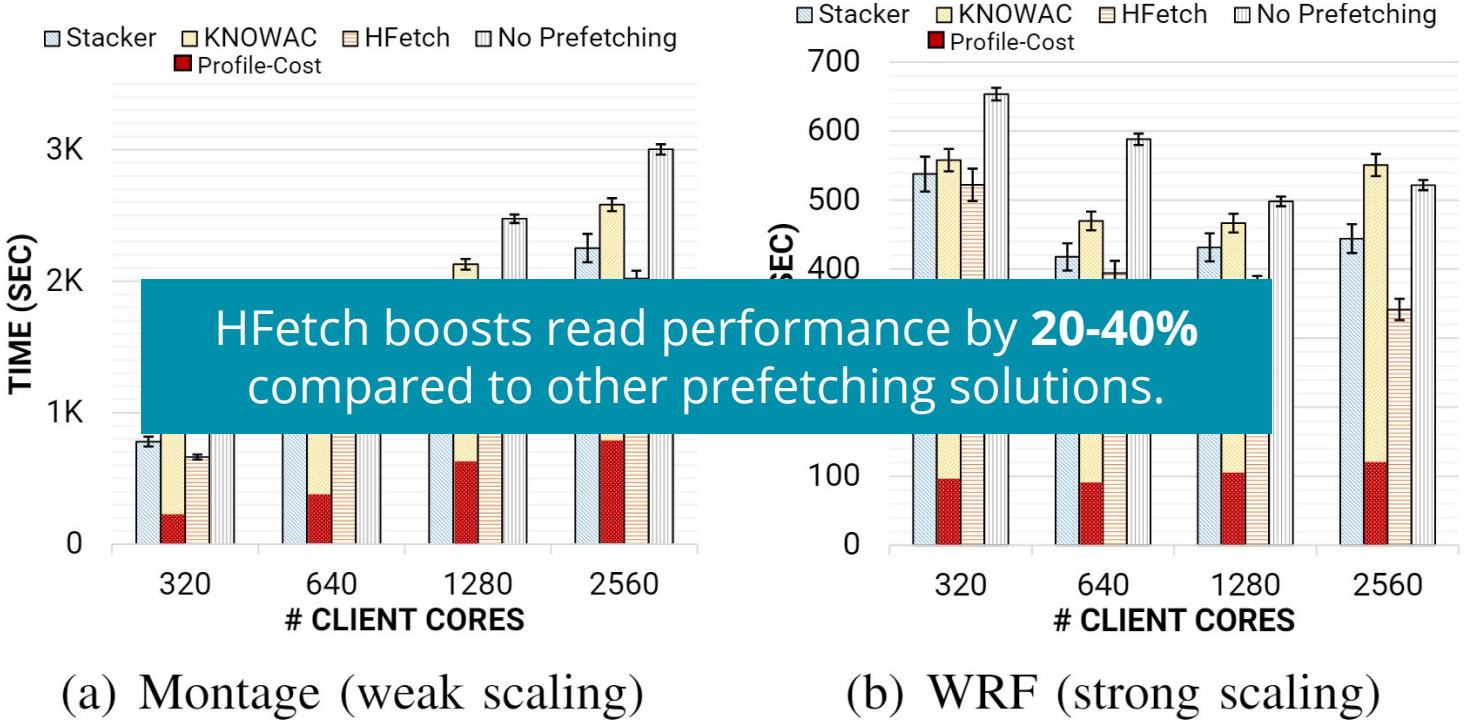

HFetch boosts read performance by **20-40%** compared to other prefetching solutions.

(a) Montage (weak scaling)

(b) WRF (strong scaling)

Questions

Hariharan Devarajan
hdevarajan@hawk.iit.edu