



**CRISP**  
Center for Research on Intelligent  
Storage and Processing in Memory

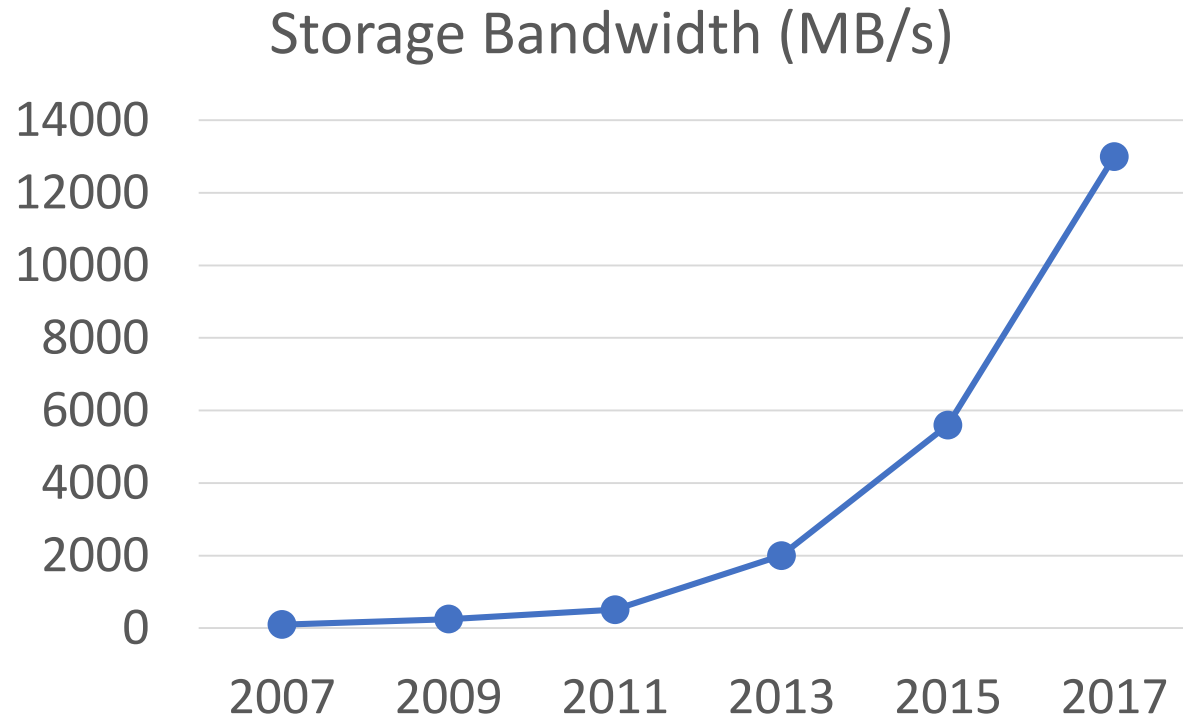
**INSIDER:**  
**Designing In-Storage Computing System  
for Emerging High-Performance Drive**

**Zain (Zhenyuan) Ruan, Tong He, Jason Cong**  
*University of California, Los Angeles*

**UCLA**

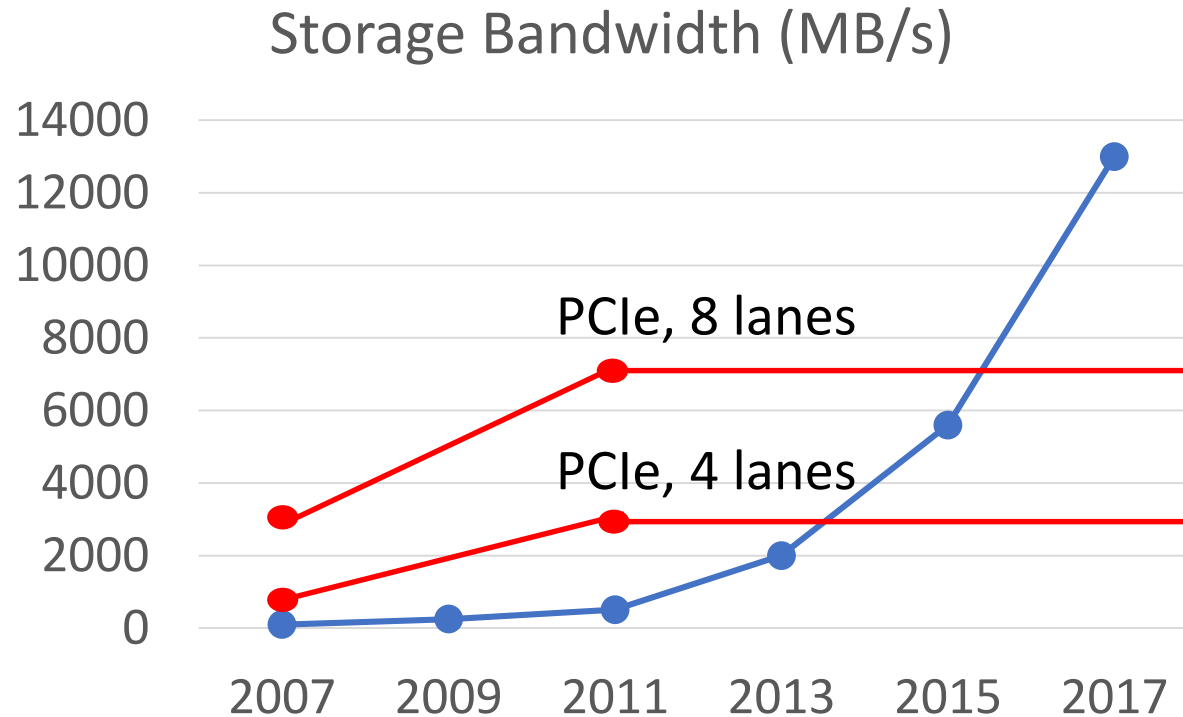
# Data Movement Bottleneck

- “Moore’s Law” of storage drive: bandwidth doubles every two years.




# Data Movement Bottleneck

- “Moore’s Law” of storage drive: bandwidth doubles every two years.
- The interconnection performance does not scale well.

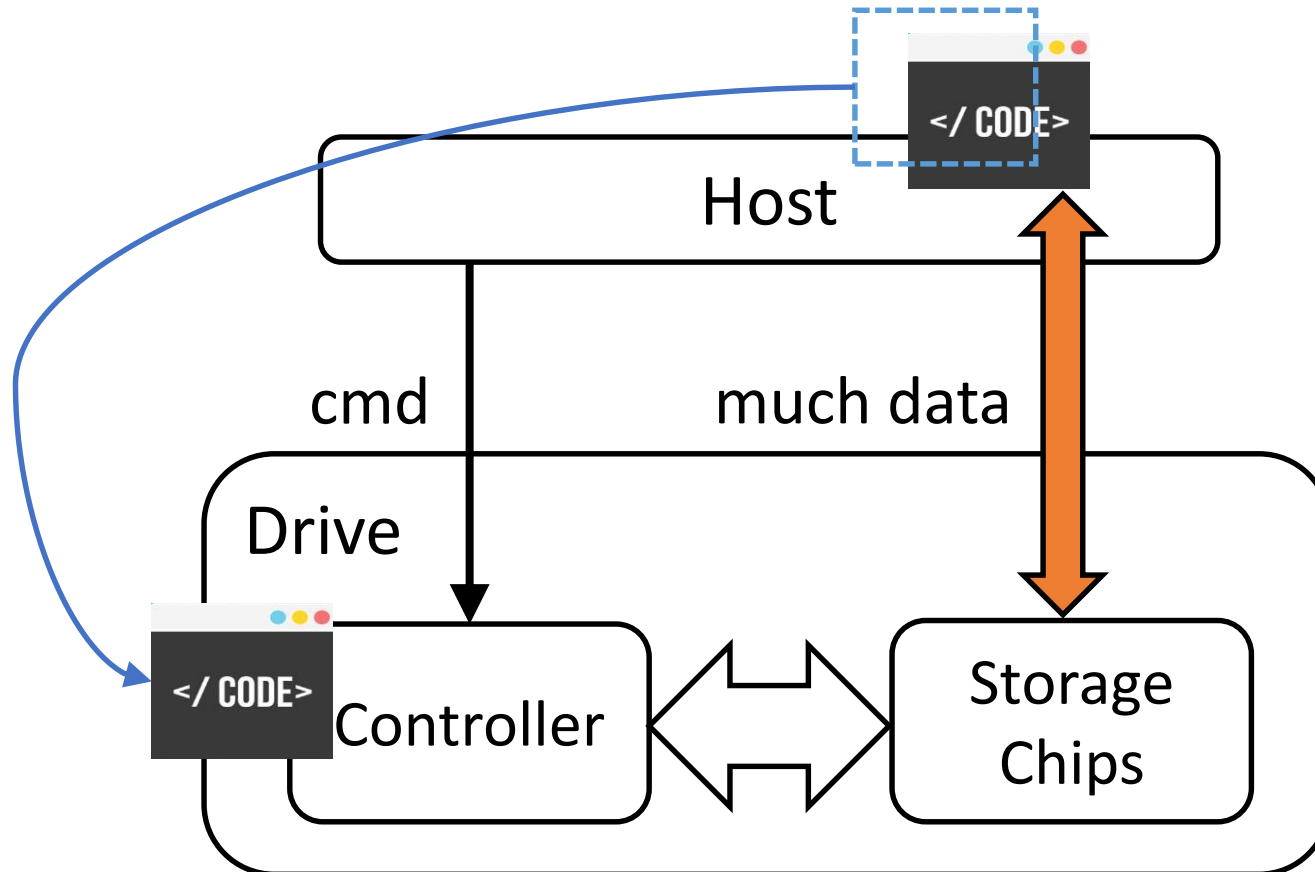


# Existing Work



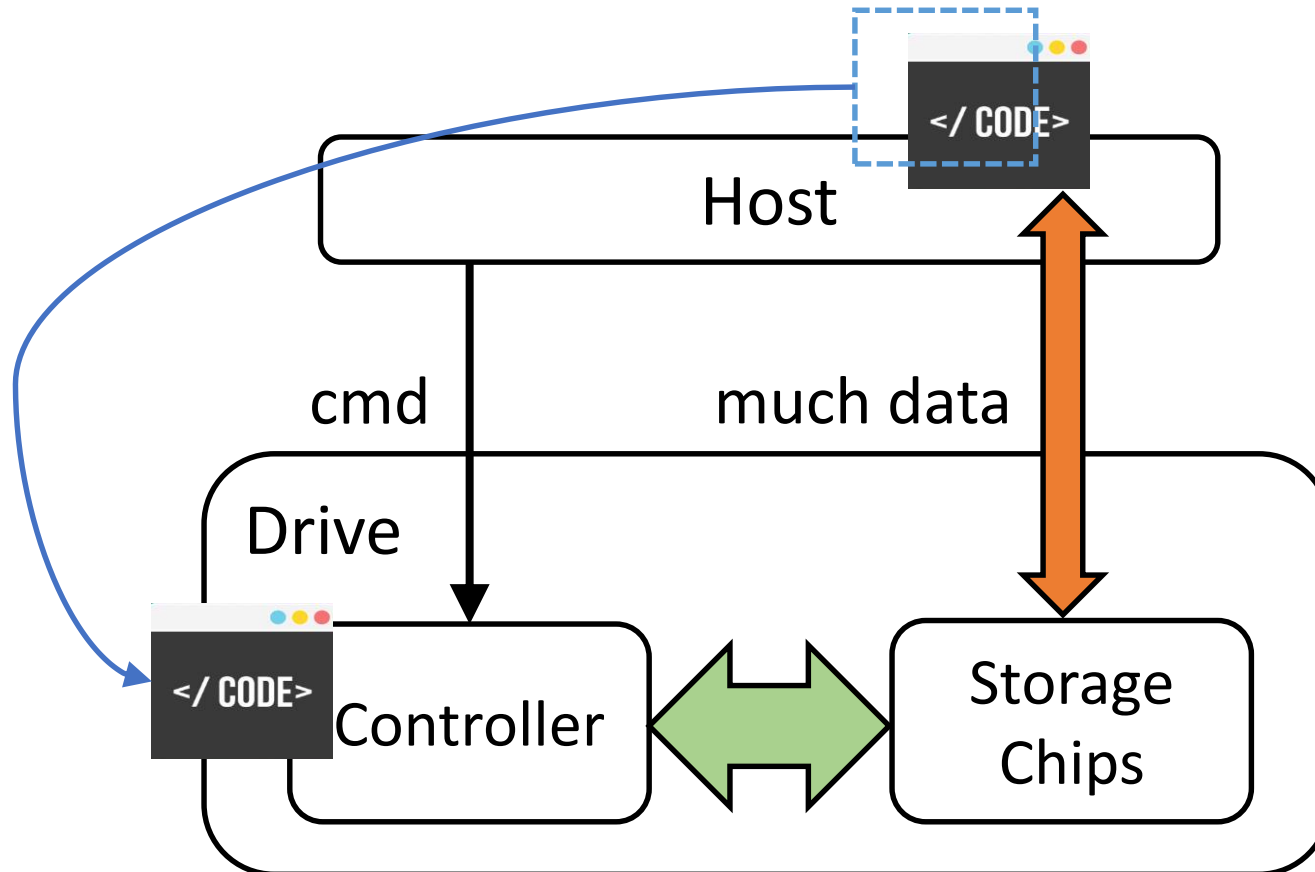
# Existing Work

- In-storage computing.



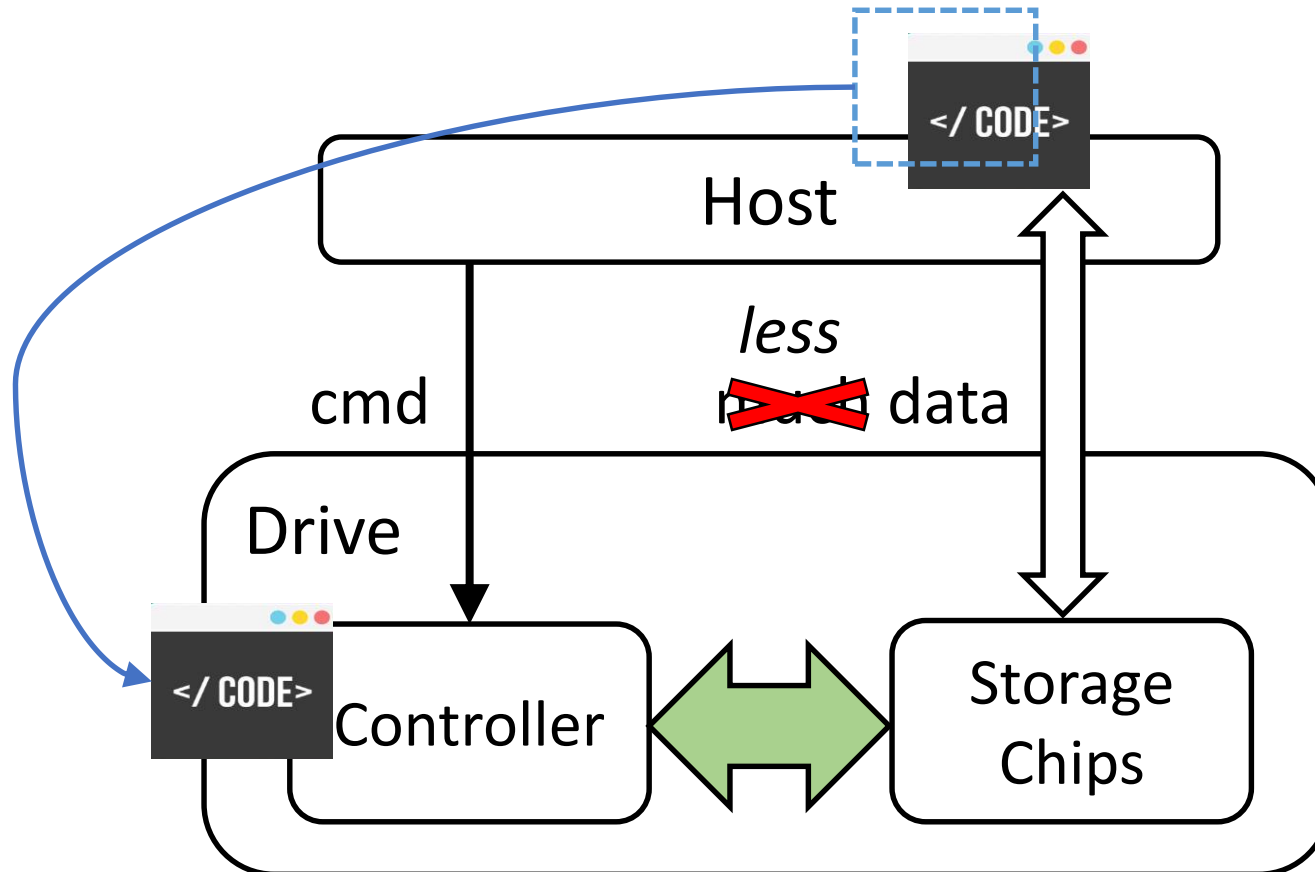
# Existing Work

- In-storage computing.




# Existing Work

- In-storage computing.



# Existing Work

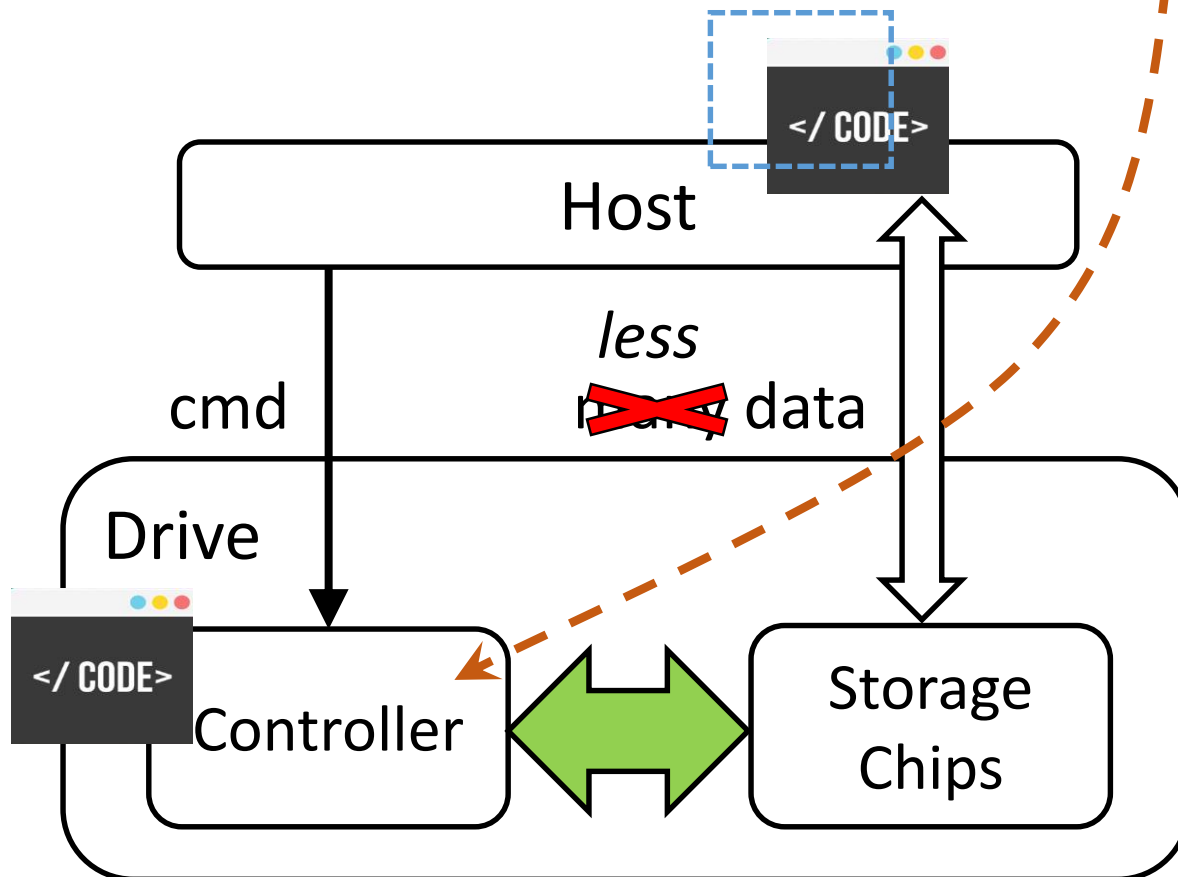
- In-storage computing.





# Existing Work

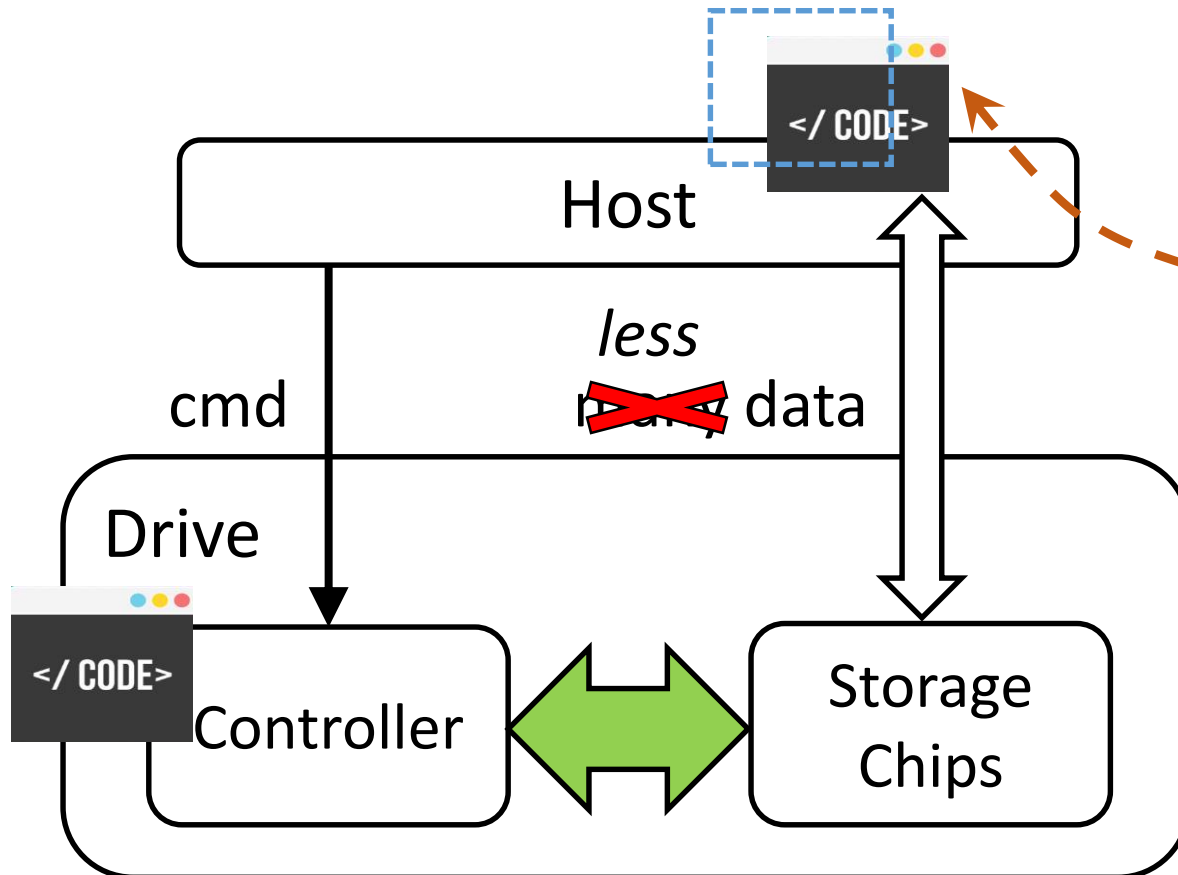
- In-storage computing.



- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.

# Existing Work


- In-storage computing.



- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.
- High programming efforts.
  - Not compatible with existing APIs.
  - Requires considerable code modifications.

# Existing Work

- In-storage computing.



- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.
- High programming efforts.
  - Not compatible with existing APIs.
  - Requires considerable code modifications.
- Lack of crucial system supports.
  - Drive prog. may access unwarranted data.
  - No scheduling among drive programs.

# INSIDER System

- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.
- High programming efforts.
  - Not compatible with existing APIs.
  - Requires considerable code modifications.
- Lack of crucial system supports.
  - Drive prog. may access unwarranted data.
  - No scheduling among drive programs.

# INSIDER System

FPGA-based.

12X perf., 31X cost efficiency.

- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.
- High programming efforts.
  - Not compatible with existing APIs.
  - Requires considerable code modifications.
- Lack of crucial system supports.
  - Drive prog. may access unwarranted data.
  - No scheduling among drive programs.

# INSIDER System

FPGA-based.

12X perf., 31X cost efficiency.

File-based abstraction for  
in-storage computing

- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.
- High programming efforts.
  - Not compatible with existing APIs.
  - Requires considerable code modifications.
- Lack of crucial system supports.
  - Drive prog. may access unwarranted data.
  - No scheduling among drive programs.

# INSIDER System

FPGA-based.

12X perf., 31X cost efficiency.

File-based abstraction for  
in-storage computing

A control plane that enforces perm.  
check and scheduling.

- Limited performance or flexibility.
  - ARM-based --- insufficient comp. speed.
  - ASIC-based --- specific to few workloads.
- High programming efforts.
  - Not compatible with existing APIs.
  - Requires considerable code modifications.
- Lack of crucial system supports.
  - Drive prog. may access unwarranted data.
  - No scheduling among drive programs.

Interested? Welcome to our talk at

*Day 2, Track 1, Session Programmable I/O Device*