Cloud Computing:

The revolution has just begun

Wolfgang Richter, Mahadev Satyanarayanan,
Jan Harkes, Benjamin Gilbert

Who?

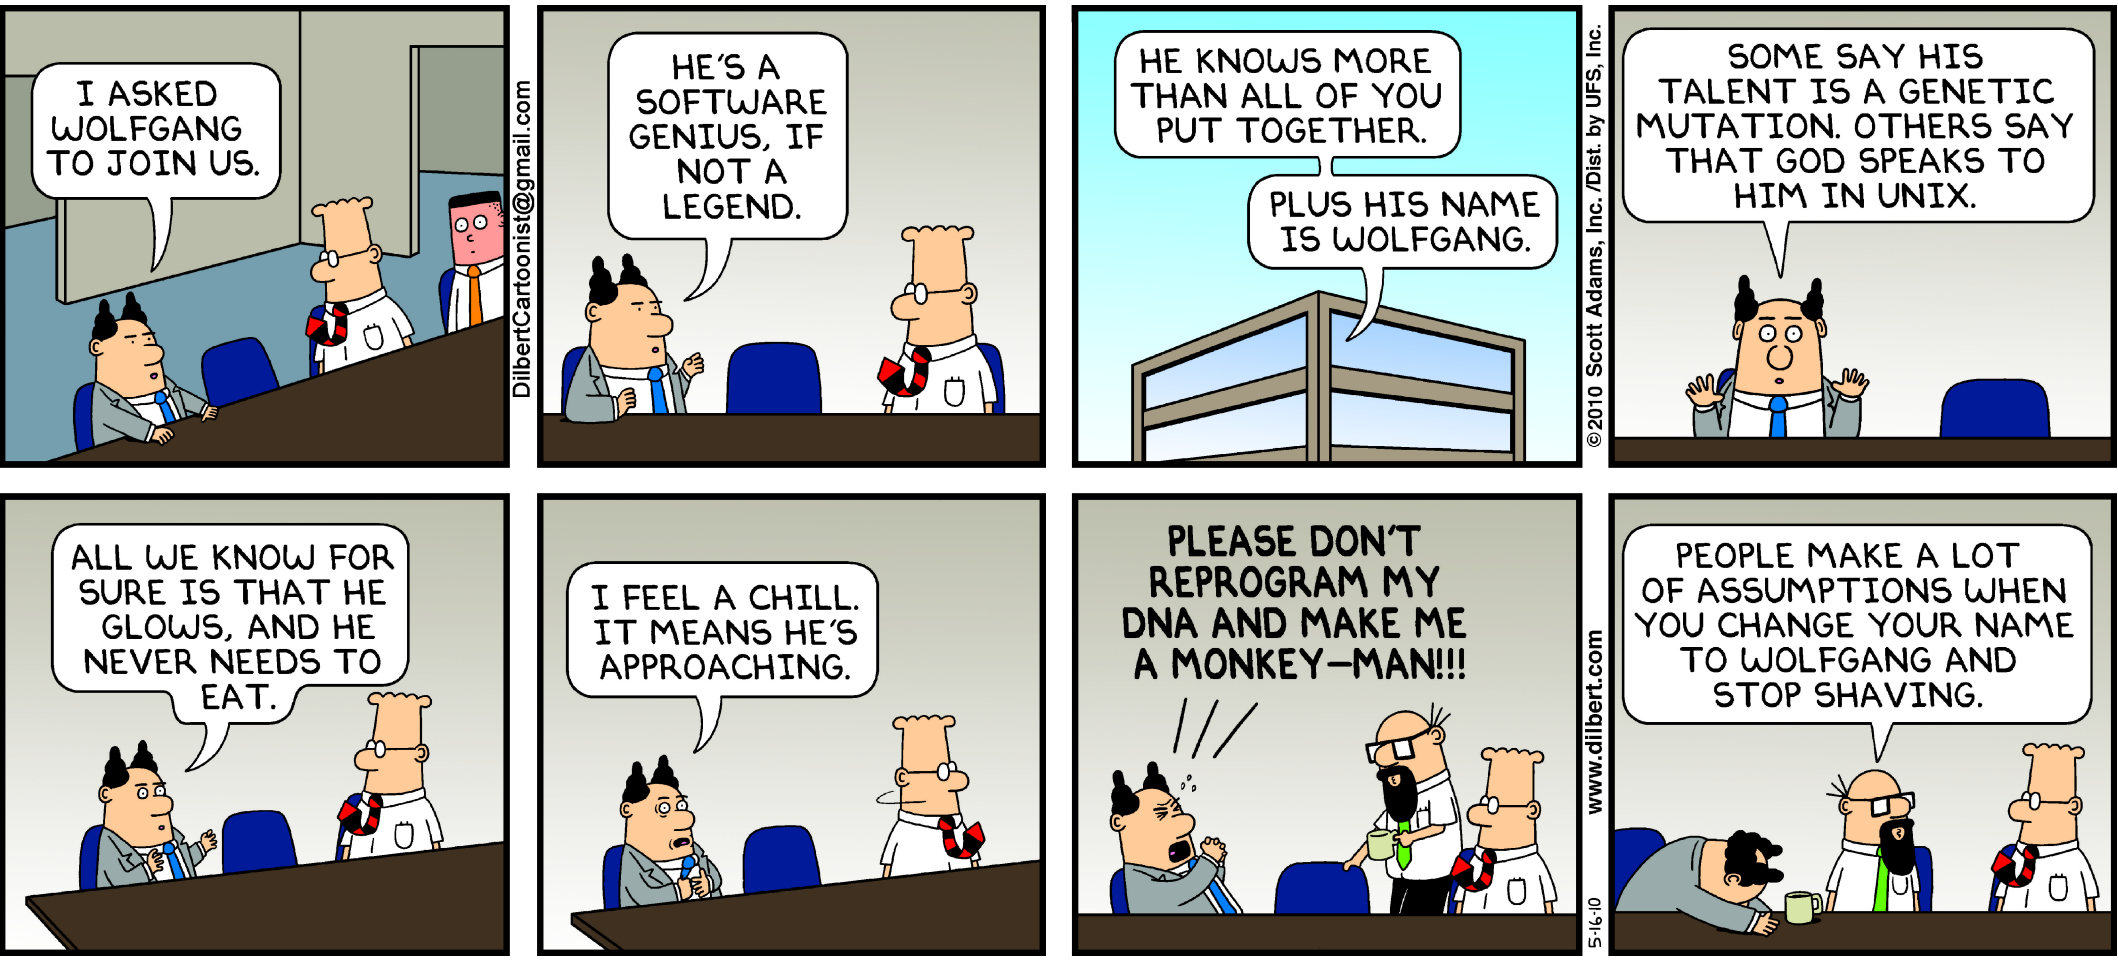
I ASKED WOLFGANG TO JOIN US.
HE'S A SOFTWARE GENIUS, IF NOT A LEGEND.
HE KNOWS MORE THAN ALL OF YOU PUT TOGETHER.
PLUS HIS NAME IS WOLFGANG.
SOME SAY HIS TALENT IS A GENETIC MUTATION. OTHERS SAY THAT GOD SPEAKS TO HIM IN UNIX.
ALL WE KNOW FOR SURE IS THAT HE GLOWS, AND HE NEVER NEEDS TO EAT.
I FEEL A CHILL. IT MEANS HE'S APPROACHING.
PLEASE DON'T REPROGRAM MY DNA AND MAKE ME A MONKEY-MAN!!!
PEOPLE MAKE A LOT OF ASSUMPTIONS WHEN YOU CHANGE YOUR NAME TO WOLFGANG AND STOP SHAVING.
DilbertCartoonist@gmail.com
©2010 Scott Adams, Inc. /Dist. by UFS, Inc.
www.dilbert.com
5-16-10

What?

Introspection

Why?

Because state-of-the-art snapshotting isn't good enough.

(1) High IO overhead
(2) High latency to file mutations

File-level mutations

create delete modify

Goals

Low IO Overhead
Low Latency
Scalability

1: Storage

- Efficient snapshotting techniques
- Focus: low IO overhead implementation
- Fast branching storage systems

1: Storage

- Efficient snapshotting techniques
- Focus: low IO overhead implementation

Simplifying Assumption:

Understanding doesn't matter

2: Smart Disk

- Semantic-based prefetching
- Semantic-based layout/reordering
- “File-aware block level storage”
- Efficient inference

2: Smart Disk

- Semantic-based prefetching
- Semantic-based layout/reordering

Simplifying Assumption:

Omniscience doesn't matter

3: Virtual Machine Introspection

- Infer file-level mutations
- Use only disk block writes
- Good: Low latency event notification
- Bad: Performance overhead
  - Slow IO
  - Single VMM system considered → not cloud

3: Virtual Machine Introspection

- Infer file-level mutations
- Use only disk block writes

Simplifying Assumption:

Performance doesn't matter

  - Slow IO
  - Single VMM system considered → not cloud

What do we want?

Performance of storage
Semantic understanding of smart disks
Virtual expertise of VMI

Idea: Cloud VMI

- Built into the cloud fabric
- Centralized/external...
  - Log monitoring
  - Software auditing
  - Configuration auditing
  - Virus scanning
- No internal agents or guest modifications

Idea: Cloud VMI

- Built into the cloud fabric
- Centralized/external...

Simplifying Assumption:

Privacy doesn't matter

  - Virus scanning
- No internal agents or guest modifications

(1) Fidelity better than state-of-the-art snapshotting

(2) Applications working with snapshotting can work on a stream of file-level mutations

Cloud VMI

- Low Overhead: asynchronous by design
- Low Latency: near-real-time file-level mutation inference
- No guest modifications
- VM instances → near-real-time data sources

Demo with Questions?

wolf@cs.cmu.edu

Technical Report: CMU-CS-12-103

File-Level Mutation Stream Δ

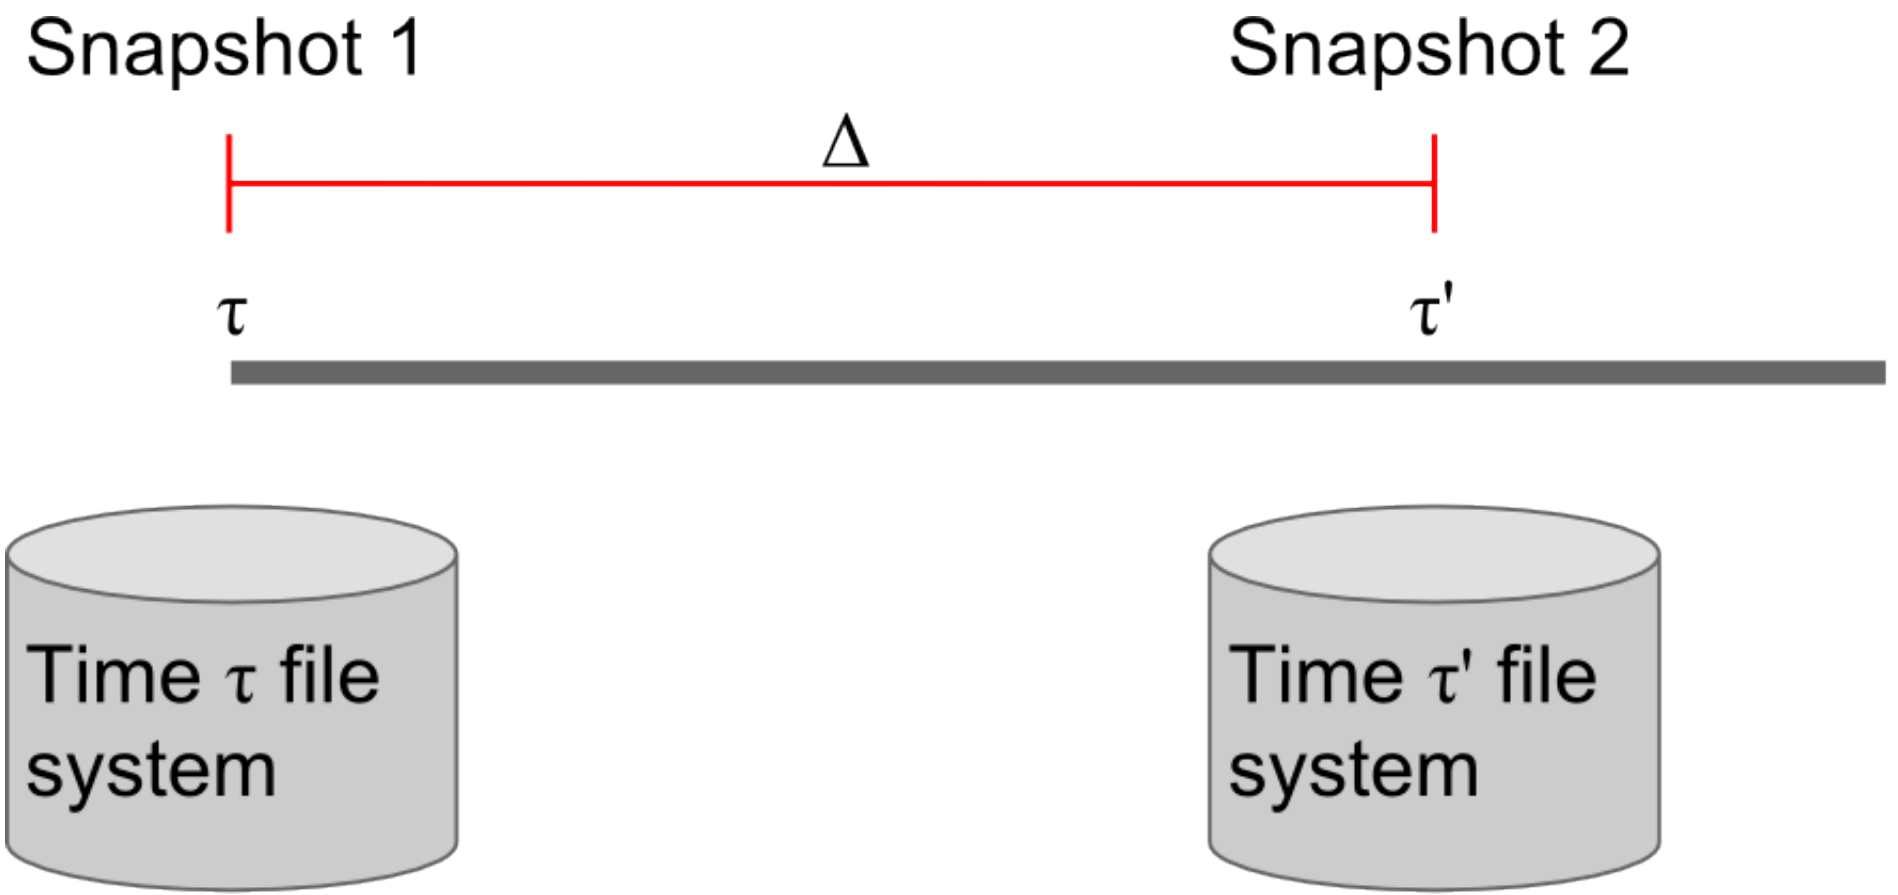

We're going whole stack...

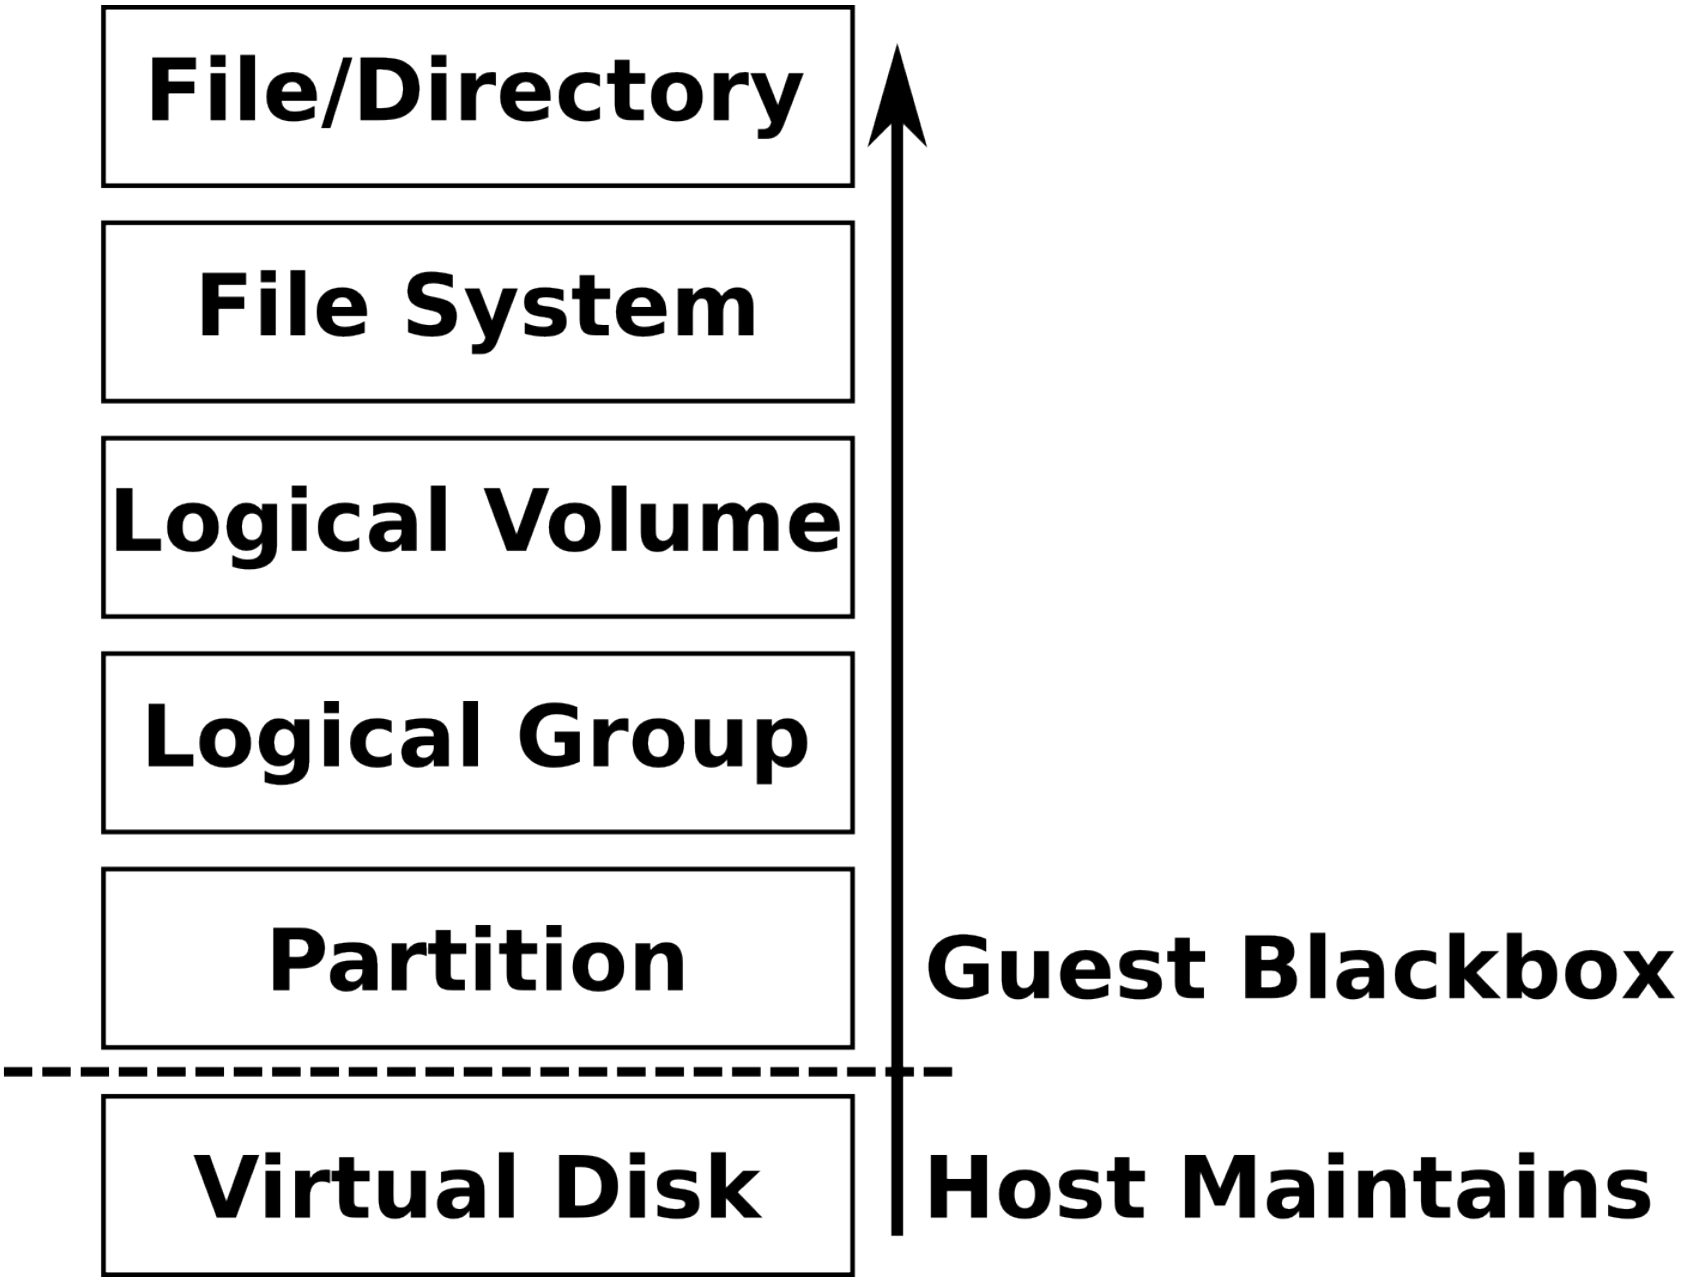

The Inference Engine

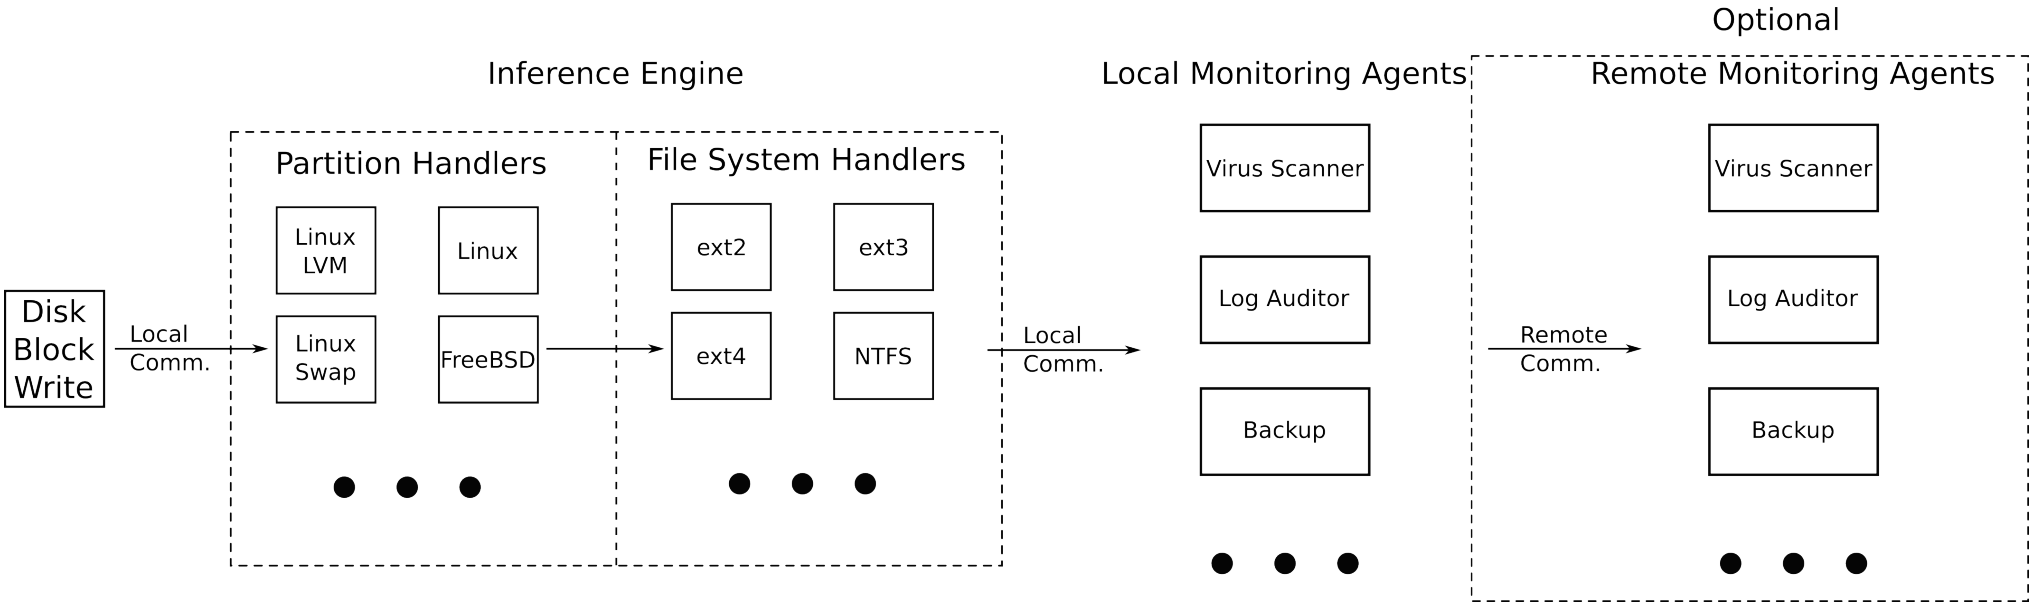

KVM + QEMU

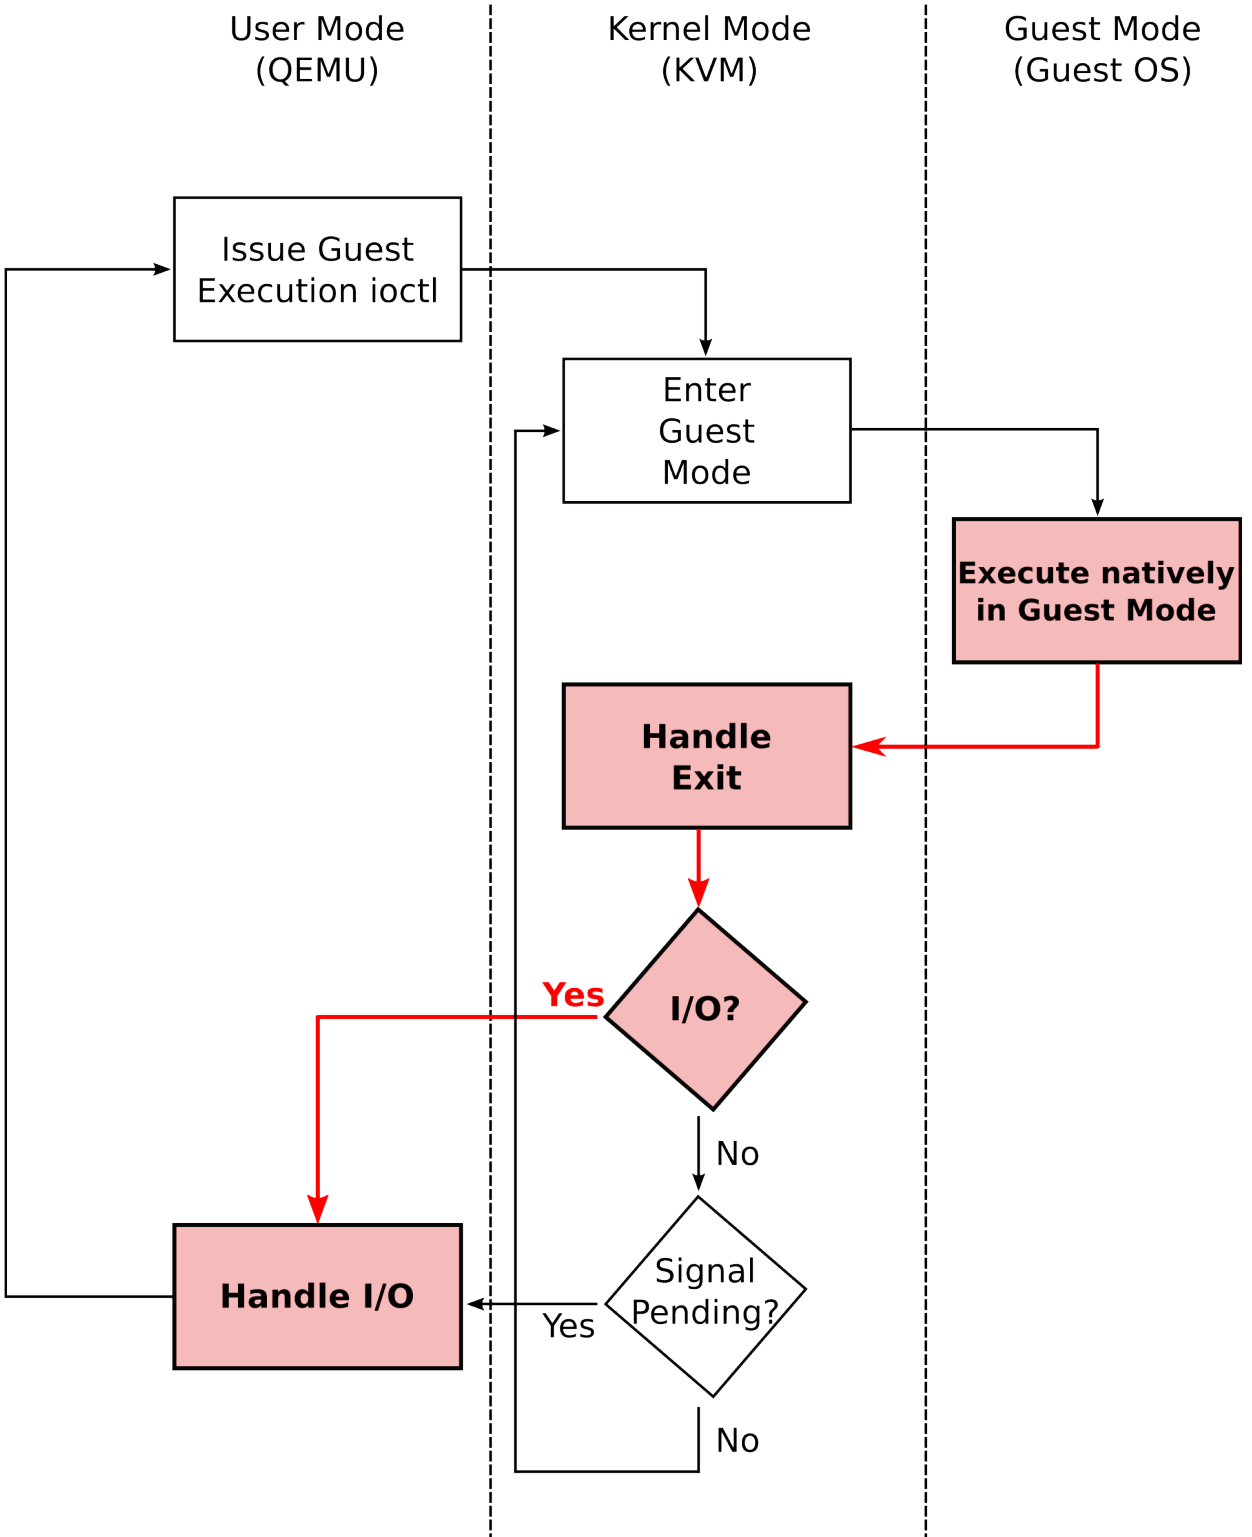
User Mode
(QEMU)
Kernel Mode
(KVM)
Guest Mode
(Guest OS)
Issue Guest
Execution ioctl
Enter
Guest
Mode
Execute natively
in Guest Mode
Handle
Exit
I/O?
Yes
No
Handle I/O
Signal
Pending?
Yes
No