γ-RAY: SCALABLE REAL-TIME EXTRACTION OF FILE UPDATES FROM VMS

Wolfgang Richter^, Canturk Isci*, Jan Harkes^, Benjamin Gilbert^, Vasanth Bala*, Mahadev Satyanarayanan^
^ Carnegie Mellon University, * IBM Research

GOALS

- Scalable: 70,000+ instances
- Real-Time: < 10 minutes
- File Updates: create, delete, modify
- Zero Guest Modifications

IMPLEMENTATION

- Hierarchical Publish-Subscribe API
- Supports live attachment
- Implemented in C
  - ext4, ext3, ext2 drivers
  - NTFS driver
- Clients may be in any language

STATIC ANALYSIS FOR BOOTSTRAPPING

1. EXTRACT METADATA

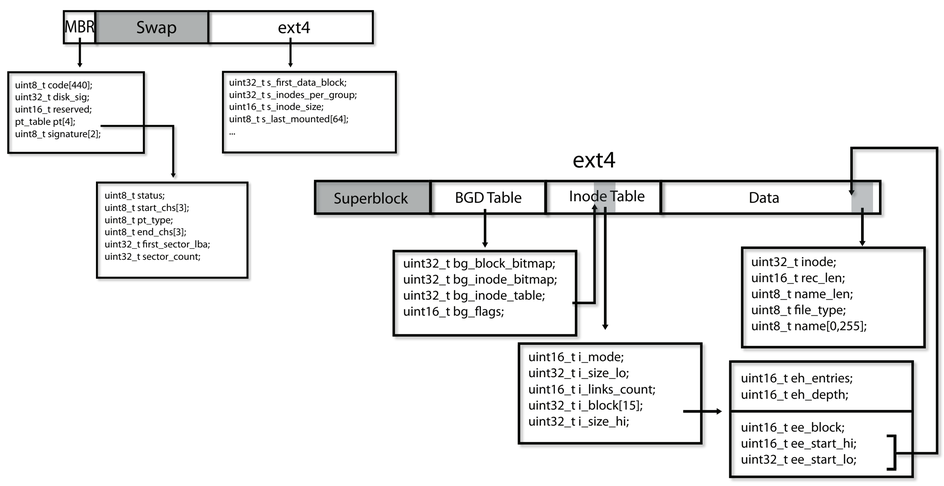

2. STORE

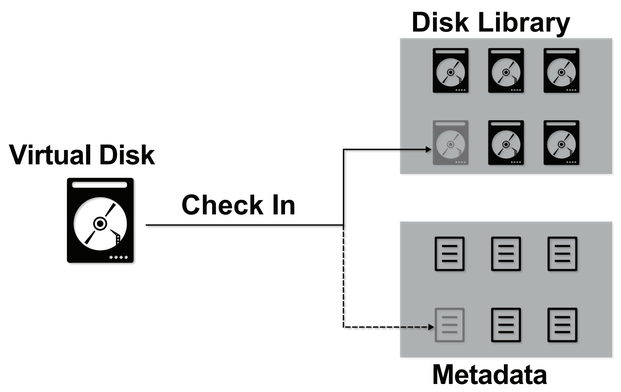

3. RETRIEVE

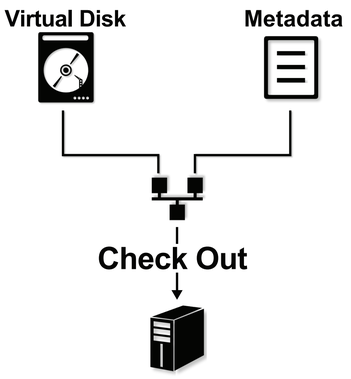

INFERENCE ARCHITECTURE

1. MODIFIED QEMU

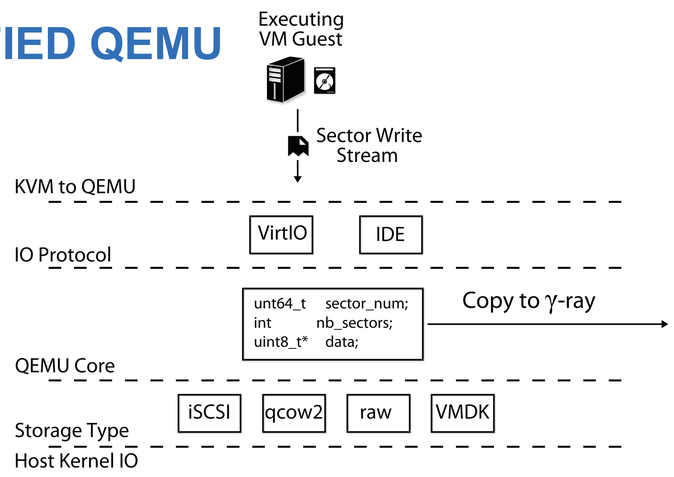

2. ASYNCHRONOUS INFERENCE

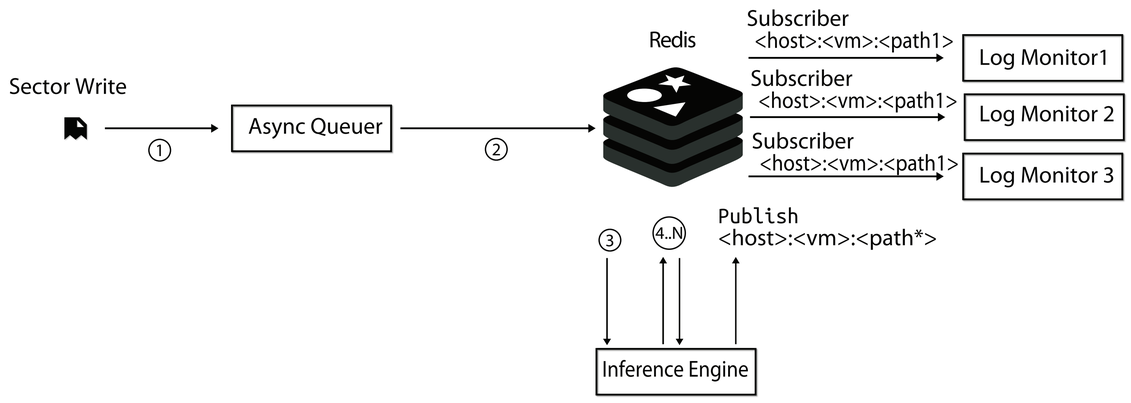

EVALUATION

METADATA

Tunable System Parameters

Tunable	Default
Unknown Write TTL	5 minutes
Async Flush Timeout	5 seconds
Async Queue Size Limit	250 MB
Async Outstanding Write Limit	16,384
Redis Max Memory	2 GB

VM SLOWDOWN

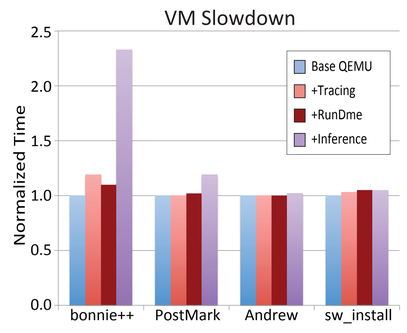

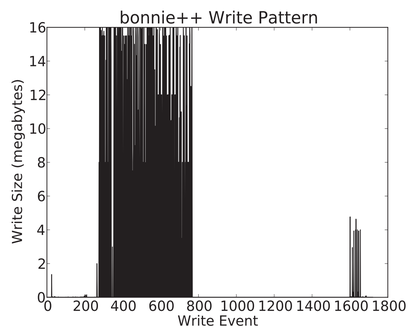

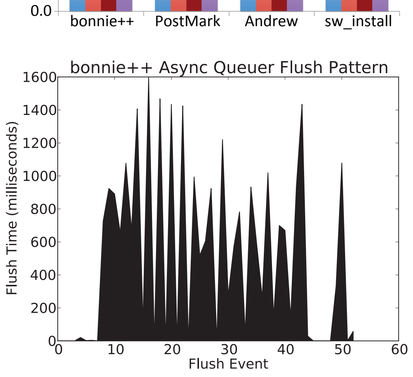

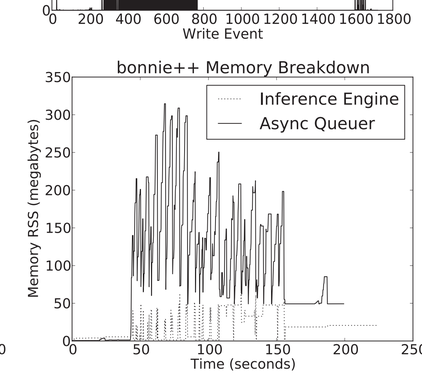

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · PRINCETON UNIVERSITY · UC Berkeley