VIRTUAL APPLIANCE DELIVERY OVER WIDE AREA NETWORKS

Yoshihisa Abe and Mahadev Satyanarayanan (Carnegie Mellon University)

BACKGROUND

- Virtual appliances
  - Pre-configured, ready-to-go virtual machine (VM) images
    - Free users from installing and configuring individual VMs, software etc.
  - Contains entire software and configurations
    - OS, applications, dependencies (libraries)
    - All necessary elements for certain purpose packaged in VM
  - Often used for cloud services
- Virtual appliances for end users have additional value
  - Try new software easily
    - No need to manage it on their own
  - Use software temporarily
    - Rent otherwise expensive software
    - Use software environment only when needed

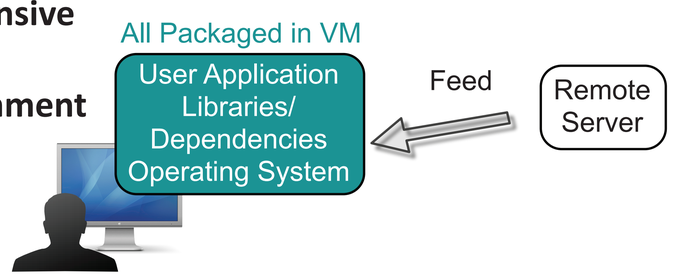

APPROACH

- Take advantage of properties specific to virtual appliances
  - Fixed delivered images
    - Starting point of VM execution is well-known
  - Specific user workloads
    - Each virtual appliance is constructed for use of specific application(s)
- Use VM execution traces to estimate working set for each virtual appliance
  - Prioritize important parts of working set upon transfer so VM can start early
- Deal with missed VM state with prefetching by read-ahead

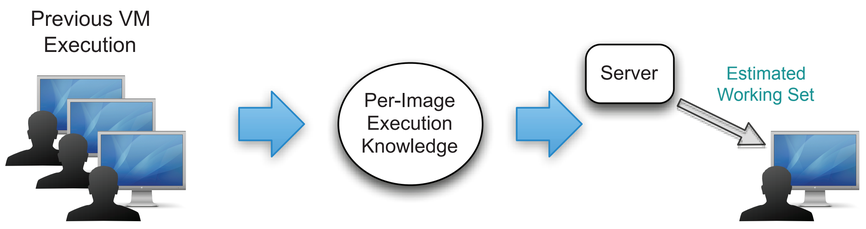

PRELIMINARY EXPERIMENTS

- Memory and disk footprint for varied workloads
  - Represents minimum VM state amount required

	Memory (unit: MB)			Disk (unit: MB)		
	Accessed	Compressed (Separately)	Compressed (Together)	Accessed	Compressed (Separately)	Compressed (Together)
Video Playback	76.4	24.6	22.9	35.8	15.4	13.9
Video Editing	97.1	29.9	27.8	25.3	9.4	8.1
LibreOffice (Writer)	101.7	32.0	29.7	122.4	48.1	42.7

- Video playback performance
  - Initial wait time traded off for good FPS preservation
  - VM state transferred without compression

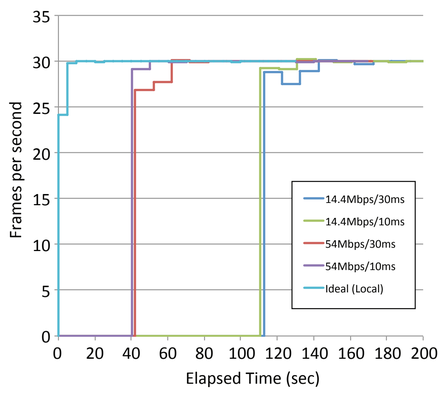

MOTIVATION

- Fast virtual appliance delivery to end users over wide-area networks (WANs)
  - Assumes provision of all VM state necessary for user tasks
    - No VM state cached on client or access through remote storage etc.
  - Work in VM state transfer has mainly targeted good network environments (= LANs)
    - Except Internet Suspend/Resume (CMU), Collective (Stanford) etc.
- Key challenges
  - Deliver virtual appliances quickly
    - Minimize user wait
  - Preserve good VM performance
    - Minimize disruption of VM execution once started
- Efficient VM state transfer to clients
  - VM execution with partial state
    - Whole VM transfer, resulting in tens of GB, is usually prohibitive over WANs
  - Start VM execution as soon as necessary state is available

SYSTEM ARCHITECTURE

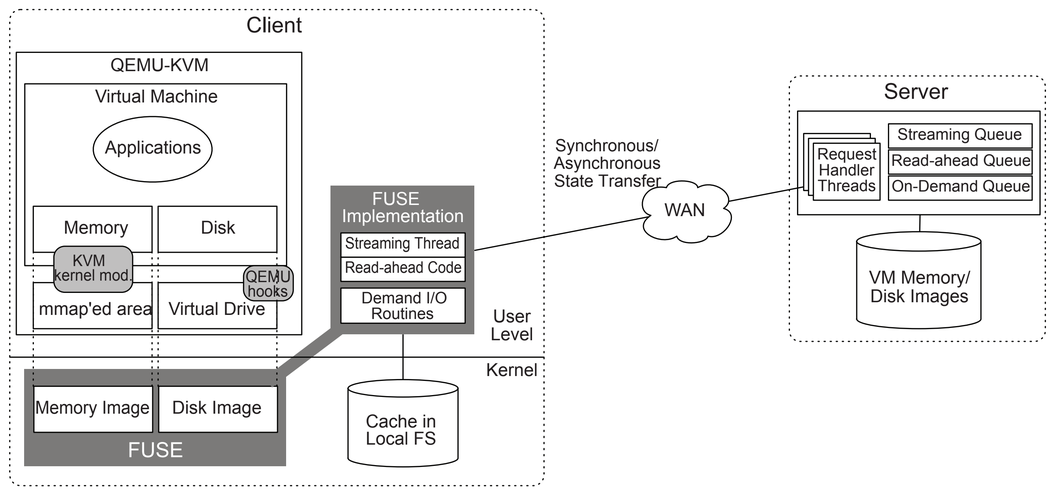

- Client is implemented as modified qemu-kvm
  - Memory and disk images passed through FUSE
- Lightweight TCP server handling state prefetch/fault-in requests

WORK IN PROGRESS

- Refining working set estimation
  - Balancing trade-off between initial wait and VM performance
    - More wait while transferring larger working set leads to better expected performance
  - Using VM execution traces gathered through collection facility
    - Automated VM launch with instrumentation, on remote host
    - Target application is used in various ways while traces are collected
- Performance evaluation with virtual appliance images
  - System-level metrics: VM stall, state hit/miss rates etc.
  - Application-level metrics: FPS, timings of interactive actions etc.

Intel Science & Technology Center for Cloud Computing

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