# Capo: Recapitulating Storage for Virtual Desktops

Mohammad Shamma, Dutch T. Meyer, Jake Wires, Maria Ivanova, Norman C. Hutchinson, and Andrew Warfield

University of British Columbia

# The World According to Gartner

- 60% of all enterprises will deploy network computers by 2001
- 5 – 30 million Windows terminals sold per year by 2005
- 40% of desktops – *49 million* – will be virtualized by 2013

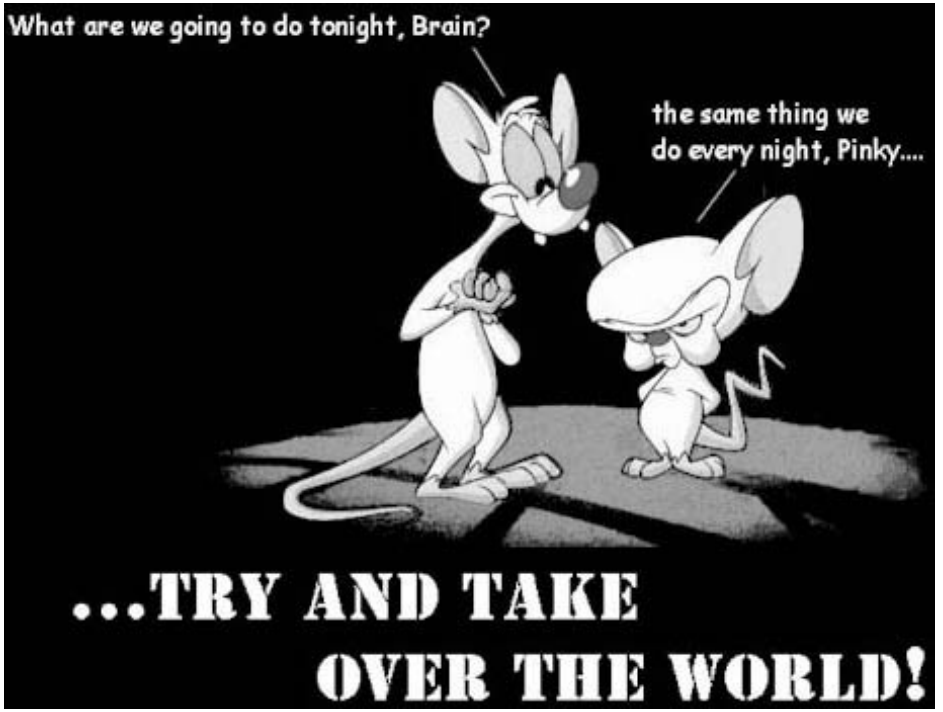


*Sooner or later they're bound to be right*

# Virtual Desktop Infrastructure (VDI)







- IT admins love it
  - Centrally administered
  - Reduced hardware and maintenance costs
- Users will embrace it (hopefully)
  - Familiar personal computing environment
  - *Performance (latency) is critical*

# Improving VDI

- How can we:
  - Reduce the cost of VDI deployments?
  - Improve the user experience?

# Outline

- Background
- How VDI Works
- UBC Workload Analysis
- Design and Implementation
- Evaluation
- Conclusion

# How VDI Works

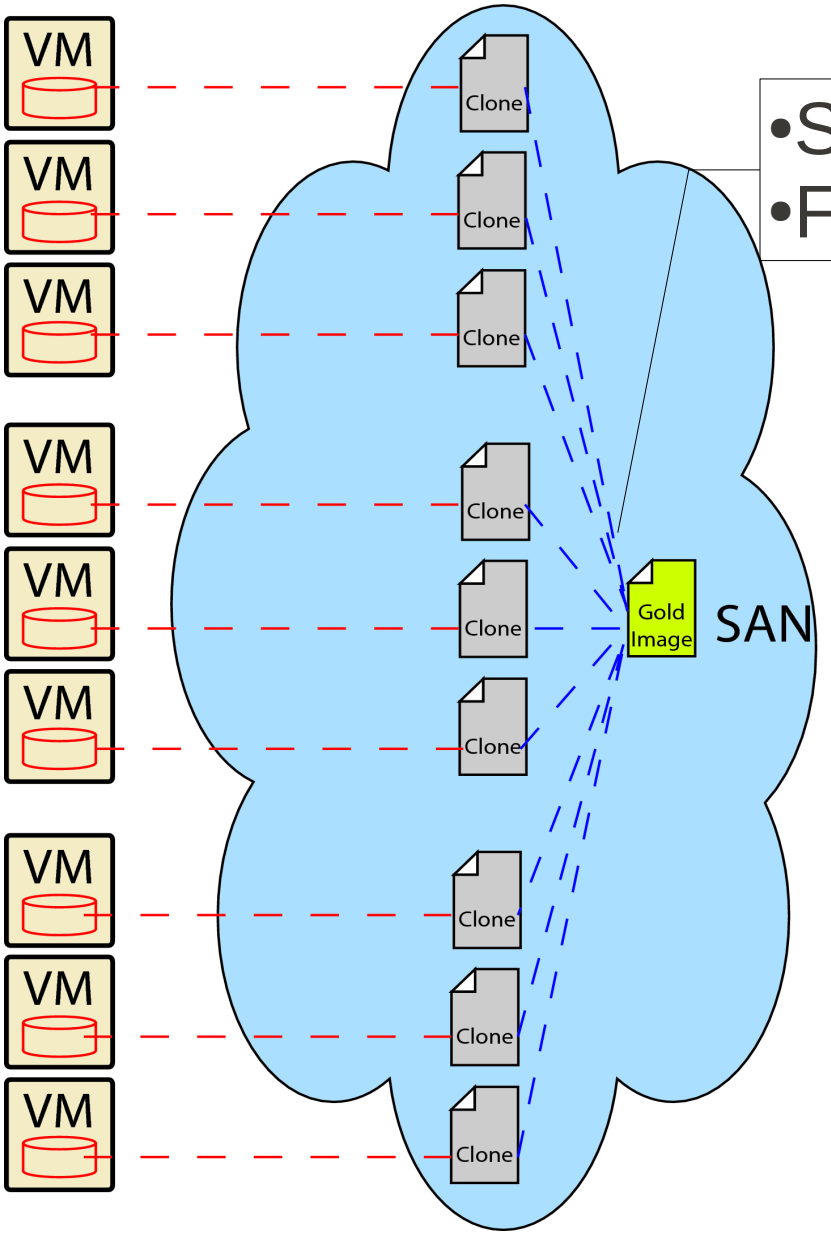


- Sparse allocation
- Fast clone

# How VDI Works

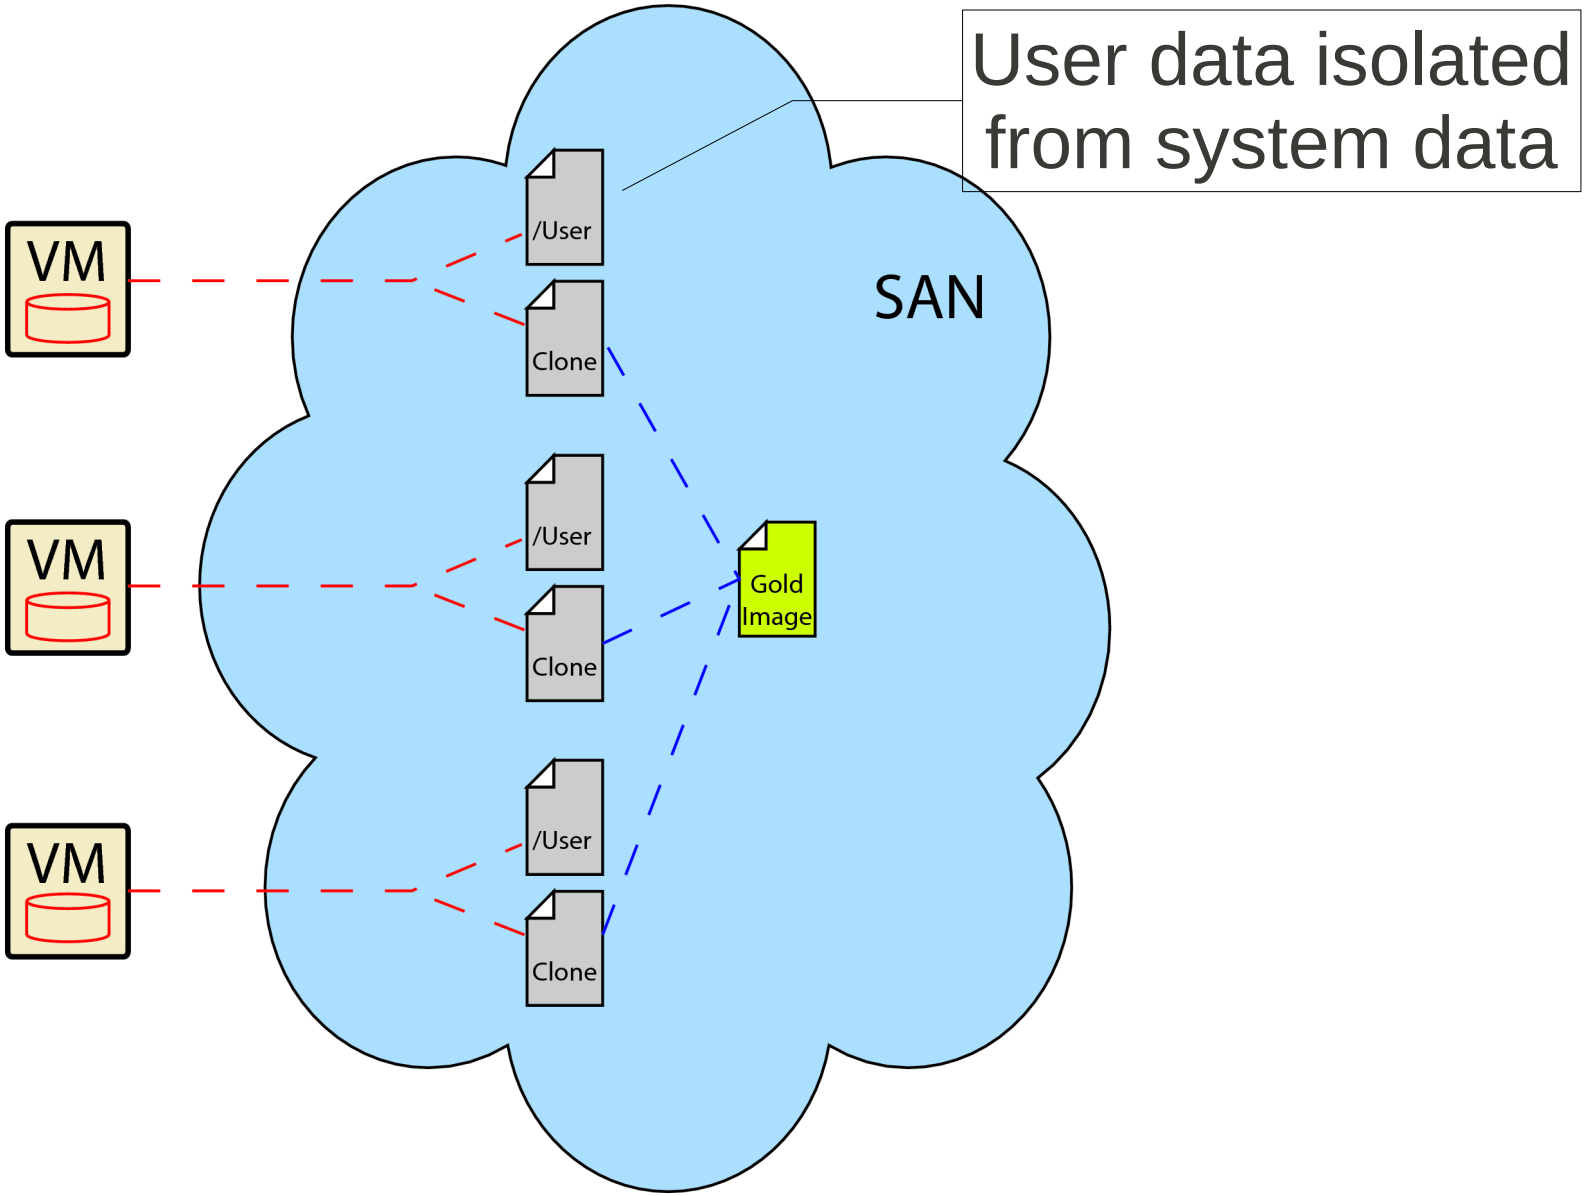

# How VDI Works

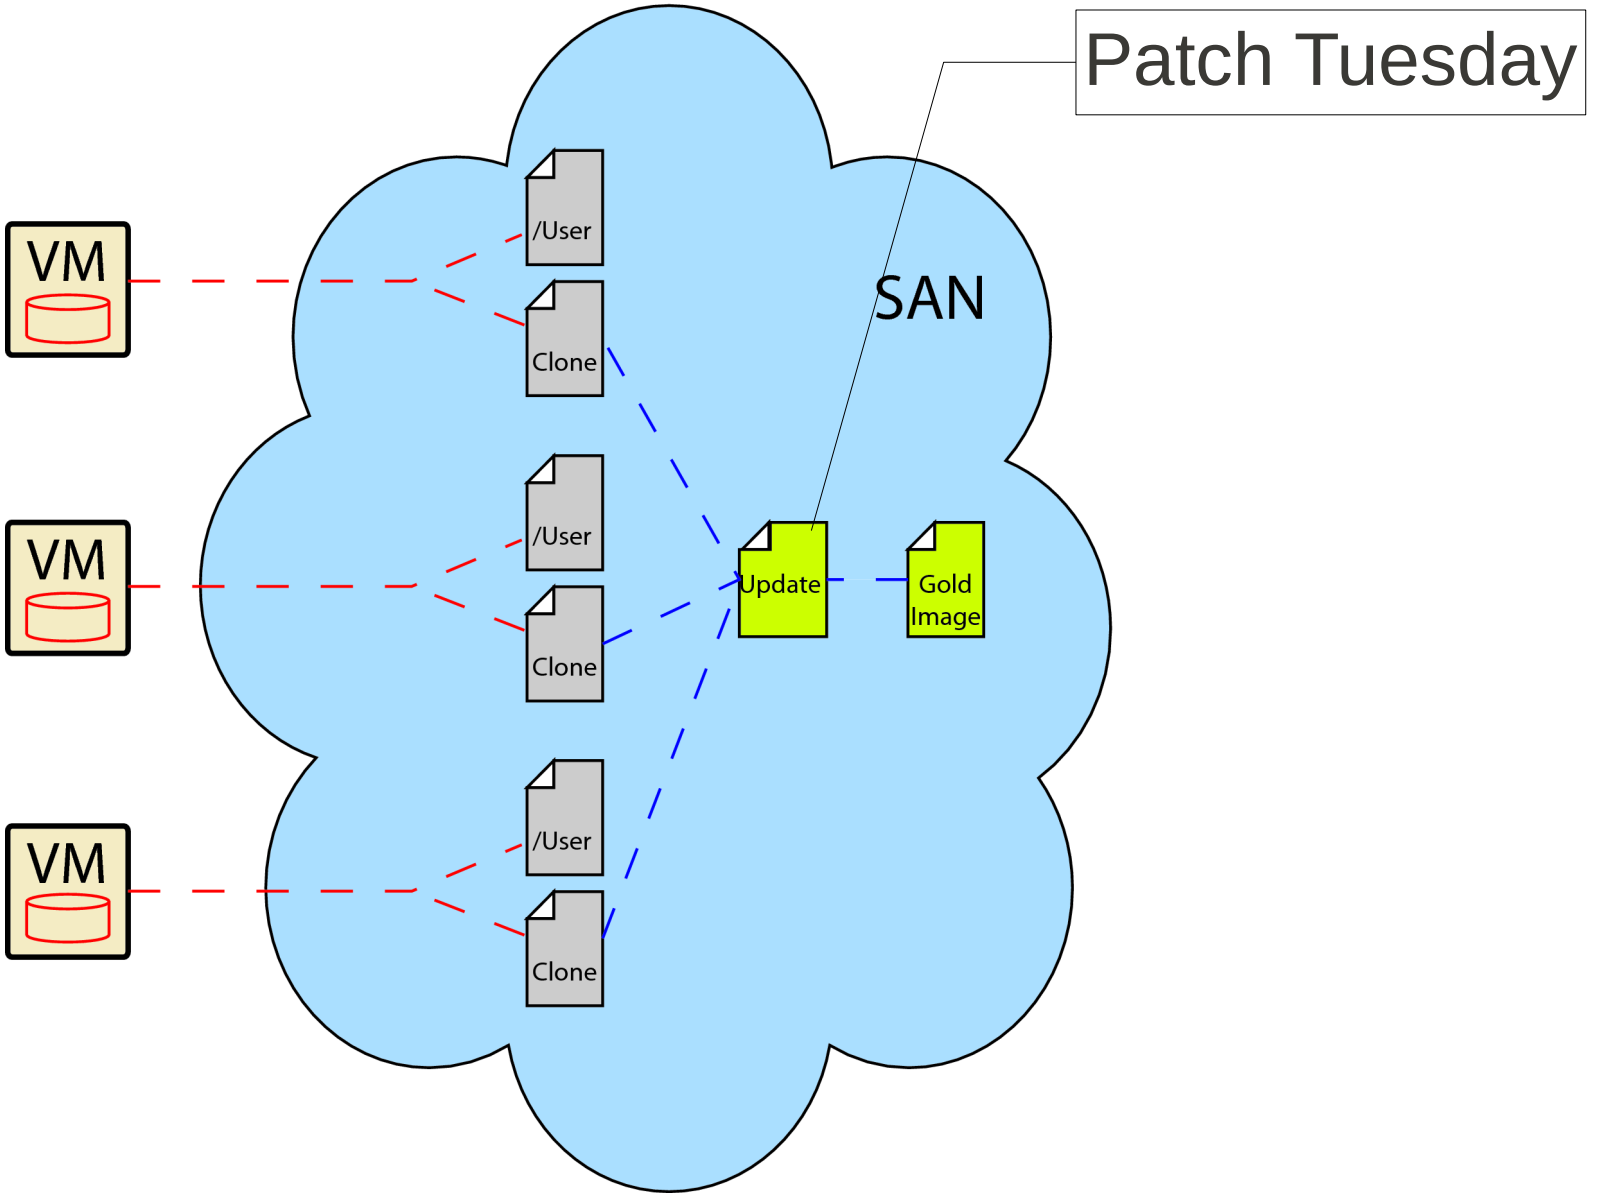

# UBC Workload Analysis


15506-41DG
'Office: 9am' Disc
© JupiterImages

Creatas

www.comstock.com

# UBC Workload Analysis

- We profiled 55 Windows Vista desktops
  - Administrative offices at UBC
  - Installed profiler during regularly-scheduled weekly update
  - Captured file- and block-level accesses
  - Collected 75 GB of compressed, binary logs

# Workload: Day-to-Day Activity

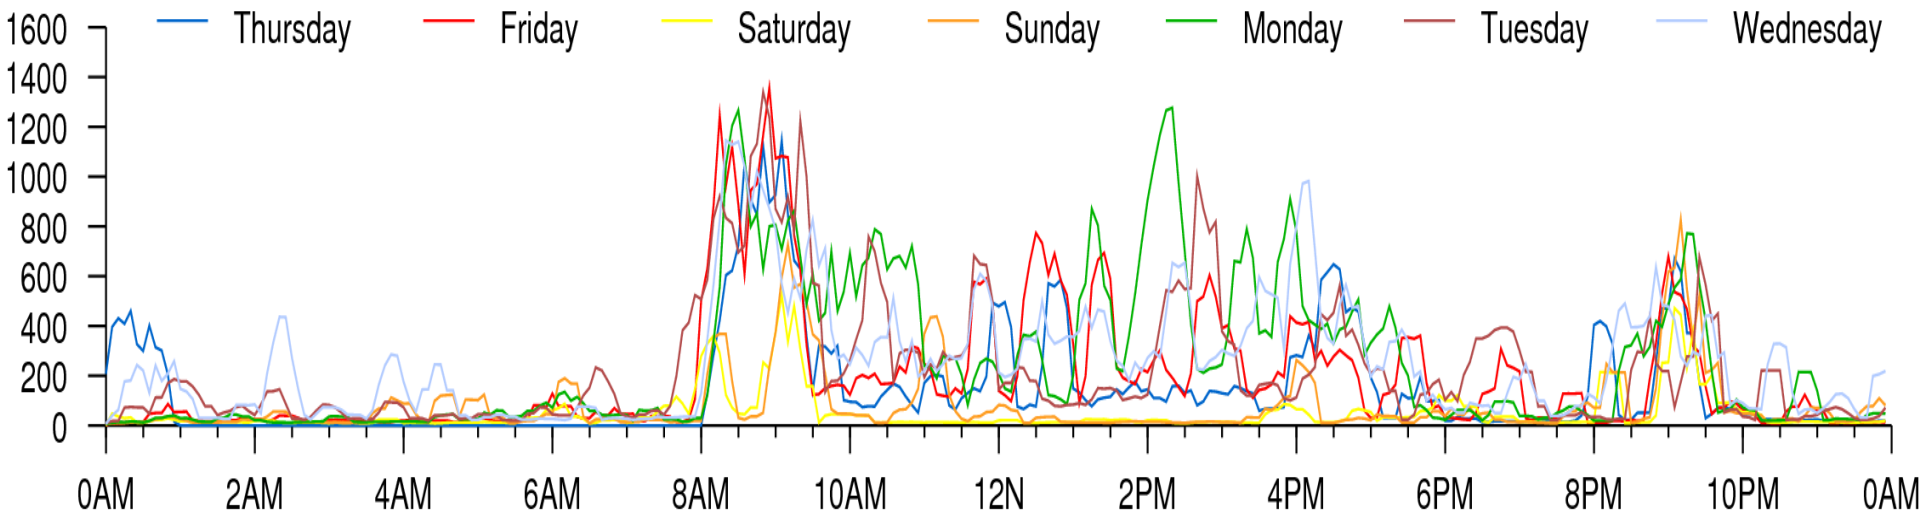


- Workload is quite bursty
- What do the requests look like?

# Workload: IO Requests

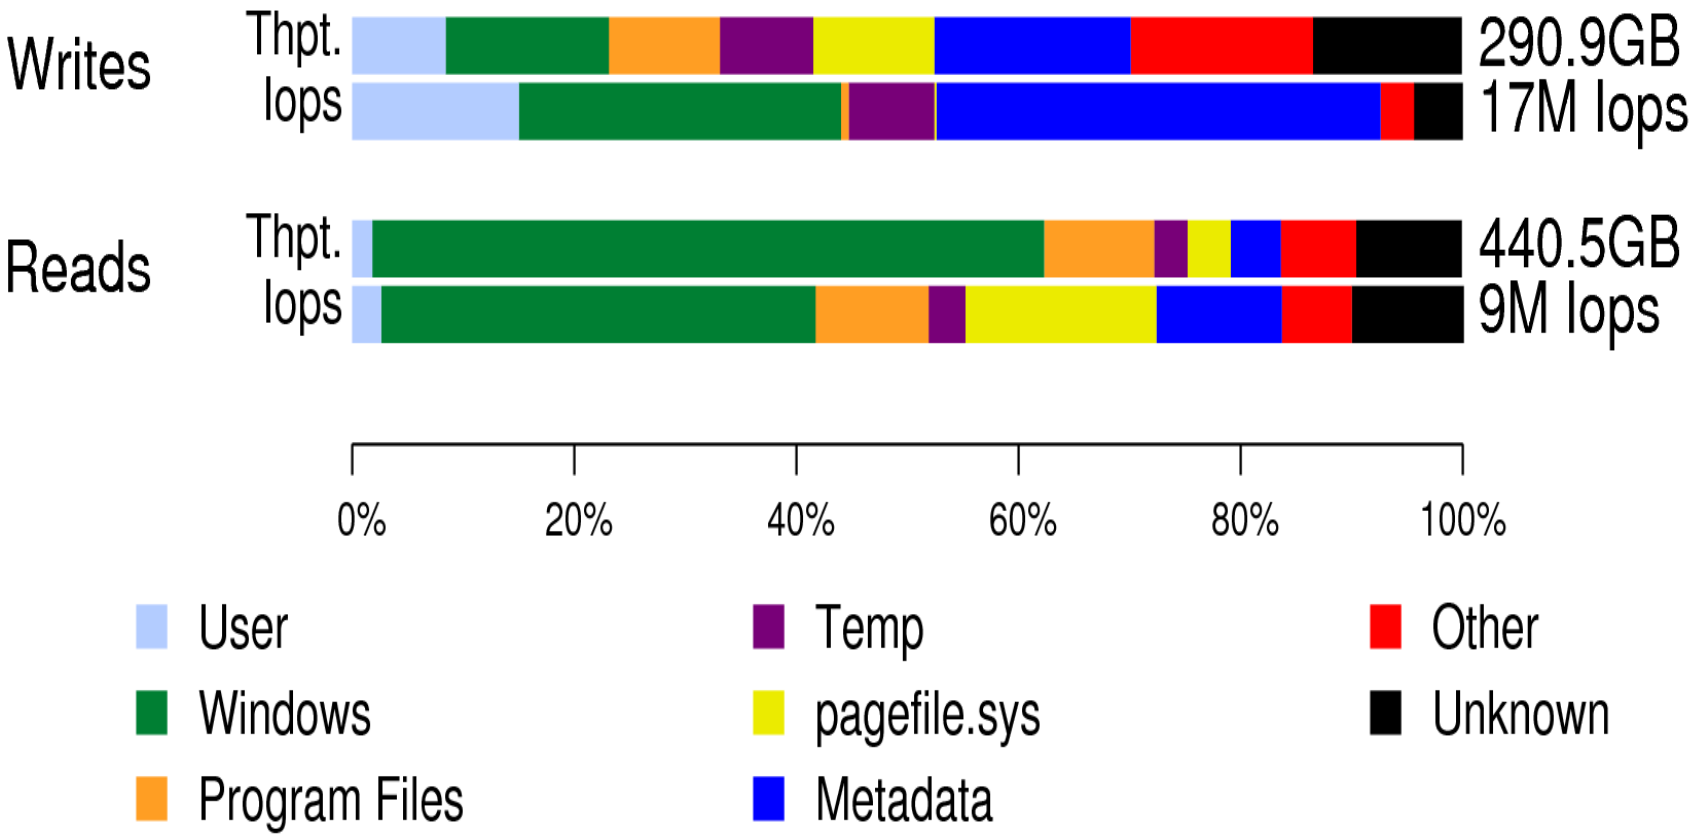

# Workload: IO Requests

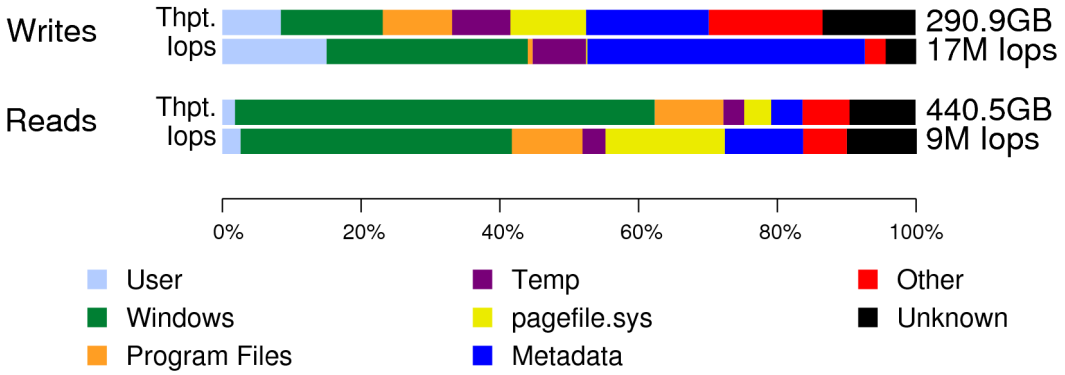


- Most accesses to system-controlled objects (\Windows, \Program Files)
- Metadata-heavy workload
- IOps: 65% writes; throughput: 65% reads
- What do these writes look like?

# Workload: Write Requests

- Fairly high churn rate
  - 8% of bytes re-written in 10 seconds
  - 50% of bytes re-written in 24 hours
- Average divergence of 1GB after about an hour
  - A large portion of this is from pagefile.sys and other Windows files

# Workload Summary

- VDI workloads are bursty
- Significant sharing among VMs
- High churn rate for hot data
- Namespace not accessed uniformly

# VDI Storage Scalability


How can we improve VDI storage scalability?

# Local Persistent Cache

- Goal: *offload IOps to local disks*
- Library in dom0 interposes on access to network files
- Cached files stored on local file system
  - Bitmaps track sparse files
- Supports write-through and write-back with adjustable window

# Multihost Preloader

- Goal: *share local caches among all hosts*
- NFS proxy snoops requests to cached files
- Shared data is multicasted to all subscribed hosts
- Basic congestion control and/or isolated network required

# Differential Durability

- Goal: *reduce writeback burden*
- Data categorized according to value
  - User-created data, installer data, temporary files, pagefile.sys
- Gold Image disk partitioned
  - Valuable data placed on a disk with an aggressive write-through policy
  - Expendable data stored on cheap local disks

# Capo Architecture Diagram

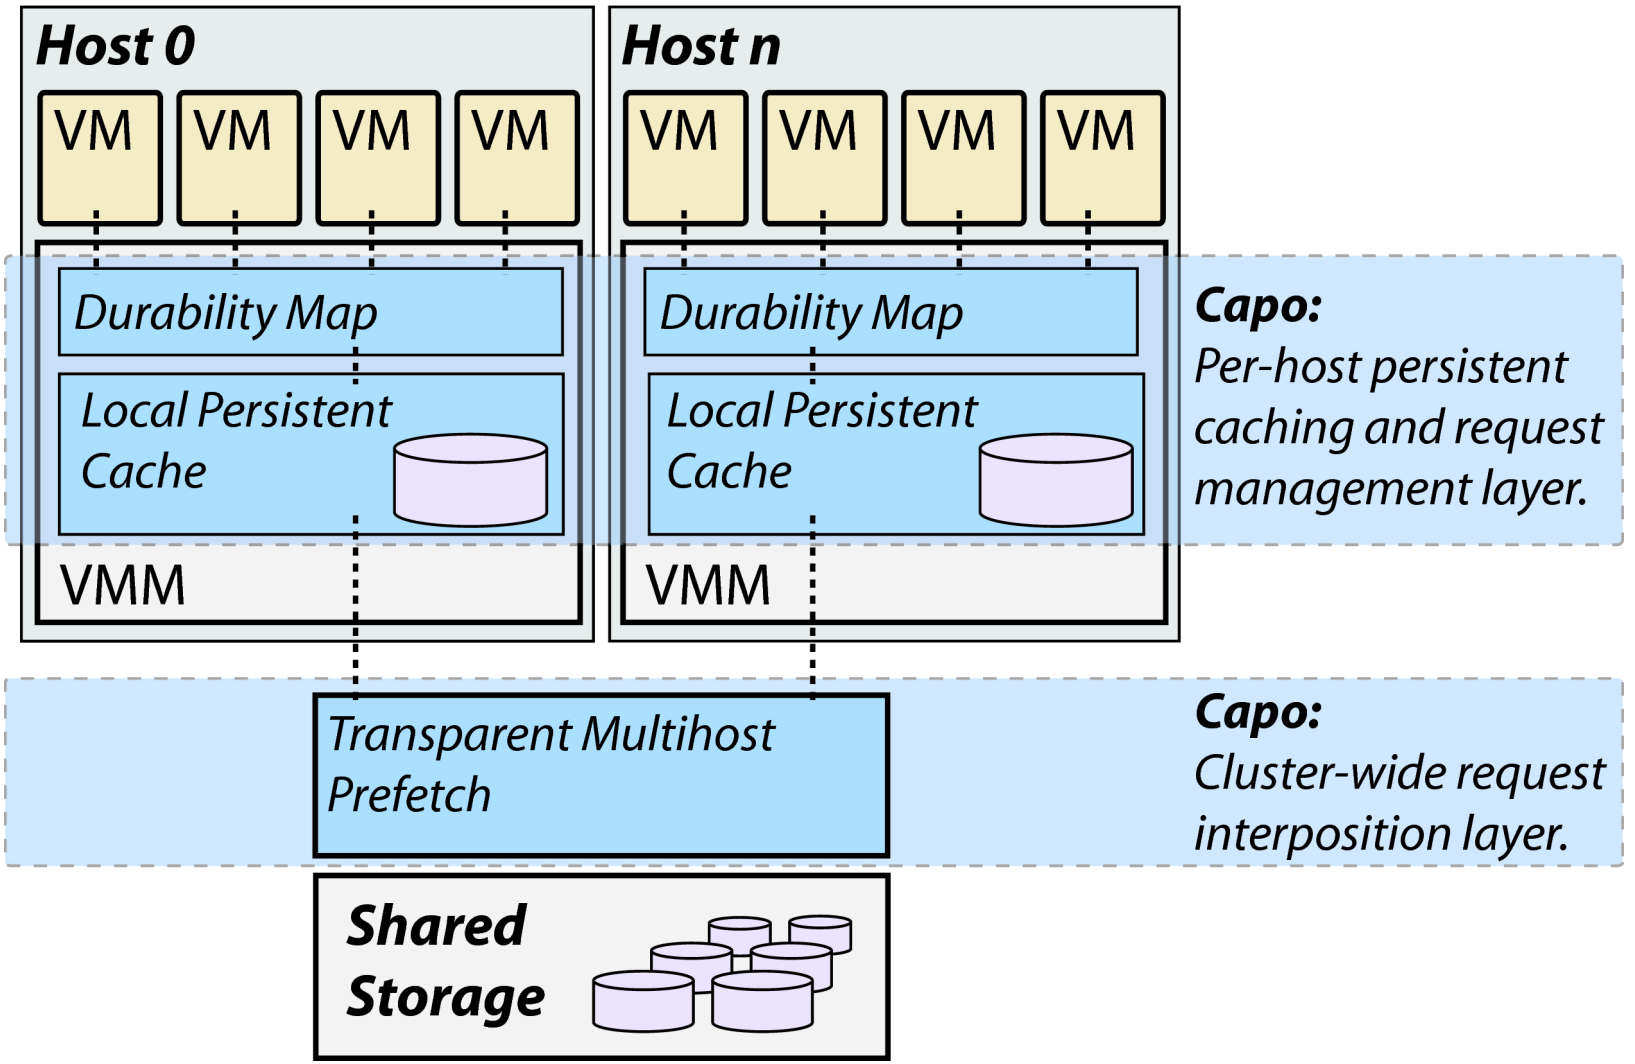

# Implementation

- Local cache prototyped during summer internship
  - 7,000 lines of C code
- Prefetcher implemented three (3) times
  - Packet capture, unfsd, RPC proxy
- IntelliCache™ now available in latest XenServer releases

# Outline

- Background
- How VDI Works
- UBC Workload Analysis
- Design and Implementation
- Evaluation
- Conclusion

# Microbenchmarks: Preloading

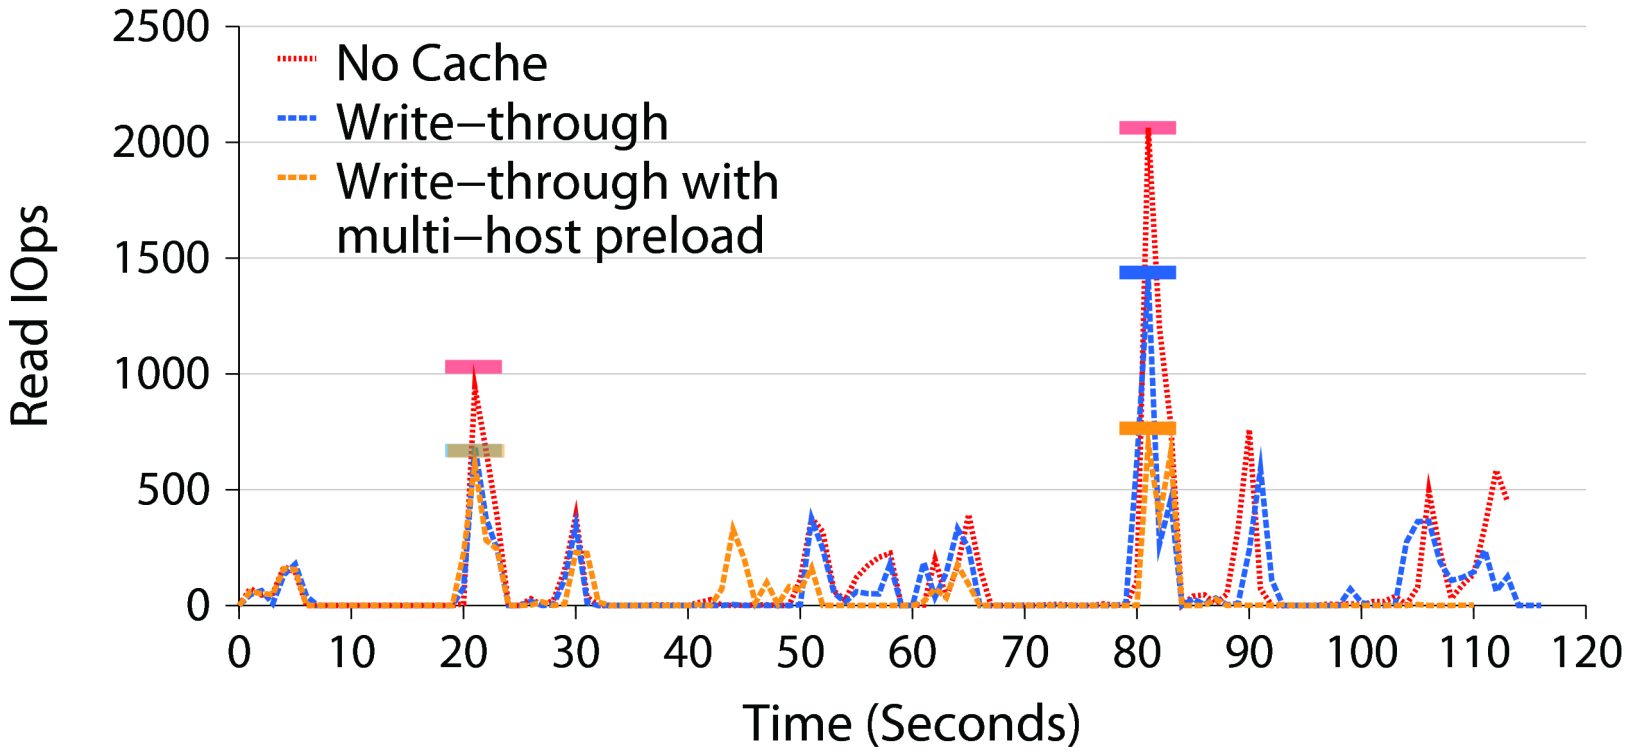

# Trace Replay: Methodology

- Original environment:
  - VMware
  - SAN
- Our lab:
  - XenServer
  - Linux NFS filer w/ 6-disk RAID 0 volume
- Replayer:
  - Simple perl application
  - Tries to match original trace request pattern

# Trace Replay: Selected Peaks

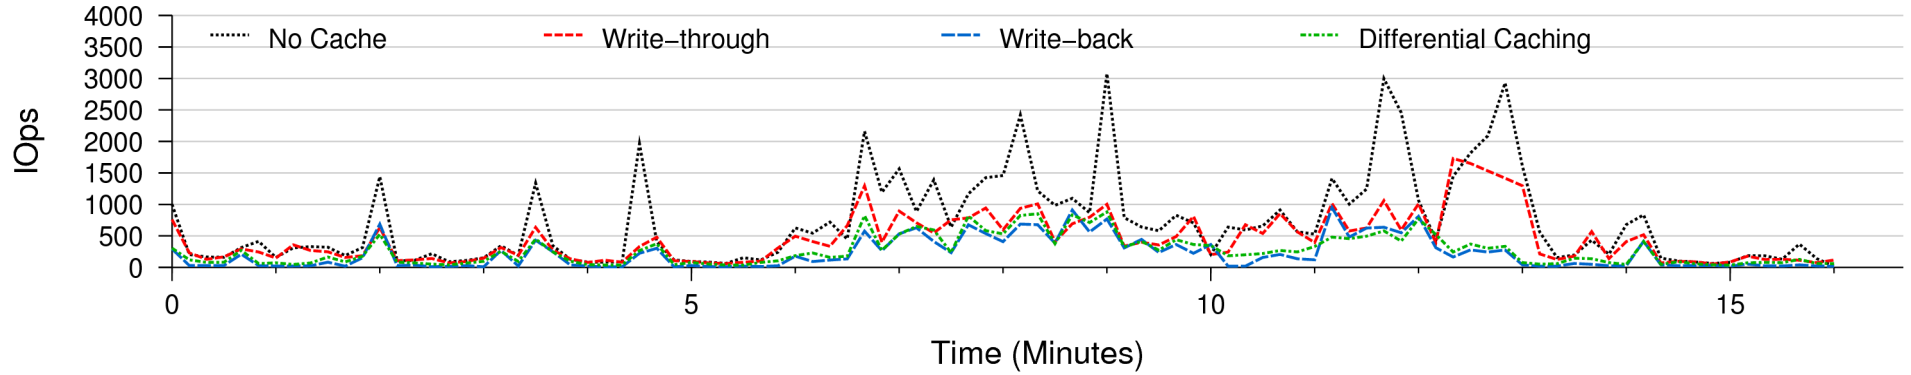


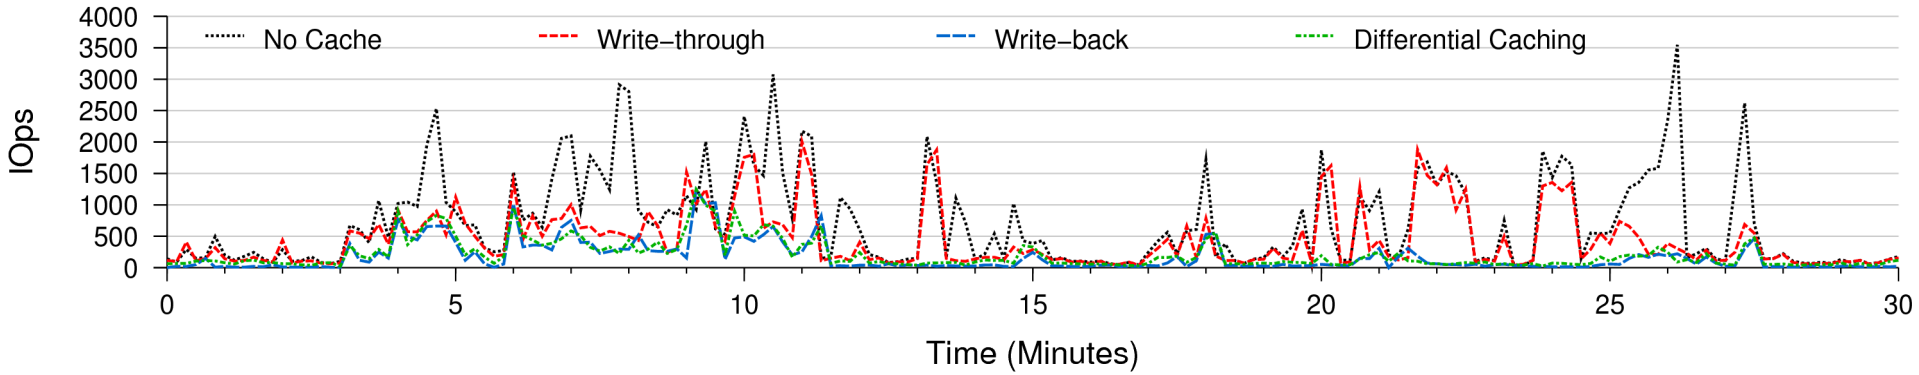

# Conclusion

- VDI presents new challenges for storage systems
- Central storage is a reasonable solution...
- But local caching and differential durability can help reduce costs and improve performance

# Questions?