

Tashi: Open-source Cloud Computing on Big Data

CMU: Jim Cipar, Greg Ganger, Garth Gibson, Julio Lopez, Michael Stroucken, Wittawat Tantisirirotj

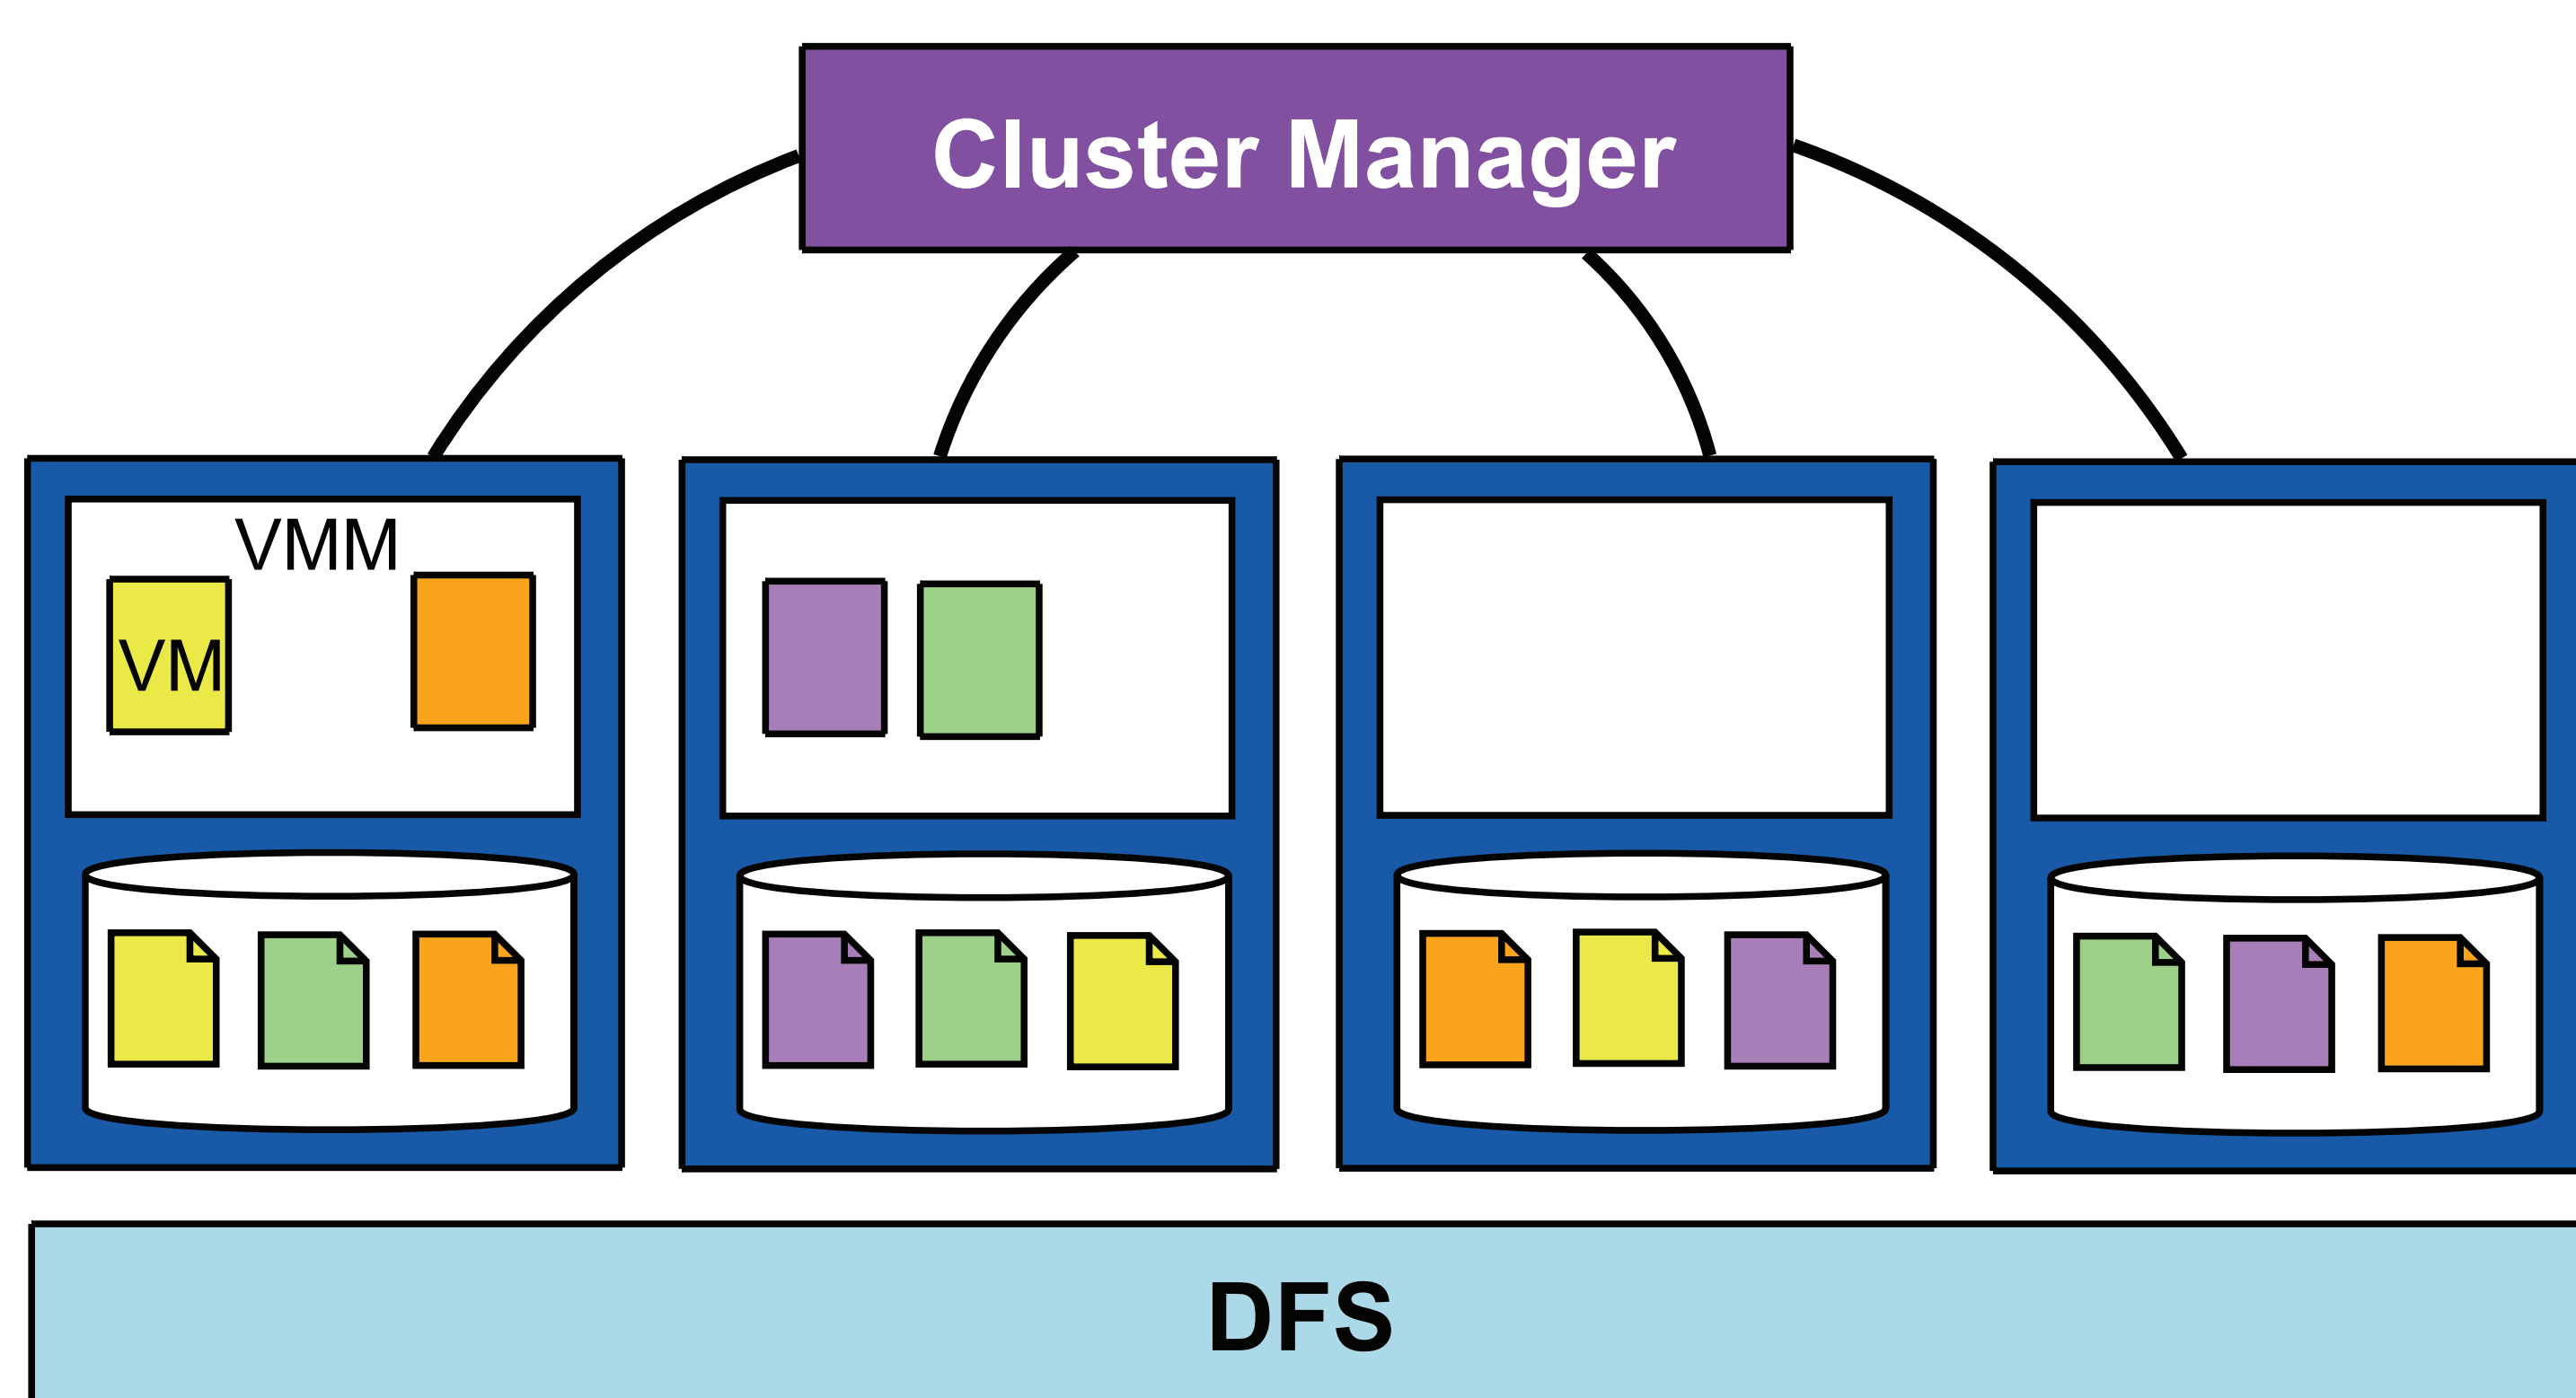
Intel: Dave O'Hallaron, Michael Kozuch, Michael Ryan, Steve Schlosser

Yahoo!: Doug Cutting, Jay Kistler, Thomas Kwan

Overview

- Cloud computing is gaining in popularity
 - An open-source solution is desirable
- Big data applications also growing in popularity
 - Poorly supported by current cloud platforms
 - Currently little coordination between compute & storage
- What do we want in a big data cloud?
 - Applications run in virtual machines
 - Data-aware VM management
 - Expose data layout to application
 - Application behavior affects placement policies

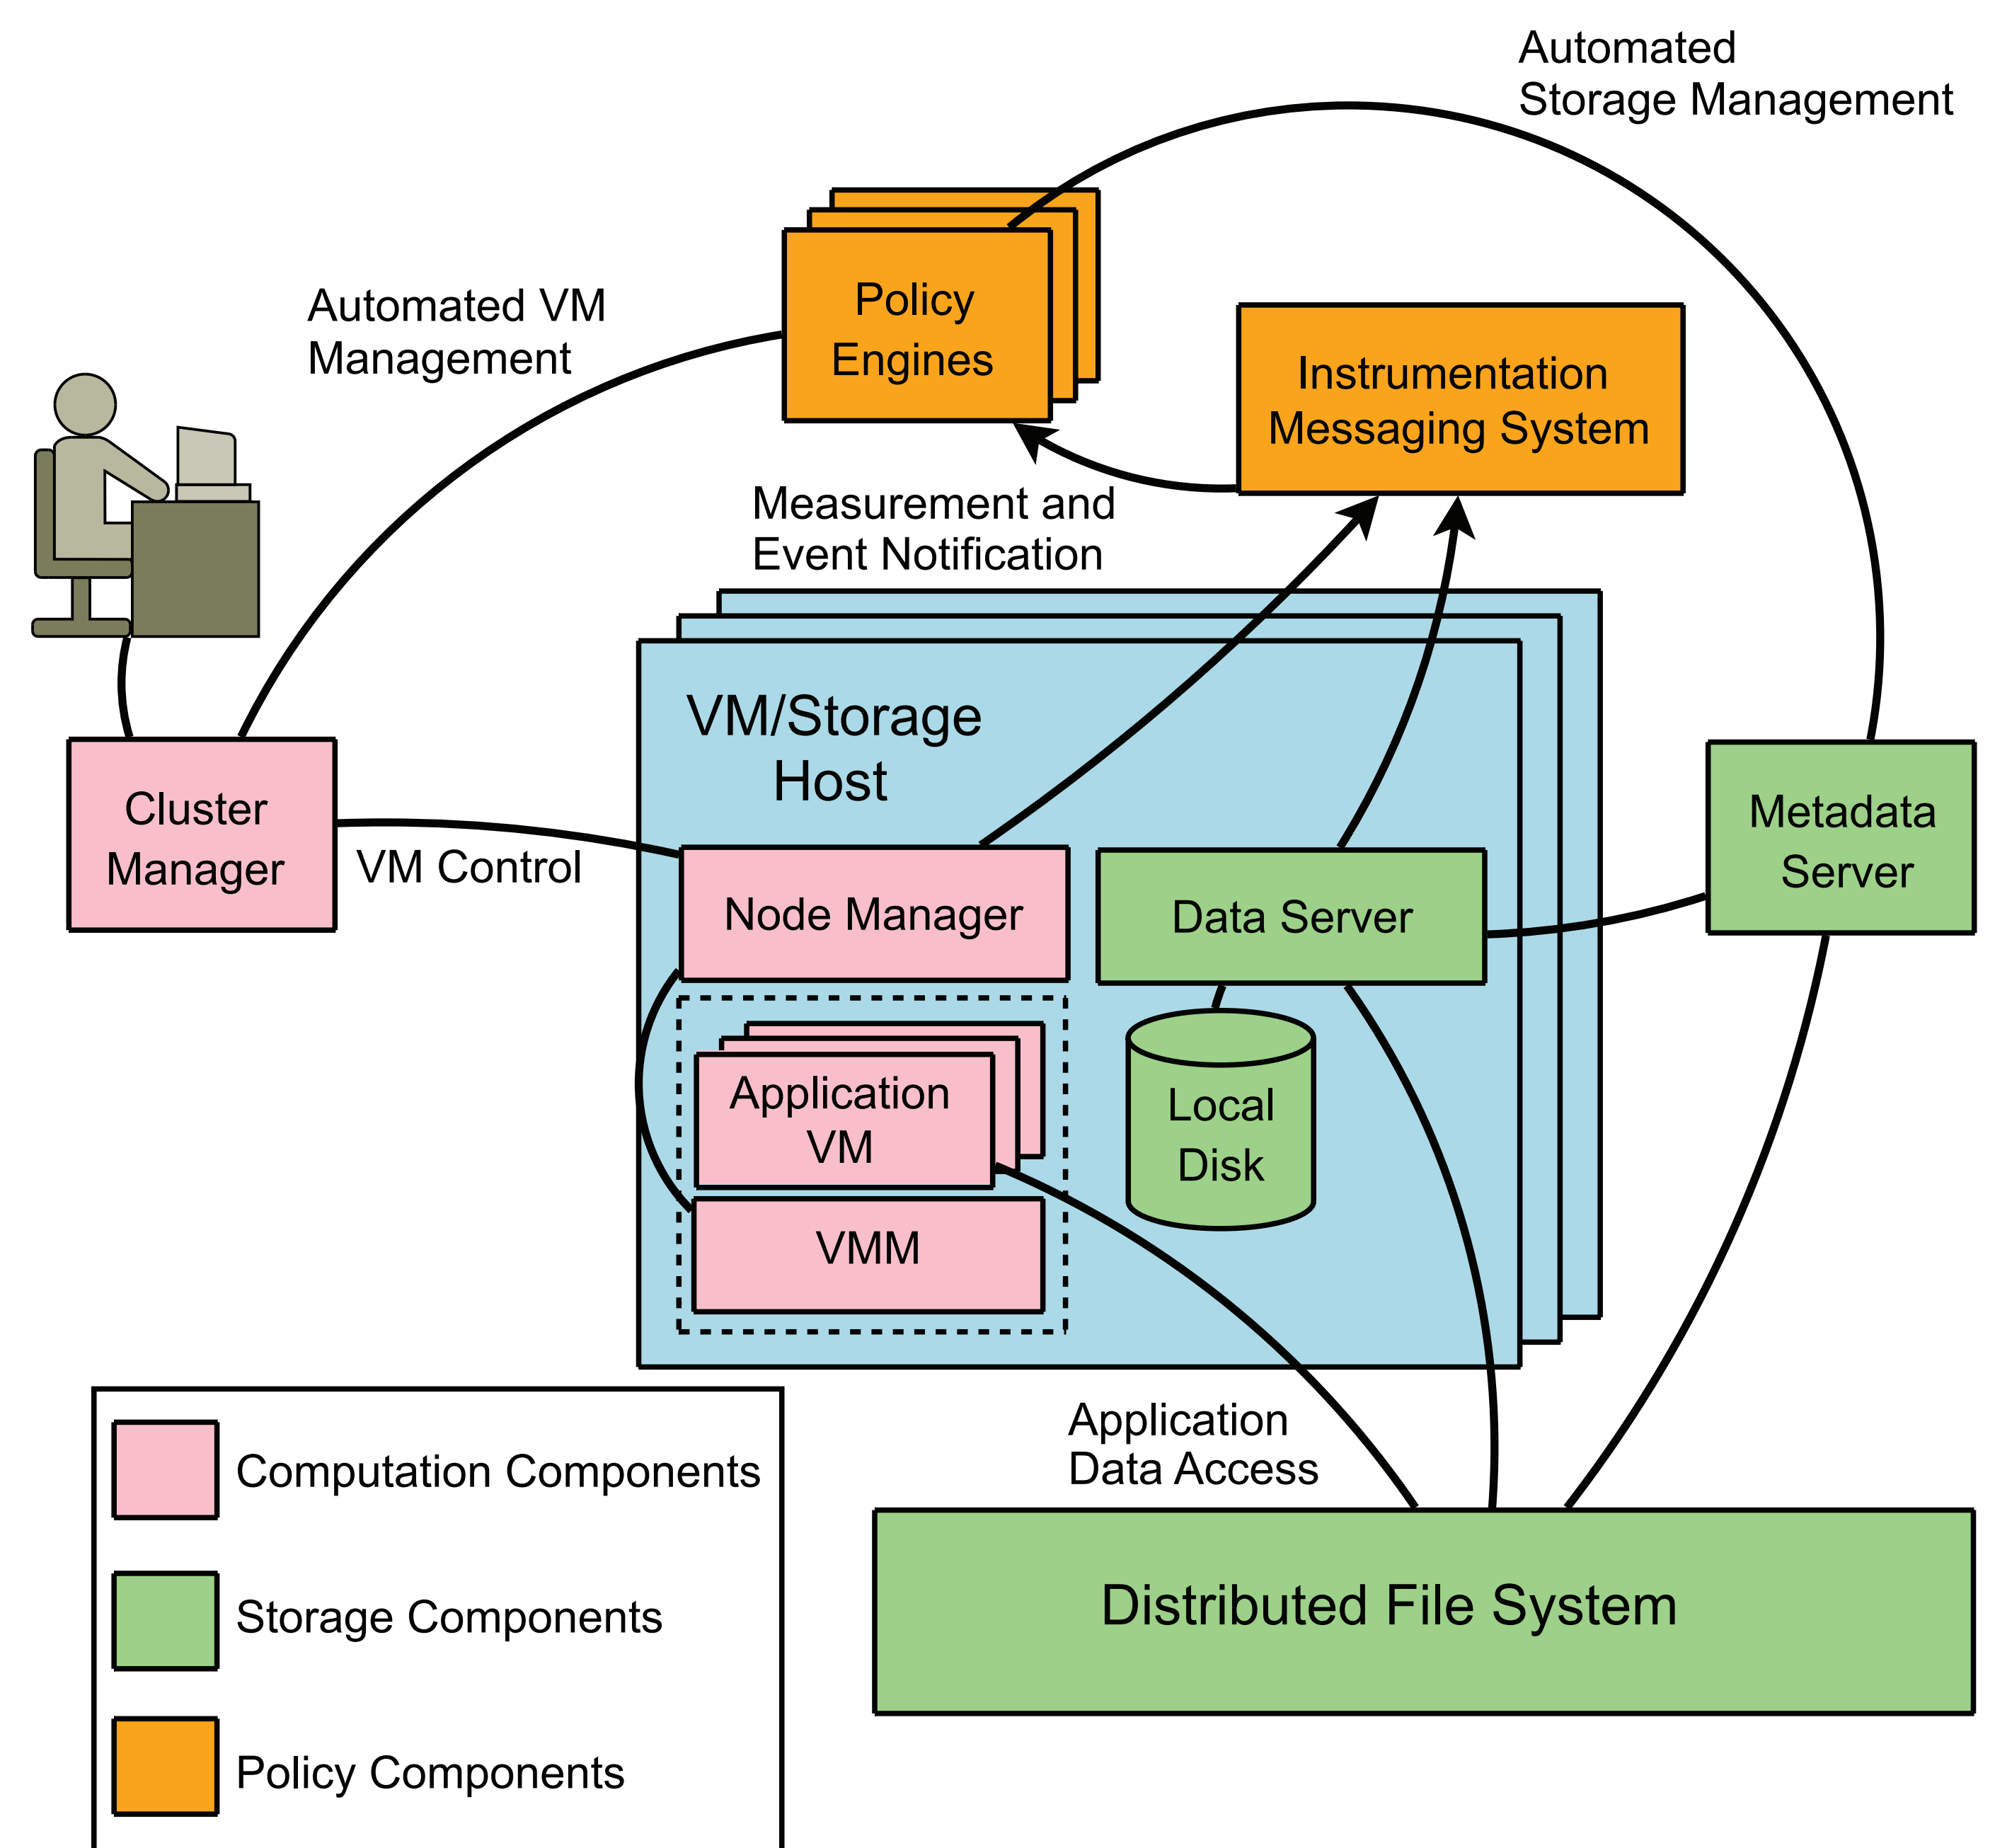
Data and Task Placement



Research Opportunities

- “What goes where?”
 - Trade-offs in co-locating storage and computation?
 - Is it better to replicate or erasure code?
 - How should data be distributed on cluster?
 - How should workload affect future data layout?
- Host availability changes
 - When can servers can be turned off?
 - Can we turn off servers that are hosting data?
 - Can we schedule jobs to keep data servers turned off?
 - When is it better to avoid virtualization?

Architecture



- Cluster manager provides mechanism for managing VMs
 - Keeps track of running VMs, watches for state changes
 - Users, apps., and policy engines all act as clients
 - Client calls CM to start, stop, pause, migrate VMs
- Node manager provides common VMM management
 - Converts calls from CM to native Xen/KVM commands
 - Observes and reports performance metrics
- Instrumentation Messaging System
 - Collects performance data and event information
 - Decouples data producers from consumers
- Policy Engines implement scheduling, migration
 - Monitor events and performance data
 - Act as clients to CM and MDS