Large-Scale Decentralized Object Store

Ricardo Vilaça Rui Oliveira

Computer Science and Technology Center
Universidade do Minho

53rd Meeting of the IFIP 10.4 Working Group
Natal, February 2008

Large-Scale Decentralized Object Store

- Motivation
  - 4x data every 18 months: processing power and storage bandwidth cannot cope with this growth.
  - corporate very large pc parks (100s) are underused: storage < 20%, cpu* <40%.
- Opportunity
  - Controlled P2P environment: same owner (even if multiple administrative domains), reasonable stable membership, particular business model.
  - Massive distributed (object) data management with in-place processing capabilities.

Large-Scale Decentralized Object Store

- Challenges
  - Dependability in face of massive distributed data: availability given by "massive" replication, consistency becomes a "massive" problem
  - Performance is highly dependent on the adaptiveness of the system:
    - data migration vs. aggregation
    - local vs. distributed execution
    - impact of reconfiguration
- Security: Resilient Computing MSc.

Large-Scale Decentralized Object Store

- Current status
  - Promising architecture leveraging the GORDA project results on consistent group-based DB replication
  - Incipient (campus wide) Last.FM prototype based on JXTA + HSQL with streaming and format conversion of MP3
- Exploitation
  - Content-aware fragmentation BT protocol
  - Sleepycat/G storage engine