Towards a Scalable, Fault-tolerant, Self-adaptive storage for the clouds

Houssem-Eddine Chihoub

INRIA Rennes, France

Houssem-eddine.chihoub @inria.fr

Advisors: Gabriel Antoniu
INRIA Rennes, France
gabriel.antoniu@inria.fr

Maria S. Pérez Hernandez
Universidad Politécnica de Madrid, spain
mperez@fi.upm.es

Introduction

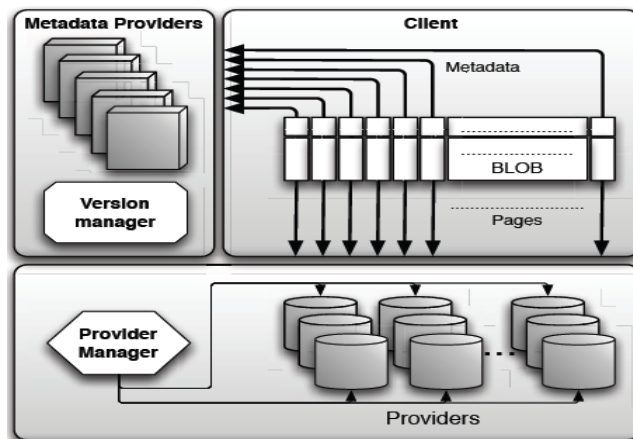**Context** Data-intensive applications on **clouds**

Build an efficient storage support for clouds

Challenges for storage

- Design scalable storage architecture
- Support huge file sharing with fine grain access
- Sustain high throughput under heavy concurrency

Bricks

BlobSeer

Key Features
- Data striping
- Distributed metadata management
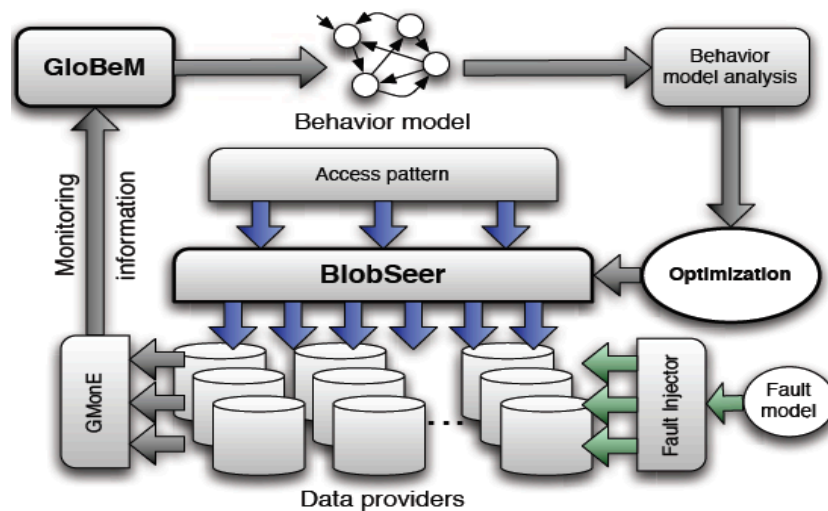- Versioning based concurrency control

GloBeM

Global **B**ehavior **M**odeling

Modeling the global behavior of a large scale distributed system :
- Observing the system
- Analyzing the data
- Building the model

Approach: BlobSeer + GloBeM

Enhance **QoS** of BlobSeer:
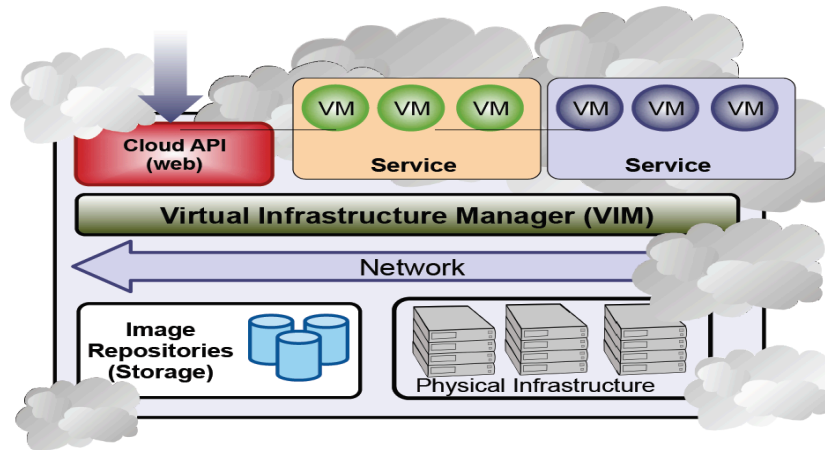Higher more stable throughput for data transfers

Methodology:
- Component monitoring
- Application-side feedback
- Behavior pattern analysis

Case study: OpenNebula Cloud

Flexible and extensible tool for building scalable Cloud environments

Storage for :
- **VM images :** NFS, SSH, LVM ...
- **Application Data :** not yet!

Roadmap
- Storage support designed for VM images
- Storage Support for Applications data
- Experimentations and comparaison with other file systems on G5K
- Refining the behavior modeling phase to fit OpenNebula requirements

Thank you for your attention!