Evaluation through Realistic Simulations of File Replication Strategies for Large Heterogeneous Distributed Systems

A. Chai, S. Camarasu-Pop, T. Glatard
H. Benoit-Cattin, F. Suter

Presented at Heteropar Workshop in EuroPar 2018

CREATIS, INSA-LYON
CC-IN2P3

ANR-11-LABX-0063/ ANR-11-IDEX-0007

Context: The **V**irtual **I**maging **P**latform

- A science gateway for medical imaging
  - >20 applications
  - >1,000 registered users
- Supported by the **BioMed** VO within EGI
  - 130 computing clusters
  - 65 sites worldwide
  - 50 different Storage Elements (SEs)

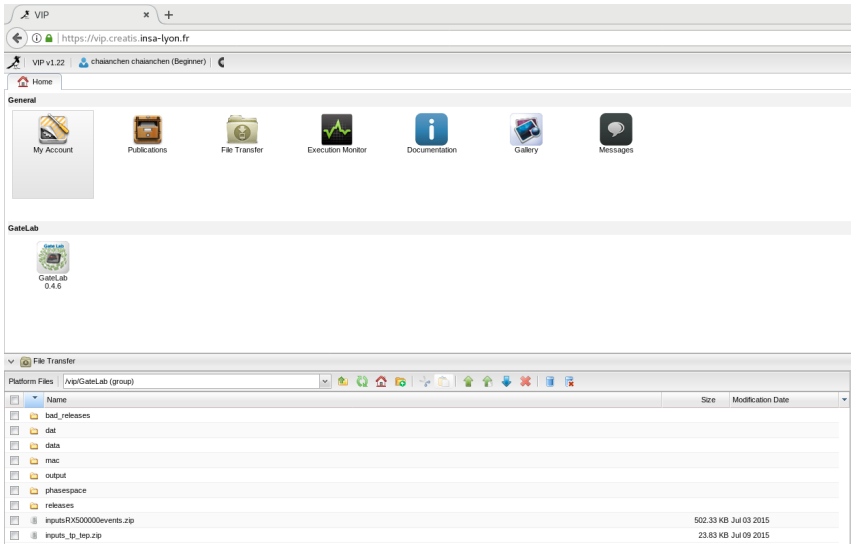

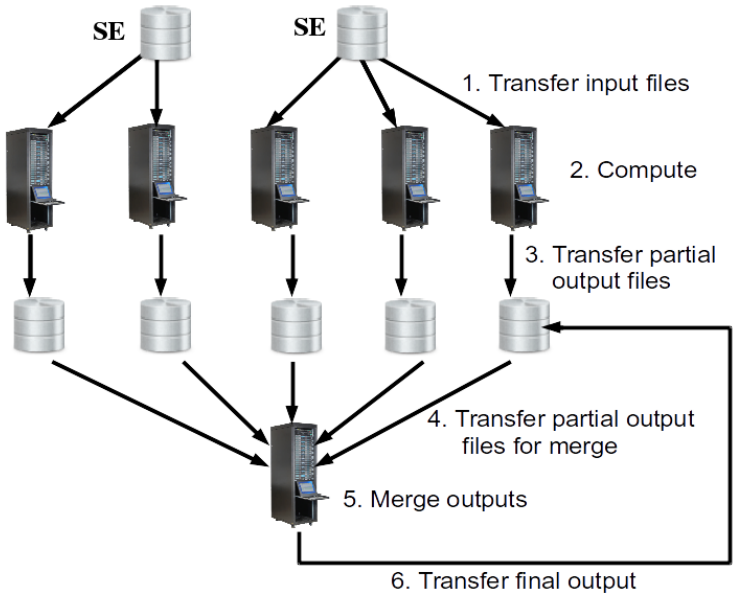

Typical execution model in VIP

- Computation distributed across multiple sites
- Numerous SEs used to transfer files
- Focus on file transfers for applications

File replication

Consists in copying the same file onto multiple SEs.

- Avoid single-point failures $\rightsquigarrow$ higher availability
- Distribute network traffic $\rightsquigarrow$ shorter transfer duration

2 aspects:

- Replica creation: where and how many times to replicate a file
- Replica selection: how to choose the "best" replica

File management on EGI

- APIs to upload, replicate, and register files
- Replica selection according to the geographical distance
- Decisions about replica creation are left to applications
  - e.g., in VIP, files replicated on 3-5 stable SEs from a predefined list

Why do we need simulation?

- Hardly to evaluate the impact of file replication directly in production
  - time-consuming, non-repeatable, harm production...

From real system to simulation

Reliability of simulation can be improved by feeding real information

- Platform model
- Instantiation of parameters

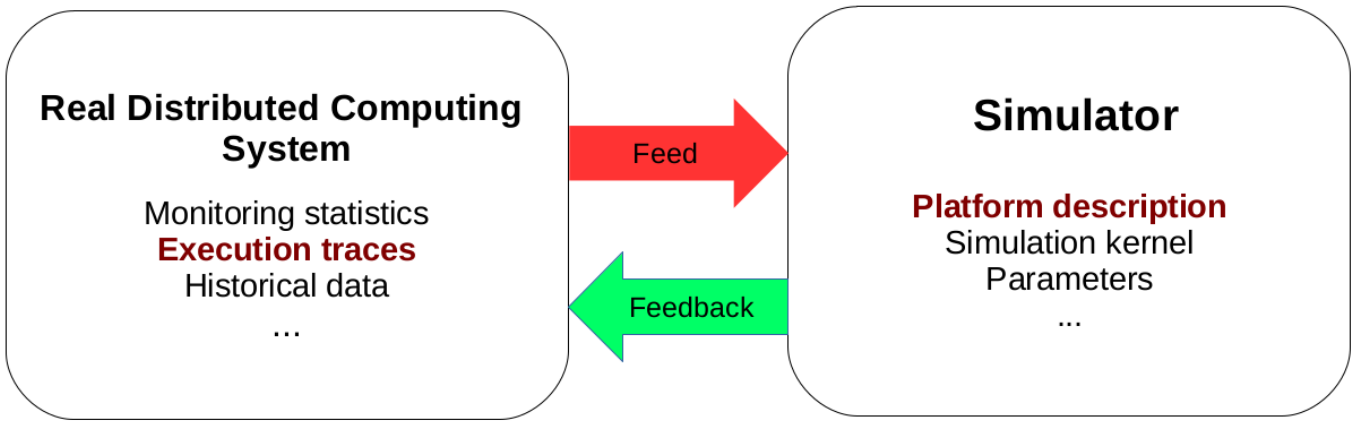

From simulation to real system

Real system can be optimized by reliable feedback from simulation

- Test different replication strategies
- Full-scale evaluation

Why do we insist on "realistic"?

Limitation of simulation studies in literature

- Validation: questionable applicability to real systems
- The most used metric is file transfer durations
  - But, platform models are often oversimplified

Our proposal

- Design a platform model built from real execution traces
- Parameter injection in simulation scenarios
- Cross-evaluation with the state-of-the-art model

Questions to answer

- What is the impact of replication strategies on file transfer duration?
- Does the impact change on different platform models?
- What would be reliable recommendations for file management in VIP?

Plan

Replication strategies

Static replication

- Decisions made before application starts and invariant during the execution of application
- How to choose SEs to have the best combination?
  - one combination of SEs = one prestaging list
- Random vs. based on statistical information on where jobs were executed

Observations

- a given site executes more than one job
- large variable queuing time on EGI

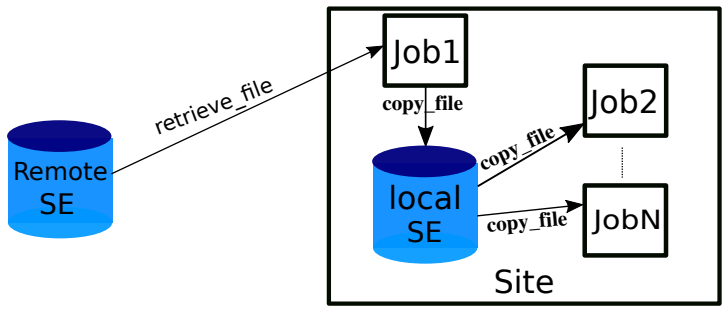

Dynamic method

Dynamic replication

- Decisions adapted to changes of the system during the execution of application
- Our proposal : a practical method inspired from "cache hit"
- The first job in a site replicates the downloaded file onto its local SE

Platform models

- 3-level hierarchical model
  - A widely used model in literature
  - 3-level network hierarchy: local, national, and inter-country
  - Unique theoretical bandwidth for each link category
  - Proposal: Enhance by average bandwidth derived from traces
    - local: 1.3 Gb/s, national: 255Mb/s, and inter-country: 100 Mb/s
  - Pros: Simplicity to build a complete network model
  - Cons: Hardly capture the heterogeneity of EGI
- Trace-based model
  - Derive network bandwidth from measured transfer durations
  - Spatial aggregation of multiple traces
    - Pros: Reproduce real-life variability
    - Cons: Construction complexity and missing links
  - Proposal: Fill missing links by leveraging existing information

Plan

Simulation of file transfers

- Realistic information extracted from execution logs
  - 15 execution instances of the typical workflow
  - Each one consists of 100 jobs
  - The queuing time and execution site of jobs (injected in simulation)
- Prestaging list = 3 SEs (among 32 SEs) for static replica creation
- 55 different prestaging lists
  - Current production setting: 3 SEs in France
  - Without a priori information: 50 random lists of 3 SEs
  - With a priori information: 4 lists based on the job distribution
- 220(2*2*55) scenarios for each workflow
  $\rightsquigarrow$ 330,000 (220*15*100) simulated file transfers

Q1: Impact of dynamic replication

- 50 random lists of 3 SEs for both models

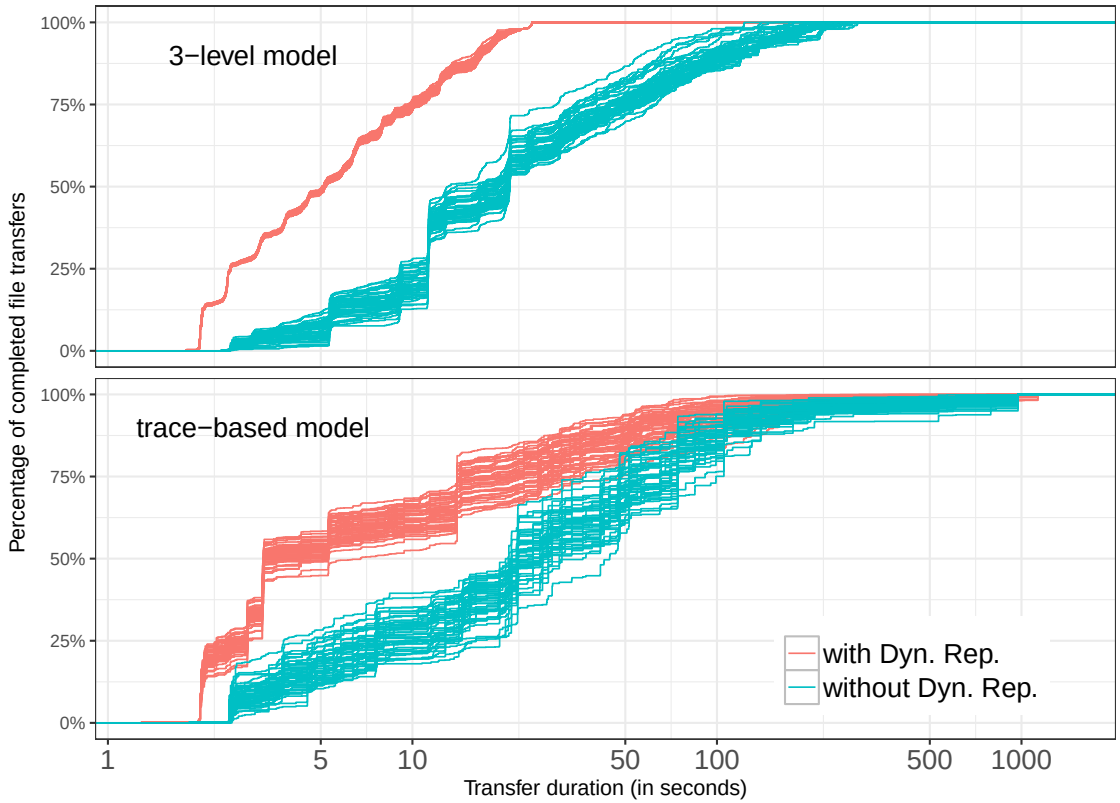

- Dynamic replication decreases transfer durations for both models
  - Larger improvement in 3-level model
- Prestaging lists have stronger impact without dynamic replication

Q1: Impact of different prestaging lists

- 50 random lists + 4 predefined list + the current production setting

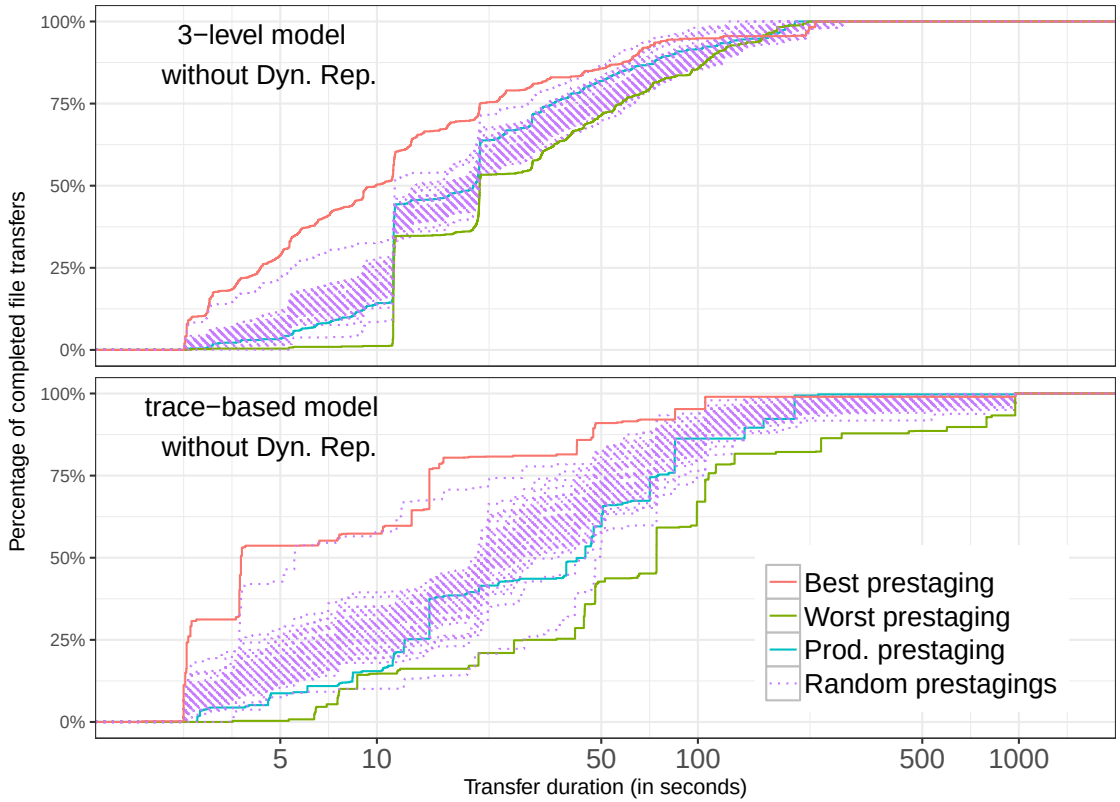

- Best prestaging list is the same for both models
- SE list for best prestaging $\rightsquigarrow$ sites with most jobs in 3 countries
  - Replicate close to jobs and scatter in different countries

Q2: Impact of platform model

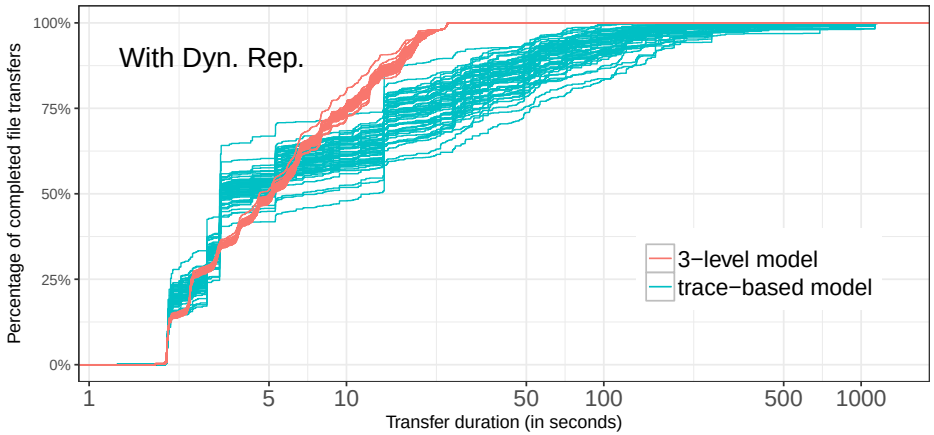

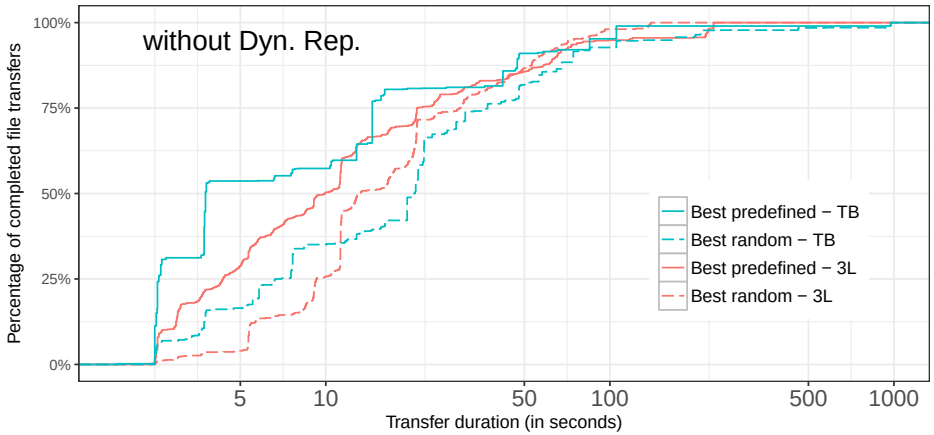

With Dynamic Replication

- 3-level: random prestaging is enough
- TB model: large performance variability depending on the SEs

Without Dynamic Replication

- With a-priori information, the best predefined list $>$ the best random list
- Improvement is larger in TB model

- More heterogeneity in the platform $\rightsquigarrow$ more gain comes from a-priori knowledge

Q3: Recommendations for file replication in VIP on EGI

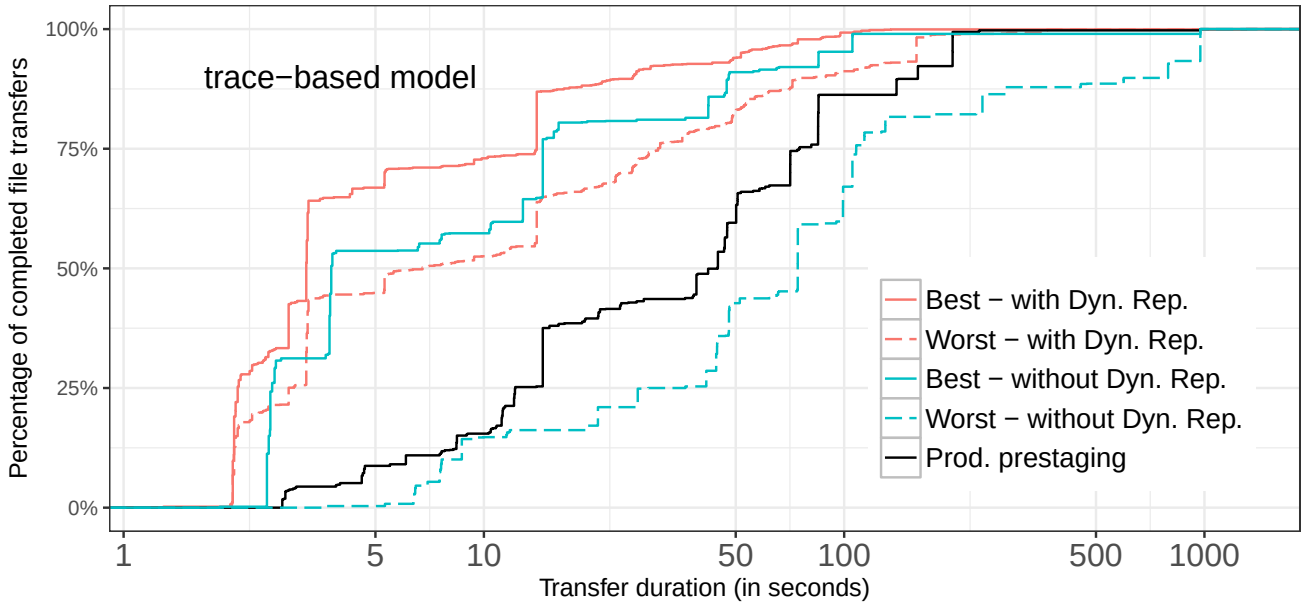

- Without Dynamic Replication:
  - Replicate files close to the sites executing most jobs
  - Scatter across different countries
  - Best results are comparable to dynamic replication
    - requires to estimate where jobs will be executed
- With Dynamic Replication:
  - Always better than Prod.prestaging
  - Reliable method to reduce the file transfer durations

Plan

Take-home messages

Realistic simulation

- Different platform models $\rightsquigarrow$ different qualitative assessment
- Realism of the platform model is key to the evaluation process
  - And also to the applicability for production systems

Recommendations for file management on EGI

- Replicate data close to jobs
  - Requires a-priori information
- Further optimization by dynamic replication
  - Non-negligible implementation cost
- Scatter across countries is a safe strategy
- A-priori information is very important for a heterogeneous platform

- Still lots of aspects to investigate

Future work

- Study the impact of the number of prestaging SEs
- Further improve the accuracy of trace-based model
  - Integrate temporal aggregation of traces
  - Investigate different methods to fill missing links
- Consider other parameters in simulation scenarios
  - Transfer failure rates
  - Available storage space
- Build probability distributions out of traces for these parameters

Thank you!

- Companion of article: http://doi.org/10.5281/zenodo.1239677
- VIPSimulator: https://github.com/frs69wq/VIPSimulator
- Simgrid: https://simgrid.gforge.inria.fr/download.php
- VIP: https://www.creatis.insa-lyon.fr/vip/
- EGI: https://www.egi.eu/

Fill missing links

Hypothesis

- For a site S_i, the connectivity to SEs follows the network hierarchy
- Measured links to/from S_i could reflect its general connectivity
- Reliability depends on the number of known links

$$C_i = \sum_c \left(\frac{|L_i^c|}{|L_i|} \cdot \frac{\widetilde{B_i^c}}{\widetilde{B^c}} \right) \tag{1}$$

- Bandwidth of a missing link of category c to/from S_i: $\widetilde{B^c} \times C_i$
- Category c: local, national, and inter-country
- $\widetilde{B^c}$: median bandwidth of all known links in category c
- $\widetilde{B_i^c}$: median bandwidth of the known links to/from S_i in category c
- $|L_i|$: the total number of known links for site S_i
- $|L_i^c|$: the number of known links to/from S_i in category c

Backup information

4 predefined prestaging list of 3 SEs

- SEs of the top 3 sites with most executed jobs in one country
- SEs of sites with the most executed jobs in 3 country
- SEs of 3 sites with no executed jobs in one country
- SEs of sites with no executed jobs in 3 country

Replica selection on EGI

- First, the local SE of the computing site
- Then, the SE in the same country as the job execution
- Randomly among all available replicas