

BtrPlace

A Flexible Consolidation Manager for
Highly Available Applications

To appear in IEEE TDSC

Fabien Hermenier

University of Nice-Sophia Antipolis

fabien.hermenier@unice.fr


Julia Lawall

INRIA/LIP6

Gilles Muller

INRIA/LIP6

N-Tiers applications




Users are looking for:

- performance
- reliability
- isolation
- ...

Where to place the VMs?

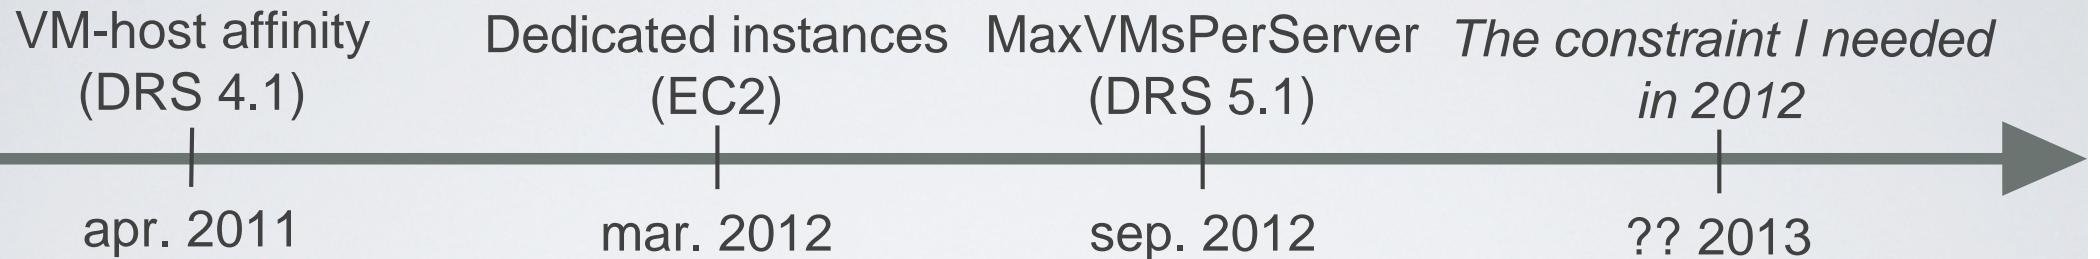
Datacenter management



Operators are looking
for:

- manageability
- security
- efficient resource usage
- ...

Placing VMs ?



- Little flexibility for the application administrator
- Solutions are provider specific
- Current algorithms are not extensible by design

Challenges in designing a flexible consolidation manager

Issues :

- Numerous specific placement constraints
- Conflicting placement constraints
- Constraints expressed by non-expert users
- Scalability:
 - thousands of applications/VMs/Hosts

BtrPlace

- Configuration scripts:
 - Application manager
 - Datacenter administrator
- Extensible library of high-level placement constraints
- VM core model
 - Memory and CPU consumption
 - Migration, instantiation, shutdown costs
- Scalable and modular constraint solver
 - VM core model + script constraints

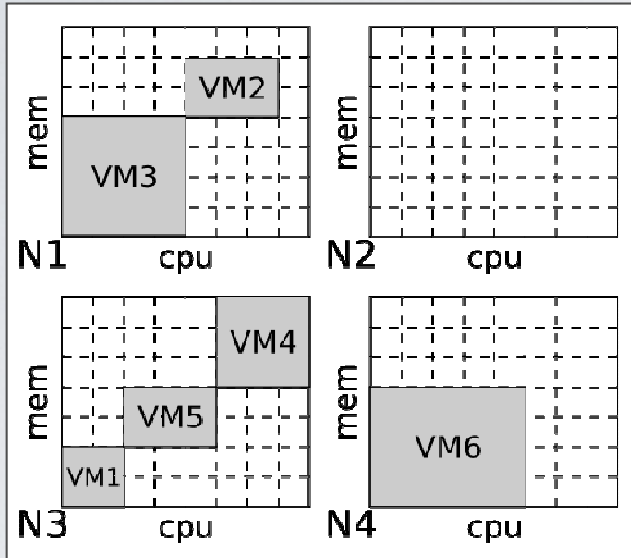
Configuration scripts

```
namespace datacenter;  
  
$servers = @N[1..12];  
$racks = { @N[1..4], @N[5..8], @N[9..12]};  
  
export $racks to *;
```

```
namespace sysadmin;  
import datacenter;  
import client.*;  
  
vmBtrplace: large;  
  
fence(vmBtrplace, @N1);  
lonely(vmBtrplace);  
ban($clients, @N5);
```

```
namespace clients.app1;  
import datacenter;  
  
VM[1..7]: small<clone, boot=5,halt=5>;  
VM[8..10]: large<clone, boot=60,halt=10>;  
$T1 = {VM1, VM2, VM3};  
$T2 = VM[4..7];  
$T3 = VM[8,10];  
  
for $t in $T[1..3] {  
    spread($t);  
}  
  
among($T3,$racks);  
export $me to sysadmin;
```

Data center administration




Btrplace

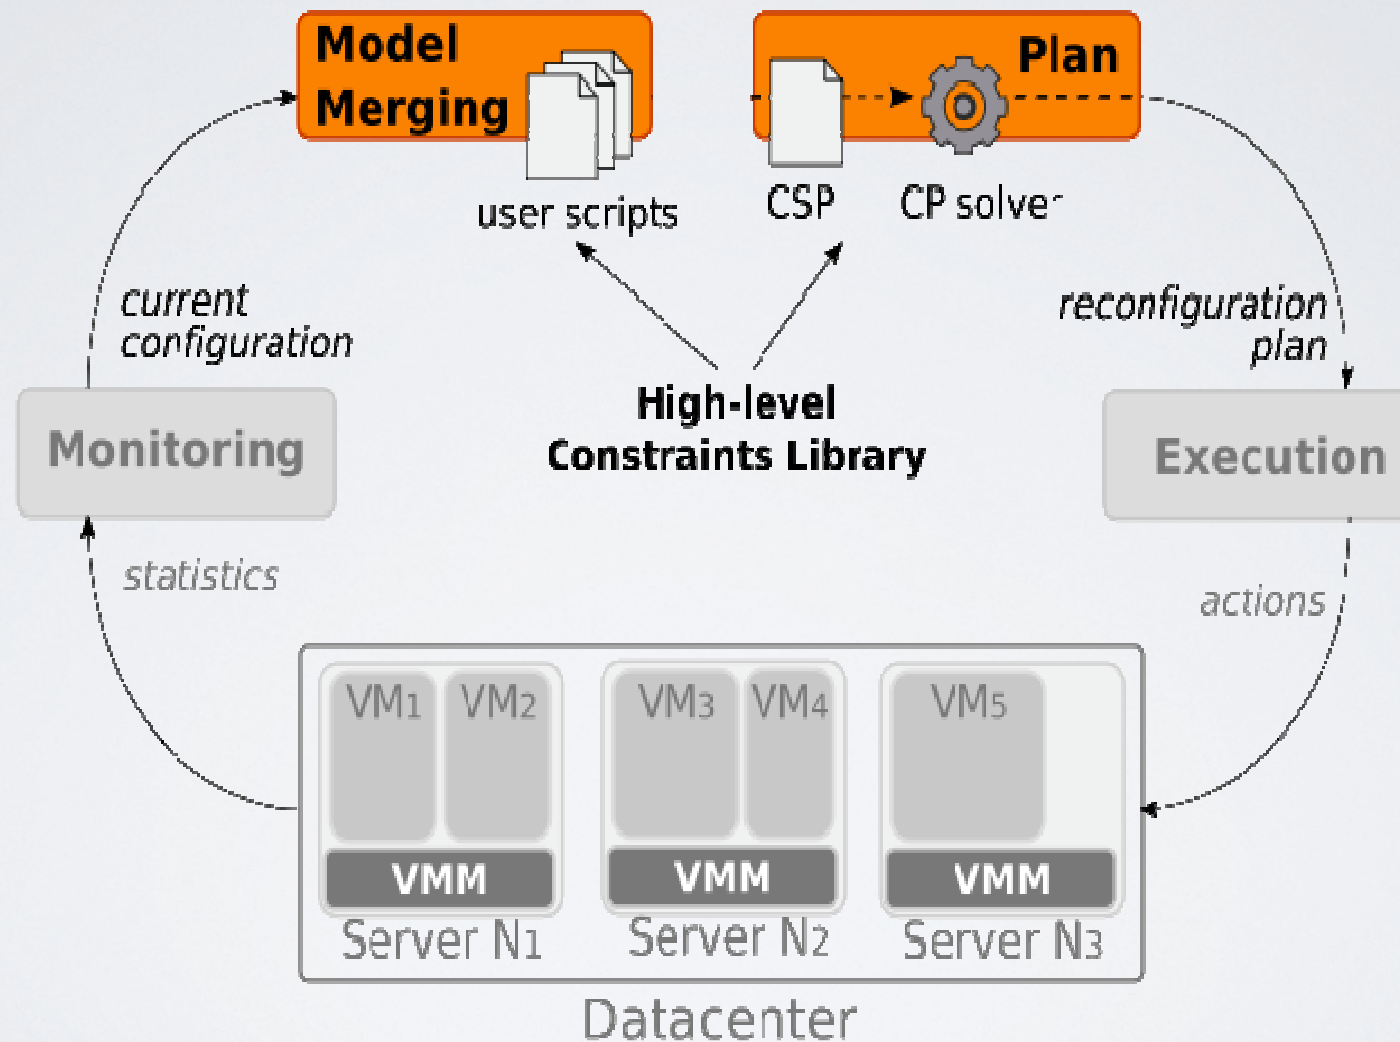
The reconfiguration plan :

```
0'00 to 0'02: relocate(VM2,N2)
0'00 to 0'04: relocate(VM6,N2)
0'02 to 0'05: relocate(VM4,N1)
0'04 to 0'08: shutdown(N4)
0'05 to 0'06: allocate(VM1,'cpu',3)
```

```
spread({VM3,VM2});
preserve({VM1},'ucpu', 3);
offline( @N4);
```



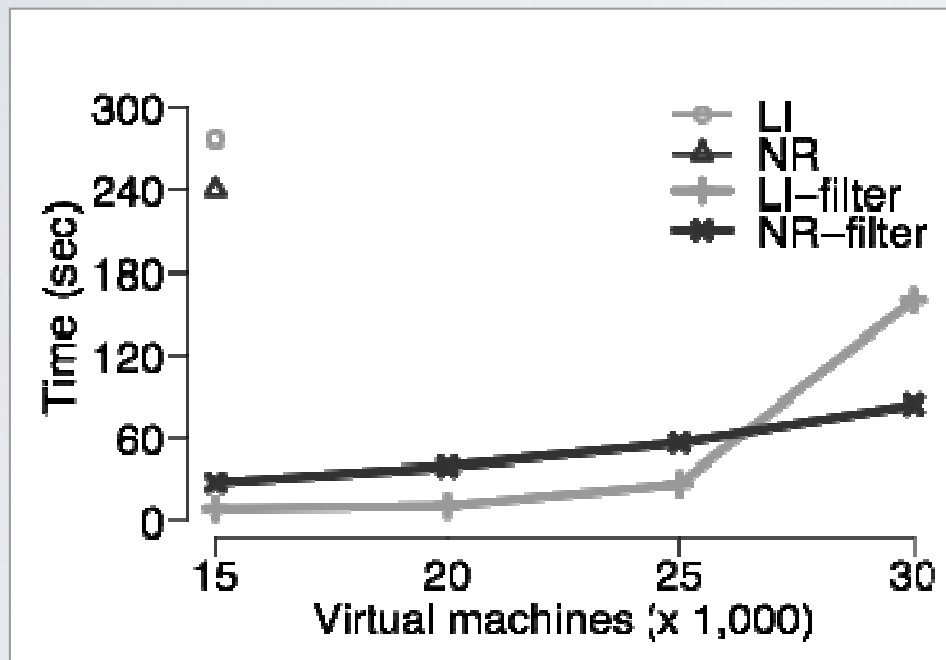
BtrPlace in practice



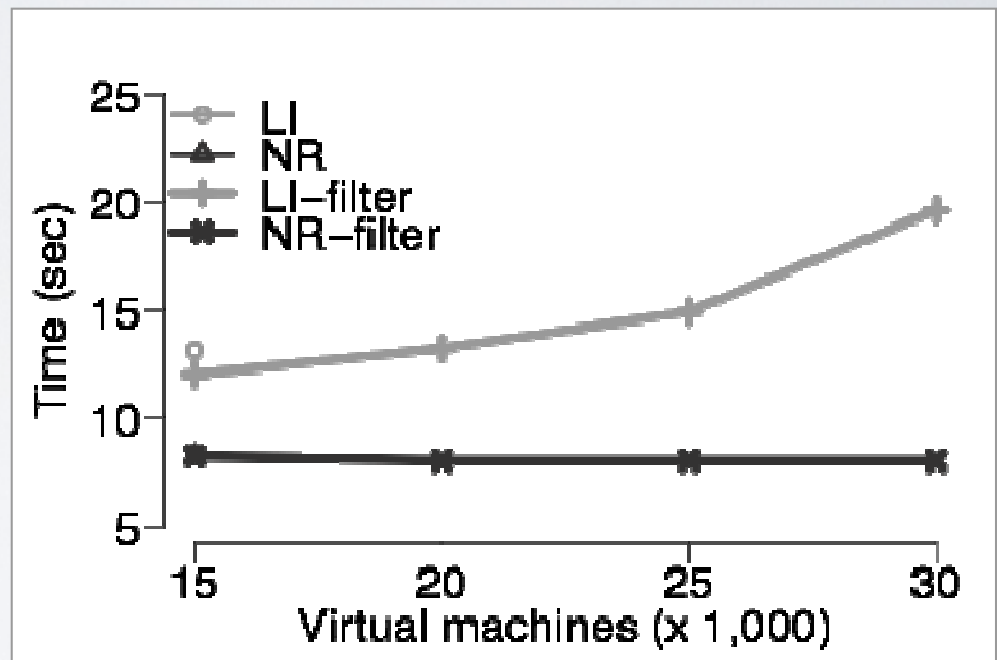
Scalability

- Simulated datacenter :
 - 5,000 servers
 - up to 1,700 3-tiers appliances (30,000 VMs)
 - a resource usage up to 73%
- 2 scenarios:
 - Load Increase (LI): 10% of the applications ask for 30% more uCPU
 - Network Rewiring (NR): 5% of the servers are turned off for a network maintenance

Performance evaluation

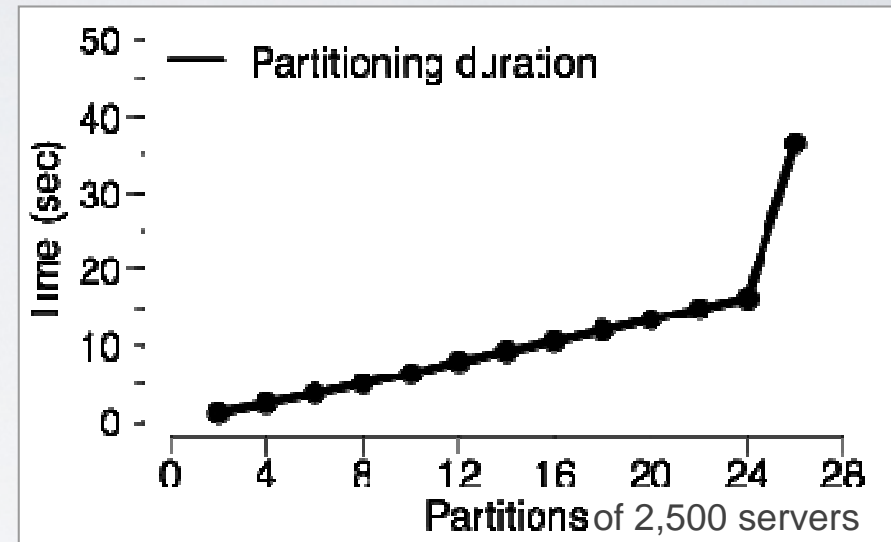
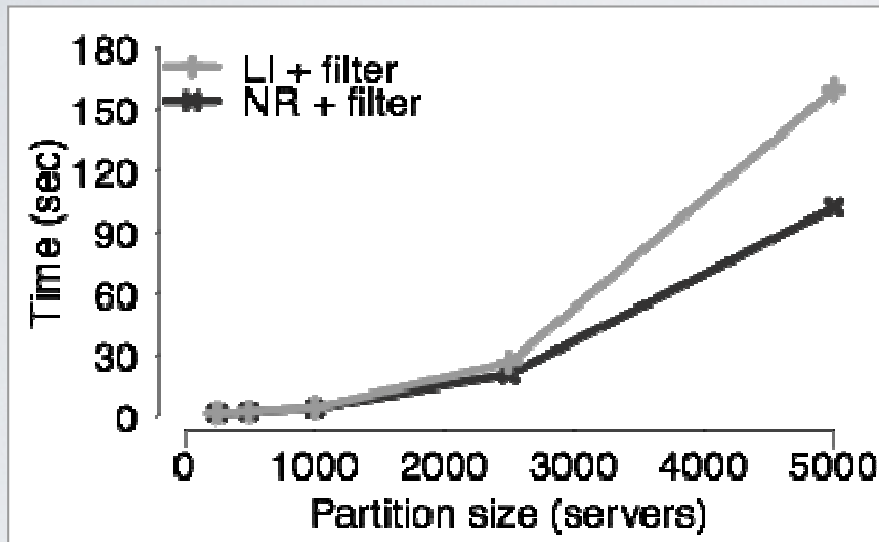


Solving duration



Reconfiguration duration

Partitioning



- the number of nodes to solve sub-RPs limits the scalability
- no impact on the quality of the reconfiguration plans
- too small partitions may alter the solvability

About BtrPlace

Online demo :

<http://btrp.inria.fr/sandbox>

Publications :

<http://sites.google.com/site/hermenierfabien/publications>