

From ZFS to BeeGFS

Oriol Mula-Valls

Institut Català de Ciències del Clima



Climate Forecasting Unit



Outline

- 1 Introduction
- 2 Old solution
- 3 New solution

Outline

1 Introduction

2 Old solution

3 New solution

Who and what

- 20 people (PostDocs, PhDs, management, IT, . . .)
- The Climate Forecasting Unit (CFU) undertakes research on the development and assessment of dynamical and statistical methods for the prediction of global and regional climate on time scales ranging from a few weeks to several years

Outline

1 Introduction

2 Old solution

3 New solution

Hardware

- 1 Supermicro server + 2 JBOD
 - $36 + 45 + 45 = 126$ x 3TB Hitachi Enterprise disks
- LSI 9280-8E RAID controllers
- 4x1Gbit Ethernet

Environment

- Set up
 - OpenIndiana and ZFS
 - RAID 0 on the controller for each disk
 - RAIDZ2 of 10 disks
 - NFSv4
- Clients
 - 20 desktop/workstations x 1Gbps
 - 2 FAT nodes x 1Gbps

Disaster recovery

- No backup, subset of data replicated
- 1 Supermicro server + 1 JBOD
 - $36 + 45 = 81$ x 3TB WD Caviar Green disks
- LSI 9280-8E RAID controllers
- 2x1Gbit Ethernet
- zrep

- Limited scalability
- Different ZFS issue (e.g. performance degradation 80%)
- Complex maintenance
 - RAID0 + RAIDZ2
 - Nested zfs + zrep

Outline

- 1 Introduction
- 2 Old solution
- 3 New solution**

Proposals

- 3 providers => 4 proposals
- HPCNow! selected
 - Presented two proposals
 - 1st and 2nd level BeeGFS support HPCNow!
 - 3rd level BeeGFS support ThinkParQ (1 year)
 - Equipment reutilization

Equipment

- 3 x Supermicro SuperStorage 6047R-E1R36N
 - 2 x 2 Intel Xeon Ivy Bridge E5-2609 V2
 - 64 GB RAM DDR3 1600 ECC
 - RAID Controller LSI 3108-H8IR
 - Intel X520 - DA2
- Reutilization of old equipment

Set up



