


Imperative Recovery

 Canonical Version Moved

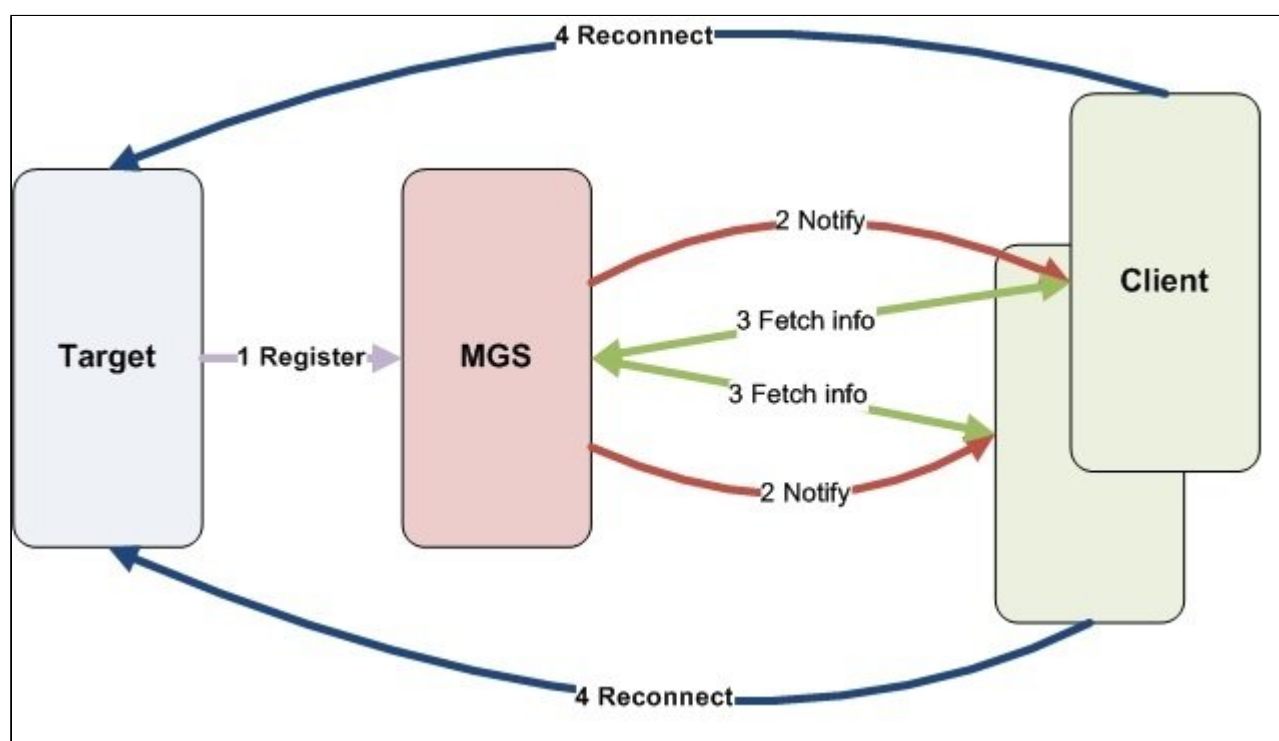
http://wiki.lustre.org/Imperative_Recovery

Imperative Recovery

1. Introduction

Imperative recovery reduces the time for clients to realize restarting of targets, so that they can do reconnection as soon as possible. The MGS is used as a reflector and it will notify clients when a target is newly registered.

This picture demonstrates the rough idea of Imperative Recovery:



2. HLD is here: [Imperative Recovery Design Document.pdf](#)

3. Testing

Scalability is very important for imperative recovery because definitely we need the capability to support at least 100K clients. Based on this situation, instead of regular unit and regression test, we also have a scalability test document to verify it works. Please check test plan here: [Imperative Recovery Test Plan Rev B.pdf](#)

4. Presentation

I did a presentation about imperative recovery on LUG2011, please check a revised PPT here: [IR-internal-1.pdf](#)

5. Results

In our simulation test on Hyperion, where it had 125 clients nodes and 600 mountpoints on each client, then we kill an OSS and investigated OST recovery time on it. It could finish recovery around 60 seconds; as a comparison, it took ~300 seconds without imperative recovery.

6. About the author

My name is Jinshan Xiong, I'm a senior software engineer at Whamcloud, Inc. Please contact with me at jinshan.xiong@whamcloud.com if you have any questions.