

Test Plan

Kerberos revival

version: 2

date: 2015/07/24

author: Sebastien Buisson – sebastien.buisson@atos.net

Introduction

Back in 2010, Lustre 2.0 was successfully kerberized at PSC. But in 2013, it was not even possible to just start Lustre 2.4 with Kerberos activated. And one year later, experimentations on Lustre 2.5 gave the same unsuccessful results: build with Kerberos was broken, and even after fixing this, an LBUG was hit right when starting Lustre with Kerberos activated.

The goal is to work on Lustre master branch to get Kerberos functional again. This work is tracked under Jira tickets LU-6020, LU-3778, and LU-6356.

Tests

Non regression tests

The full set of Auster test suites will be run to verify there are no functional regression.

Functional tests

As Kerberos provides a way to authenticate nodes and users that access Lustre file systems, we will ensure authentication is properly enforced:

- a server node with no or improper principal will not mount;
- a server node with proper principal will mount;
- a client node with no or improper principal will not mount;
- a client node with proper principal will mount;
- a user with no or improper principal will not be able to access files;
- a user with proper principal will be able to access files depending on the Unix rights.

The sanity-krb5 test suite will be run to validate that Kerberos support is functional again. But as Maloo does not have any proper Kerberos environment, sanity-krb5 will have to be launched on a dedicated, regression-test specific file system.

Performances

The performance impact of Kerberos will be evaluated, for all Kerberos flavors krb5n, krb5a, krb5i,

krb5p, in terms of bandwidth and metadata performance.

Upgrade and downgrade

Upgrade and downgrade of the Lustre client will be verified.

Upgrade and downgrade of the Lustre server will be verified.

Results

Non regression tests

Complete Auster test suite has been triggered by Maloo when patches from LU-6020, LU-3778, and LU-6356 were submitted. All tests passed.

Functional tests

Here are the results of the basic functional tests, ran with all Kerberos flavors krb5n, krb5a, krb5i, krb5p:

- a server node with no or improper principal does not mount: OK
- a server node with proper principal does mount: OK
- a client node with no or improper principal does not mount: OK
- a client node with proper principal does mount: OK
- a user with no or improper principal is not able to access files: OK
- a user with proper principal is able to access files depending on the Unix rights: OK

The sanity-krb5 test suite has been launched on an in-house test system where Kerberos environment is set up. The configuration is the following:

- software
 - RHEL 6
 - *Kerberos* MIT v5
 - Lustre 2.7.0 + patches
- hardware: 1 node per Lustre role, to 'ease' Kerberos setup
 - 1 MGS
 - 1 MDS
 - 1 OSS
 - 1 client

} ramdisk storage

 - 12 cores
 - 24 GB RAM
- Interconnect: Infiniband QDR

With all patches pushed under Jira tickets LU-6020, LU-3778, and LU-6356, all sanity-krb5 tests pass.

Performances

To evaluate Kerberos impact over performance, we set up the same test configuration as for sanity-krb5 run.

To sum up, the bandwidth performance tests carried out with IOR show that krb5n and krb5a flavors have no impact. With krb5i, write bandwidth performance drops by 50% and with krb5p performance drops by between 75% and 95%, depending on the encryption algorithm chosen: aes efficiency surpasses des3. Read performance impact is pretty similar, with a drop of 60% when using krb5i, and a drop of between 80% and 95% with krb5p.

The metadata performance tests have been carried out with mdtest. We show that even krb5n flavor has an impact: around 5%. Then we observe a drop in performance of about 20% with krb5a, 25% with krb5i, and finally between 40% and 60% with krb5p. Again when using krb5p flavor, the smallest performance hit is obtained with aes256 encryption algorithm. So aes256 is definitely the right choice, because it also provides the strongest security.

For complete details and explanations on the performance figures, please see the presentation given at LUG 2015 (performance study starts at slide 20):
http://cdn.opensfs.org/wp-content/uploads/2015/04/Lustre-and-Kerberos_Buisson.pdf

Upgrade and downgrade

Upgrade and downgrade tests have not been run yet.