DDN® STORAGE

Lustre Audit with Changelogs

October, 5th, 2017

Sebastien Buisson
sbuisson@ddn.com

What is Audit?

Information technology audits are used to evaluate the organization's ability to protect its information assets and to properly dispense information to authorized parties.

https://en.wikipedia.org/wiki/Information_technology_audit

Need for audit in Lustre

▶ **Support of rich security features:**

- authentication with Kerberos
- mandatory access control with SELinux
- isolation
- etc.

⇒ **Audit as a proof of security in place**

Need for audit in Lustre

- Lustre outside of traditional HPC field
- e.g. Life science
  - data privacy is crucial

⇒ Audit as a regulation compliance

Audit with SELinux

Pros	Cons
• integrated logging and auditing facility • proven	• on client side • need to consolidate

Audit with Changelogs

Pros	Cons
• integrated in Lustre • centralized • transactional	• lacks some info

Audit with Changelogs

▶ Lustre activity as seen by MDS

- file system namespace
- file metadata

▶ Store in Changelog records

- internal Lustre files

▶ Read from audit nodes

- dedicated clients
- move outside for later analysis

```
5 01CREAT 15:44:32.385864793 2017.07.18 0x0 t=[0x200000402:0x3:0x0]
          ef=0x1 p=[0x200000402:0x2:0x0] fileA
```

Lustre needs for proper audit

- **Identify subject of action**
  - uid/gid
  - NID
- **Record all actions**
  - open
  - close
  - xattr
  - denied accesses

Changelogs enhancements: LU-9727

▶ Mandatory base patch

- capability for additional extra fields in changelog entries

LU-9727 lustre: Add an additional set of 64 changelog flags

https://review.whamcloud.com/28045

Patch in collaboration with ANU

Changelogs enhancements: LU-9727

Subject identification

- add uid/gid

LU-9727 lustre: add uid/gid to Changelogs entries
https://review.whamcloud.com/28114

- add client NID

LU-9727 lustre: add client NID to Changelogs entries
https://review.whamcloud.com/28213

Changelogs enhancements: LU-9727

- **Record all actions**

  LU-9727 lustre: implement CL_OPEN for Changelogs
  https://review.whamcloud.com/28214

  LU-9727 lustre: record CLOSE if OPEN was recorded
  https://review.whamcloud.com/27929

  LU-9727 lustre: add CL_GETXATTR for Changelogs
  https://review.whamcloud.com/28251

  LU-9727 lustre: record denied OPEN in Changelogs
  https://review.whamcloud.com/288125

Changelogs enhancements: LU-9727

- **optimizations for audit**

  LU-9727 lustre: limit OPEN and CLOSE rates in Changelogs

  https://review.whamcloud.com/28299

  LU-9727 nodemap: add audit_mode flag to nodemap

  https://review.whamcloud.com/28313

  LU-9727 lustre: record if enable_audit is set on nodemap

  https://review.whamcloud.com/28314

Audit with Changelogs: impact study

Benchmark testbed

- MDS
  - SuperMicro SuperServer (2 x E5-2698 v4, 256GB memory, FDR)
  - SFA 7700X, Toshiba RI SSD
  - 2 x RAID1 with LVM for RAID10
- Client x 32
  - Intel server S2600KPR (E5-2650 v4 @ 2.20GHz, 24 cores, 128GB memory)

Audit with Changelogs: impact study

- Changelogs space consumption evaluation

	# changelog entries	changelog size
After 10 000 files created	30000	3755824
After 10 000 files read	50000	6096448
After 10 000 files removed	60000	7461440

- Rule of thumb: provision 125 B / entry on MDT

15

Audit with Changelogs: impact study

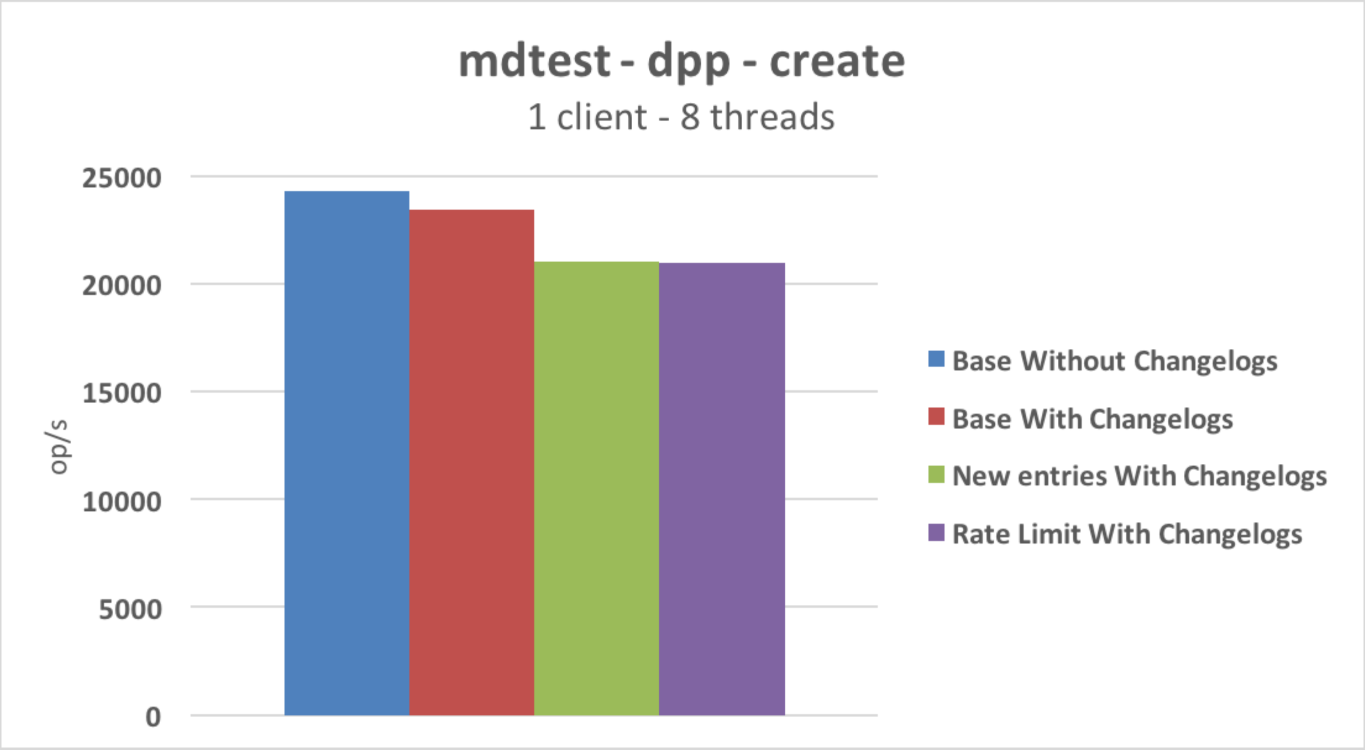

Audit with Changelogs: impact study

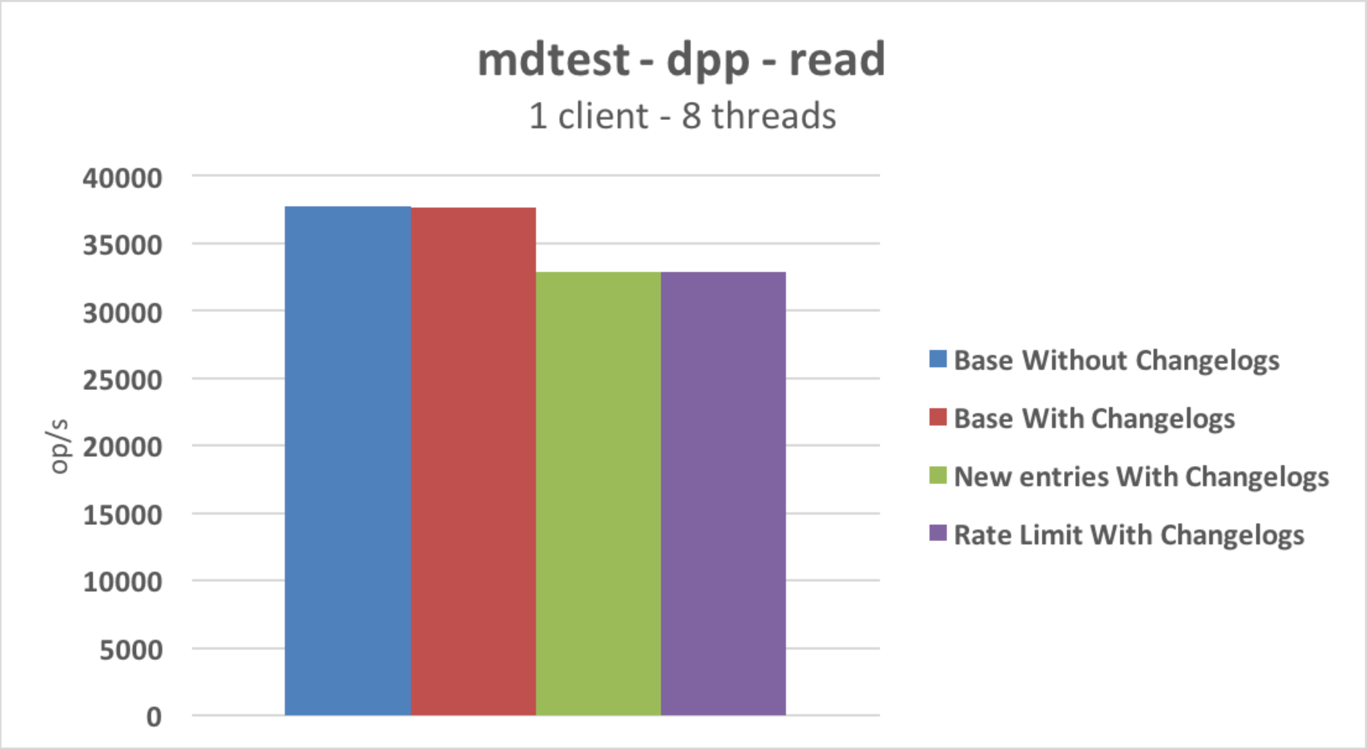

Audit with Changelogs: impact study

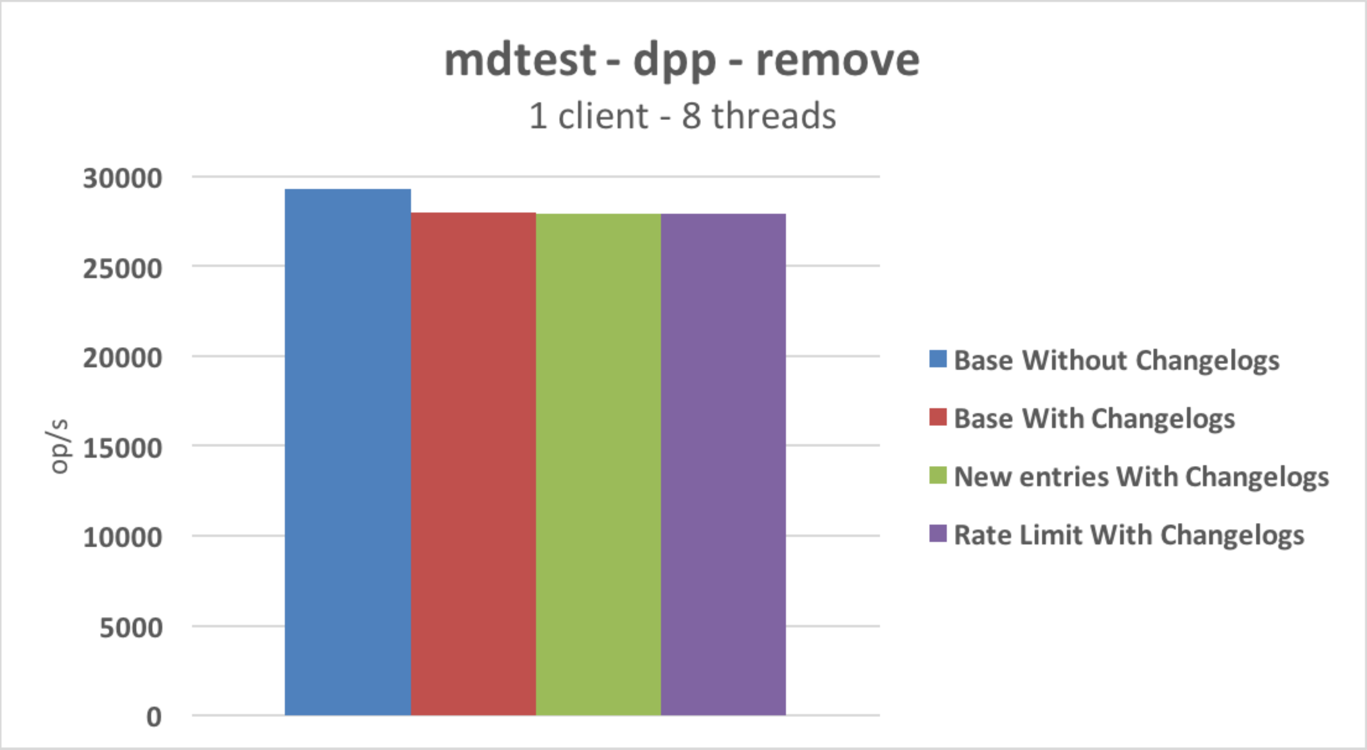

Audit with Changelogs: impact study

- **32 clients**
  - saturated MDT
  - with Audit: Changelogs file is bottleneck
- **For metadata intensive workload**
  - Changelogs on separate device
  - Changelogs in memory

Lustre Audit with Changelogs

- **Changelogs usable for audit**
- **All patches contributed to Community**
  - review in progress, help welcome ☺

- **next step**
  - Changelogs consumer specific for Audit

Thank You!

Keep in touch with us

sales@ddn.com

@ddn_limitless

company/datadirect-networks

9351 Deering Avenue,
Chatsworth, CA 91311

1.800.837.2298
1.818.700.4000

DDN® STORAGE