
JuiceFS 元数据性能分析

- 背景
- 系统调用下发到 Fuse 的请求
 - opendir()
 - readdir()
 - closedir()
 - stat()
 - fstat()
 - fstatat()
 - chmod()
 - open()
 - open(..., O_CREATE, ...)
 - read()
 - write()
 - close()
 - mkdir()
 - rmdir()
 - unlink()
- 缓存一致性测试
 - 测试场景 1: close-to-open 一致性测试 (读写轮询)
 - 测试场景 2: close-to-open 一致性测试 (全写)
- 缓存性能测试
 - 测试场景 1: 关闭 attr-cache, 关闭 entry-cache
 - 测试场景 2: 打开 attr-cache, 关闭 entry-cache
 - 测试场景 3: 关闭 attr-cache, 打开 entry-cache
 - 测试场景 4: 打开 attr-cache, 打开 entry-cache
- Fuse 对应 redis 请求
 - getattr
 - lookup
 - readdir
- 总结

背景

本文主要分析为什么 JuiceFS 在开启 close-to-open (默认) 配置下 ls 会比 CurveFS 性能好很多

系统调用下发到 Fuse 的请求

挂载点为 /mnt/jfs

opendir()

```
| DIR = opendir("/mnt/jfs/dir1/dir2")
```

```
2023.02.10 14:42:02.644483 [uid:0,gid:0,pid:15197] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.003496>
2023.02.10 14:42:02.648441 [uid:0,gid:0,pid:15197] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003751>
2023.02.10 14:42:02.654635 [uid:0,gid:0,pid:15197] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.005978>
2023.02.10 14:42:02.661527 [uid:0,gid:0,pid:15197] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.006693>
2023.02.10 14:42:02.665421 [uid:0,gid:0,pid:15197] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003739>
2023.02.10 14:42:02.665619 [uid:0,gid:0,pid:15197] opendir (707): OK [fh:5] <0.000020>
2023.02.10 14:42:02.670575 [uid:0,gid:0,pid:15197] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004794>
```

readdir()

readdir(DIR)

```
2023.02.10 14:42:59.645282 [uid:0,gid:0,pid:15197] readdir (707,4096,0): OK (4) <0.010503>
2023.02.10 14:43:09.697267 [uid:0,gid:0,pid:15197] readdir (707,4096,4): OK (0) <0.000005>
```

closedir()

| closedir(DIR)

```
2023.02.10 14:43:09.697817 [uid:0,gid:0,pid:0] releasedir (707): OK <0.000014>
```

stat()

| stat("/mnt/jfs/dir1/dir2/fl", ...)

```
2023.02.10 11:06:50.105309 [uid:0,gid:0,pid:14524] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.003318>
2023.02.10 11:06:50.109652 [uid:0,gid:0,pid:14524] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.004200>
2023.02.10 11:06:50.113628 [uid:0,gid:0,pid:14524] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003847>
2023.02.10 11:06:50.116382 [uid:0,gid:0,pid:14524] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002607>
2023.02.10 11:06:50.121037 [uid:0,gid:0,pid:14524] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004473>
2023.02.10 11:06:50.123726 [uid:0,gid:0,pid:14524] lookup (707,f1): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.002495>
2023.02.10 11:06:50.128384 [uid:0,gid:0,pid:14524] getattr (1201): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.004502>
```

fstat()

| fstat(fd, ...)

```
2023.02.10 11:09:18.153126 [uid:0,gid:0,pid:14527] getattr (1201): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.006537>
```

fstatat()

| fstatat(int dirfd, ...)

```
2023.02.10 15:37:27.749677 [uid:0,gid:0,pid:15374] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003039>
2023.02.10 15:37:27.753601 [uid:0,gid:0,pid:15374] lookup (707,f1): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.003745>
2023.02.10 15:37:27.761961 [uid:0,gid:0,pid:15374] getattr (1201): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.008165>
```

chmod()

| chmod("/mnt/jfs/dir1/dir2/f1", ...)

```
2023.02.14 20:05:38.017446 [uid:0,gid:0,pid:952965] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1676362192,1676362192,4096]) <0.000513>

2023.02.14 20:05:38.017861 [uid:0,gid:0,pid:952965] lookup (1,dir1): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000340>

2023.02.14 20:05:38.018352 [uid:0,gid:0,pid:952965] getattr (1305): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000341>

2023.02.14 20:05:38.018832 [uid:0,gid:0,pid:952965] lookup (1305,dir2): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375813,1676375813,4096]) <0.000325>

2023.02.14 20:05:38.019095 [uid:0,gid:0,pid:952965] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375813,1676375813,4096]) <0.000182>

2023.02.14 20:05:38.019481 [uid:0,gid:0,pid:952965] lookup (1306,f1): OK
(1328,[-rwxr--rwx:0100747,1,0,0,1676375813,1676375813,1676376296,0]) <0.000282>

2023.02.14 20:05:38.019968 [uid:0,gid:0,pid:952965] setattr (1328,0x1,[mode=-rwxr--rwx:00747]): OK
(1328,[-rwxr--rwx:0100747,1,0,0,1676375813,1676375813,1676376296,0]) <0.000414>
```

open()

| open("/mnt/jfs/dir1/dir2/f1", ...)

```
2023.02.10 11:08:23.228302 [uid:0,gid:0,pid:14527] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.004908>
2023.02.10 11:08:23.230793 [uid:0,gid:0,pid:14527] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002362>
2023.02.10 11:08:23.233989 [uid:0,gid:0,pid:14527] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002950>
2023.02.10 11:08:23.238151 [uid:0,gid:0,pid:14527] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003985>
2023.02.10 11:08:23.241274 [uid:0,gid:0,pid:14527] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002900>
2023.02.10 11:08:23.243799 [uid:0,gid:0,pid:14527] lookup (707,f1): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.002341>
2023.02.10 11:08:23.246869 [uid:0,gid:0,pid:14527] getattr (1201): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.002837>
2023.02.10 11:08:23.249586 [uid:0,gid:0,pid:14527] open (1201): OK [fh:3] <0.002553>
```

open(..., O_CREATE, ...)

| open("/mnt/jfs/dir1/dir2/f1", 0 | RDWR | O_CREATE, ...)

```
2023.02.14 19:56:53.527220 [uid:0,gid:0,pid:932390] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1676362192,1676362192,4096]) <0.000458>

2023.02.14 19:56:53.527694 [uid:0,gid:0,pid:932390] lookup (1,dir1): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000250>

2023.02.14 19:56:53.528094 [uid:0,gid:0,pid:932390] getattr (1305): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000175>

2023.02.14 19:56:53.528484 [uid:0,gid:0,pid:932390] lookup (1305,dir2): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375802,1676375802,4096]) <0.000252>

2023.02.14 19:56:53.528798 [uid:0,gid:0,pid:932390] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375802,1676375802,4096]) <0.000245>

2023.02.14 19:56:53.529018 [uid:0,gid:0,pid:932390] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375802,1676375802,4096]) <0.000119>

2023.02.14 19:56:53.529589 [uid:0,gid:0,pid:932390] lookup (1306,f006): no such file or directory <0.000475>

2023.02.14 19:56:53.530445 [uid:0,gid:0,pid:932390] create (1306,f006,-rw-r--r--:0100644): OK
(1328,[-rw-r--r--:0100644,1,0,0,1676375813,1676375813,1676375813,0]) [fh:421] <0.000743>
2023.02.14 19:56:53.530747 [uid:0,gid:0,pid:932390] getattr (1328): OK
(1328,[-rw-r--r--:0100644,1,0,0,1676375813,1676375813,1676375813,0]) <0.000190>
```

read()

| read(fd, ...)

```
2023.02.14 20:19:48.994594 [uid:0,gid:0,pid:987800] getattr (1332): OK
(1332,[-rw-r--r--:0100644,1,0,0,1676377160,1676377160,1676377160,0]) <0.000499>
2023.02.14 20:19:48.994759 [uid:0,gid:0,pid:987800] read (1332,4096,0): OK (0) <0.000011>
```

write()

| write(fd, ...)

```
2023.02.14 20:16:54.092888 [uid:0,gid:0,pid:980544] write (1331,11,0,432): OK <0.000127>
```

close()

| close(fd)

```
2023.02.10 11:09:18.153636 [uid:0,gid:0,pid:14527] flush (1201,3): OK <0.000008>
2023.02.10 11:09:18.153839 [uid:0,gid:0,pid:0] release (1201): OK <0.000027>
```

mkdir()

| mkdir("/mnt/jfs/dir1/dir2/d1")

```
2023.02.14 19:49:06.834595 [uid:0,gid:0,pid:913340] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1676362192,1676362192,4096]) <0.000542>

2023.02.14 19:49:06.835104 [uid:0,gid:0,pid:913340] lookup (1,dir1): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000293>

2023.02.14 19:49:06.835581 [uid:0,gid:0,pid:913340] getattr (1305): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000214>

2023.02.14 19:49:06.836043 [uid:0,gid:0,pid:913340] lookup (1305,dir2): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375114,1676375114,4096]) <0.000199>

2023.02.14 19:49:06.836563 [uid:0,gid:0,pid:913340] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375114,1676375114,4096]) <0.000341>

2023.02.14 19:49:06.836903 [uid:0,gid:0,pid:913340] lookup (1306,d1): no such file or directory <0.000268>

2023.02.14 19:49:06.837184 [uid:0,gid:0,pid:913340] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,3,0,0,1676362192,1676375114,1676375114,4096]) <0.000179>

2023.02.14 19:49:06.838684 [uid:0,gid:0,pid:913340] mkdir (1306,d1,?rwxr-xr-x:00755): OK
(1322,[drwxr-xr-x:0040755,2,0,0,1676375346,1676375346,1676375346,4096]) <0.001420>
```

rmdir()

```
| rmdir("/mnt/jfs/dir1/dir2/d1")
```

```
2023.02.14 19:51:31.522663 [uid:0,gid:0,pid:919279] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1676362192,1676362192,4096]) <0.000493>

2023.02.14 19:51:31.523237 [uid:0,gid:0,pid:919279] lookup (1,dir1): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000373>

2023.02.14 19:51:31.523841 [uid:0,gid:0,pid:919279] getattr (1305): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000379>

2023.02.14 19:51:31.524466 [uid:0,gid:0,pid:919279] lookup (1305,dir2): OK
(1306,[drwxr-xr-x:0040755,4,0,0,1676362192,1676375346,1676375346,4096]) <0.000379>

2023.02.14 19:51:31.524814 [uid:0,gid:0,pid:919279] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,4,0,0,1676362192,1676375346,1676375346,4096]) <0.000202>

2023.02.14 19:51:31.525332 [uid:0,gid:0,pid:919279] lookup (1306,d1): OK
(1322,[drwxr-xr-x:0040755,2,0,0,1676375346,1676375346,1676375346,4096]) <0.000347>

2023.02.14 19:51:31.525829 [uid:0,gid:0,pid:919279] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,4,0,0,1676362192,1676375346,1676375346,4096]) <0.000279>

2023.02.14 19:51:31.527952 [uid:0,gid:0,pid:919279] rmdir (1306,d1): OK <0.001943>
```

unlink()

```
| unlink("/mnt/jfs/dir1/dir2/fl")
```

```
2023.02.14 19:41:51.505990 [uid:0,gid:0,pid:895598] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1676362192,1676362192,4096]) <0.000533>

2023.02.14 19:41:51.506466 [uid:0,gid:0,pid:895598] lookup (1,dir1): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000286>

2023.02.14 19:41:51.506924 [uid:0,gid:0,pid:895598] getattr (1305): OK
(1305,[drwxr-xr-x:0040755,3,0,0,1676362192,1676362192,1676362192,4096]) <0.000214>

2023.02.14 19:41:51.507291 [uid:0,gid:0,pid:895598] lookup (1305,dir2): OK
(1306,[drwxr-xr-x:0040755,2,0,0,1676362192,1676374733,1676374733,4096]) <0.000180>

2023.02.14 19:41:51.507537 [uid:0,gid:0,pid:895598] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,2,0,0,1676362192,1676374733,1676374733,4096]) <0.000125>

2023.02.14 19:41:51.507963 [uid:0,gid:0,pid:895598] lookup (1306,f1): OK
(1310,[-rw-r--r--:0100644,1,0,0,1676366706,1676366706,1676366706,0]) <0.000324>

2023.02.14 19:41:51.508166 [uid:0,gid:0,pid:895598] getattr (1306): OK
(1306,[drwxr-xr-x:0040755,2,0,0,1676362192,1676374733,1676374733,4096]) <0.000131>

2023.02.14 19:41:51.508872 [uid:0,gid:0,pid:895598] unlink (1306,f1): OK <0.000607>
```

缓存一致性测试

A, B 2 台机器挂在相同的 JuiceFS 进行测试

测试场景 1: close-to-open 一致性测试 (读写轮询)

```
| juicefs mount redis://10.221.103.160:16379/1 /mnt/jfs --debug --attr-cache 3600 --entry-cache 3600 --dir-entry-cache 3600
```

```
A: echo "1" >> /mnt/jfs/dir1/dir2/f1
A: cat /mnt/jfs/dir1/dir2/f1
B: echo "2" >> /mnt/jfs/dir1/dir2/f1
B: cat /mnt/jfs/dir1/dir2/f1
```

```
A: echo "3" >> /mnt/jfs/dir1/dir2/fl
A: cat /mnt/jfs/dir1/dir2/fl
  A {1 2} {1,2,3} inode attritube

A: echo "1" >> /mnt/jfs/dir1/dir2/fl
2023.02.17 11:19:42.758515 [uid:0,gid:0,pid:26811] lookup (1,dir1): OK
(1346,[drwxr-xr-x:0040755,3,0,0,1676452174,1676452174,1676452174,4096]) <0.004064>
2023.02.17 11:19:42.762161 [uid:0,gid:0,pid:26811] lookup (1346,dir2): OK
(1347,[drwxr-xr-x:0040755,2,0,0,1676452174,1676603750,1676603750,4096]) <0.003482>
2023.02.17 11:19:42.768212 [uid:0,gid:0,pid:26811] lookup (1347,f1): no such file or directory <0.005889>
2023.02.17 11:19:42.799507 [uid:0,gid:0,pid:26811] create (1347,f1,-rw-r--r--:0100644): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676603982,1676603982,0]) [fh:2] <0.031082>
2023.02.17 11:19:42.799728 [uid:0,gid:0,pid:26811] flush (2201,2): OK <0.000010>
2023.02.17 11:19:42.799938 [uid:0,gid:0,pid:26811] write (2201,2,0,2): OK <0.000088> # write (ino, size, off,
fh)
2023.02.17 11:19:42.878080 [uid:0,gid:0,pid:26811] flush (2201,2): OK <0.078084>
2023.02.17 11:19:42.878239 [uid:0,gid:0,pid:0] release (2201): OK <0.000011>

A: cat /mnt/jfs/dir1/dir2/fl
2023.02.17 11:20:10.680637 [uid:0,gid:0,pid:26857] getattr (1347): OK
(1347,[drwxr-xr-x:0040755,2,0,0,1676452174,1676603982,1676603982,4096]) <0.003445>
2023.02.17 11:20:10.684178 [uid:0,gid:0,pid:26857] getattr (2201): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676603982,1676603982,2]) <0.003266>
2023.02.17 11:20:10.687633 [uid:0,gid:0,pid:26857] open (2201): OK [fh:3] <0.003292>
2023.02.17 11:20:10.692177 [uid:0,gid:0,pid:26857] read (2201,4096,0): OK (2) <0.004284>
2023.02.17 11:20:10.695852 [uid:0,gid:0,pid:26857] getattr (2201): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676603982,1676603982,2]) <0.003488>
2023.02.17 11:20:10.696022 [uid:0,gid:0,pid:26857] flush (2201,3): OK <0.000002>
2023.02.17 11:20:10.696159 [uid:0,gid:0,pid:0] release (2201): OK <0.000034>

B: echo "2" >> /mnt/jfs/dir1/dir2/fl
2023.02.17 03:20:37.028283 [uid:0,gid:0,pid:5640] lookup (1,dir1): OK
(1346,[drwxr-xr-x:0040755,3,0,0,1676452174,1676452174,1676452174,4096]) <0.003963>
2023.02.17 03:20:37.031748 [uid:0,gid:0,pid:5640] lookup (1346,dir2): OK
(1347,[drwxr-xr-x:0040755,2,0,0,1676452174,1676603982,1676603982,4096]) <0.003137>
2023.02.17 03:20:37.035460 [uid:0,gid:0,pid:5640] lookup (1347,f1): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676603982,1676603982,2]) <0.003487>
2023.02.17 03:20:37.043663 [uid:0,gid:0,pid:5640] open (2201): OK [fh:2] <0.007964>
```

```
2023.02.17 03:20:37.043894 [uid:0,gid:0,pid:5640] flush (2201,2): OK <0.000040>
2023.02.17 03:20:37.044298 [uid:0,gid:0,pid:5640] write (2201,2,2,2): OK <0.000177>
2023.02.17 03:20:37.106039 [uid:0,gid:0,pid:5640] flush (2201,2): OK <0.061658>
2023.02.17 03:20:37.106206 [uid:0,gid:0,pid:0] release (2201): OK <0.000024>
```

```
B: cat /mnt/jfs/dir1/dir2/f1
```

```
2023.02.17 03:21:02.835069 [uid:0,gid:0,pid:5892] getattr (2201): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676604037,1676604037,4]) <0.003456>
2023.02.17 03:21:02.839851 [uid:0,gid:0,pid:5892] open (2201): OK [fh:3] <0.004576>
2023.02.17 03:21:02.862047 [uid:0,gid:0,pid:5892] read (2201,4096,0): OK (4) <0.021954>
2023.02.17 03:21:02.865591 [uid:0,gid:0,pid:5892] getattr (2201): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676604037,1676604037,4]) <0.003307>
2023.02.17 03:21:02.865833 [uid:0,gid:0,pid:5892] flush (2201,3): OK <0.000003>
2023.02.17 03:21:02.865986 [uid:0,gid:0,pid:0] release (2201): OK <0.000009>
```

```
A: echo "3" >> /mnt/jfs/dir1/dir2/f1
```

```
2023.02.17 11:21:29.973904 [uid:0,gid:0,pid:26811] open (2201): OK [fh:4] <0.003867>
2023.02.17 11:21:29.974171 [uid:0,gid:0,pid:26811] flush (2201,4): OK <0.000022>
2023.02.17 11:21:29.974629 [uid:0,gid:0,pid:26811] write (2201,2,2,4): OK <0.000307>
2023.02.17 11:21:30.015865 [uid:0,gid:0,pid:26811] flush (2201,4): OK <0.041041>
2023.02.17 11:21:30.016059 [uid:0,gid:0,pid:0] release (2201): OK <0.000021>
```

```
A: cat /mnt/jfs/dir1/dir2/f1
```

```
2023.02.17 11:21:53.234619 [uid:0,gid:0,pid:26858] getattr (2201): OK
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676604090,1676604090,4]) <0.012815>
2023.02.17 11:21:53.238300 [uid:0,gid:0,pid:26858] open (2201): OK [fh:5] <0.003462>
2023.02.17 11:21:53.241335 [uid:0,gid:0,pid:26858] read (2201,4096,0): OK (4) <0.002824>
2023.02.17 11:21:53.246976 [uid:0,gid:0,pid:26858] getattr (2201): OK
```

```
(2201,[-rw-r--r--:0100644,1,0,0,1676603982,1676604090,1676604090,4]) <0.005422>
2023.02.17 11:21:53.247131 [uid:0,gid:0,pid:26858] flush (2201,5): OK <0.000002>
2023.02.17 11:21:53.247236 [uid:0,gid:0,pid:0] release (2201): OK <0.000022>
```

测试场景 2: close-to-open 一致性测试 (全写)

```
| juicefs mount redis://10.221.103.160:16379/1 /mnt/jfs --debug --attr-cache 3600 --entry-cache 3600 --dir-entry-cache 3600
```

```
A: echo "1" >> /mnt/jfs/dir1/dir2/f1
B: echo "2" >> /mnt/jfs/dir1/dir2/f1
A: echo "3" >> /mnt/jfs/dir1/dir2/f1
A: echo "4" >> /mnt/jfs/dir1/dir2/f1
A: cat /mnt/jfs/dir1/dir2/f1
B: cat /mnt/jfs/dir1/dir2/f1
A B {1,2,3,4}
```

```
A: echo "1" >> /mnt/jfs/dir1/dir2/f1
2023.02.17 11:35:11.717416 [uid:0,gid:0,pid:26811] lookup (1,dir1): OK
(1346,[drwxr-xr-x:0040755,3,0,0,1676452174,1676452174,1676452174,4096]) <0.003514>
2023.02.17 11:35:11.724886 [uid:0,gid:0,pid:26811] lookup (1346,dir2): OK
(1347,[drwxr-xr-x:0040755,2,0,0,1676452174,1676604804,1676604804,4096]) <0.007259>
2023.02.17 11:35:11.729444 [uid:0,gid:0,pid:26811] lookup (1347,f1): no such file or directory <0.004447>
2023.02.17 11:35:11.765362 [uid:0,gid:0,pid:26811] create (1347,f1,-rw-r--r--:0100644): OK
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604911,1676604911,0]) [fh:2] <0.035714>
2023.02.17 11:35:11.765517 [uid:0,gid:0,pid:26811] flush (2301,2): OK <0.000007>
2023.02.17 11:35:11.765915 [uid:0,gid:0,pid:26811] write (2301,2,0,2): OK <0.000178>
2023.02.17 11:35:11.846431 [uid:0,gid:0,pid:26811] flush (2301,2): OK <0.080389>
2023.02.17 11:35:11.846530 [uid:0,gid:0,pid:0] release (2301): OK <0.000011>
```

```
B: echo "2" >> /mnt/jfs/dir1/dir2/f1
2023.02.17 03:35:46.343648 [uid:0,gid:0,pid:5640] lookup (1,dir1): OK
(1346,[drwxr-xr-x:0040755,3,0,0,1676452174,1676452174,1676452174,4096]) <0.002804>
2023.02.17 03:35:46.348171 [uid:0,gid:0,pid:5640] lookup (1346,dir2): OK
(1347,[drwxr-xr-x:0040755,2,0,0,1676452174,1676604911,1676604911,4096]) <0.004296>
2023.02.17 03:35:46.351598 [uid:0,gid:0,pid:5640] lookup (1347,f1): OK
```

```
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604911,1676604911,2]) <0.003269>
2023.02.17 03:35:46.358470 [uid:0,gid:0,pid:5640] open (2301): OK [fh:2] <0.006614>
2023.02.17 03:35:46.358733 [uid:0,gid:0,pid:5640] flush (2301,2): OK <0.000031>
2023.02.17 03:35:46.359146 [uid:0,gid:0,pid:5640] write (2301,2,2,2): OK <0.000122>
2023.02.17 03:35:46.423230 [uid:0,gid:0,pid:5640] flush (2301,2): OK <0.064004>
2023.02.17 03:35:46.423450 [uid:0,gid:0,pid:0] release (2301): OK <0.000022>
```

```
A: echo "3" >> /mnt/jfs/dir1/dir2/fl1
```

```
2023.02.17 11:36:10.811525 [uid:0,gid:0,pid:26811] getattr (1347): OK
(1347,[drwxr-xr-x:0040755,2,0,0,1676452174,1676604911,1676604911,4096]) <0.003361>
2023.02.17 11:36:10.818069 [uid:0,gid:0,pid:26811] getattr (2301): OK
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604946,1676604946,4]) <0.006342>
2023.02.17 11:36:10.822910 [uid:0,gid:0,pid:26811] open (2301): OK [fh:3] <0.004721>
2023.02.17 11:36:10.823078 [uid:0,gid:0,pid:26811] flush (2301,3): OK <0.000009>
2023.02.17 11:36:10.823308 [uid:0,gid:0,pid:26811] write (2301,2,4,3): OK <0.000084>
2023.02.17 11:36:10.851441 [uid:0,gid:0,pid:26811] flush (2301,3): OK <0.028029>
2023.02.17 11:36:10.851661 [uid:0,gid:0,pid:0] release (2301): OK <0.000029>
```

```
A: cat /mnt/jfs/dir1/dir2/fl1
```

```
2023.02.17 11:40:52.759010 [uid:0,gid:0,pid:26910] getattr (2301): OK
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604970,1676604970,6]) <0.003771>
2023.02.17 11:40:52.764608 [uid:0,gid:0,pid:26910] open (2301): OK [fh:4] <0.005383>
2023.02.17 11:40:52.787226 [uid:0,gid:0,pid:26910] read (2301,4096,0): OK (6) <0.022263>
2023.02.17 11:40:52.791313 [uid:0,gid:0,pid:26910] getattr (2301): OK
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604970,1676604970,6]) <0.003495>
2023.02.17 11:40:52.791530 [uid:0,gid:0,pid:26910] flush (2301,4): OK <0.000003>
2023.02.17 11:40:52.791637 [uid:0,gid:0,pid:0] release (2301): OK <0.000009>
```

```
B: cat /mnt/jfs/dir1/dir2/fl1
```

```
2023.02.17 03:41:23.859603 [uid:0,gid:0,pid:5968] getattr (2301): OK
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604970,1676604970,6]) <0.003565>
2023.02.17 03:41:23.862869 [uid:0,gid:0,pid:5968] open (2301): OK [fh:3] <0.003035>
2023.02.17 03:41:23.886538 [uid:0,gid:0,pid:5968] read (2301,4096,0): OK (6) <0.023365>
2023.02.17 03:41:23.891627 [uid:0,gid:0,pid:5968] getattr (2301): OK
```

```
(2301,[-rw-r--r--:0100644,1,0,0,1676604911,1676604970,1676604970,6]) <0.004940>
2023.02.17 03:41:23.891928 [uid:0,gid:0,pid:5968] flush (2301,3): OK <0.000039>
2023.02.17 03:41:23.892077 [uid:0,gid:0,pid:0] release (2301): OK <0.000024>
```

缓存性能测试

```
| ls -l /mnt/jfs/dir1/dir2
```

```
total 13
drwxr-xr-x 3 root root 4096 Feb 10 10:54 .
drwxr-xr-x 3 root root 4096 Feb  9 17:15 ..
drwxr-xr-x 2 root root 4096 Feb  9 17:46 d1
-rw-r--r-- 1 root root    2 Feb 10 10:55 f1
```

测试场景 1: 关闭 attr-cache, 关闭 entry-cache

```
| juicefs mount redis://10.221.103.160:16379/1 /mnt/jfs --debug --attr-cache 0 --entry-cache 0 --dir-entry-cache 0
```

```
real    0m0.158s
user    0m0.004s
sys     0m0.002s
```

```
lookup*12
releasedir*1
getattr*20
readdir*2
opendir*1
```

```
2023.02.10 16:58:58.001690 [uid:0,gid:0,pid:16062] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.008249>
```

```
2023.02.10 16:58:58.005151 [uid:0,gid:0,pid:16062] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003262>
```

```
2023.02.10 16:58:58.017125 [uid:0,gid:0,pid:16062] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.011738>

2023.02.10 16:58:58.021199 [uid:0,gid:0,pid:16062] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003946>

2023.02.10 16:58:58.023686 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002331>

2023.02.10 16:58:58.023870 [uid:0,gid:0,pid:16062] opendir (707): OK [fh:2] <0.000008>

2023.02.10 16:58:58.028552 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004565>

2023.02.10 16:58:58.037266 [uid:0,gid:0,pid:16062] readdir (707,4096,0): OK (4) <0.008439>
2023.02.10 16:58:58.040094 [uid:0,gid:0,pid:16062] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.002572>
2023.02.10 16:58:58.043715 [uid:0,gid:0,pid:16062] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003396>
2023.02.10 16:58:58.050074 [uid:0,gid:0,pid:16062] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.006138>
2023.02.10 16:58:58.054131 [uid:0,gid:0,pid:16062] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003854>
2023.02.10 16:58:58.057766 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003427>
2023.02.10 16:58:58.060449 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002479>
2023.02.10 16:58:58.063969 [uid:0,gid:0,pid:16062] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.003181>
2023.02.10 16:58:58.066242 [uid:0,gid:0,pid:16062] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002051>
2023.02.10 16:58:58.069967 [uid:0,gid:0,pid:16062] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003475>
2023.02.10 16:58:58.073792 [uid:0,gid:0,pid:16062] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003565>
2023.02.10 16:58:58.076382 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002344>
2023.02.10 16:58:58.079553 [uid:0,gid:0,pid:16062] getattr (706): OK
```

```
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002912>
2023.02.10 16:58:58.082582 [uid:0,gid:0,pid:16062] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.002763>
2023.02.10 16:58:58.086005 [uid:0,gid:0,pid:16062] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003175>
2023.02.10 16:58:58.088493 [uid:0,gid:0,pid:16062] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002340>
2023.02.10 16:58:58.096050 [uid:0,gid:0,pid:16062] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.007365>
2023.02.10 16:58:58.102423 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.006178>
2023.02.10 16:58:58.106639 [uid:0,gid:0,pid:16062] lookup (707,d1): OK
(901,[drwxr-xr-x:0040755,2,0,0,1675936014,1675936014,1675936014,4096]) <0.004014>
2023.02.10 16:58:58.110358 [uid:0,gid:0,pid:16062] getattr (901): OK
(901,[drwxr-xr-x:0040755,2,0,0,1675936014,1675936014,1675936014,4096]) <0.003532>
2023.02.10 16:58:58.117477 [uid:0,gid:0,pid:16062] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.006859>
2023.02.10 16:58:58.121839 [uid:0,gid:0,pid:16062] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.004223>
2023.02.10 16:58:58.127693 [uid:0,gid:0,pid:16062] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.005651>
2023.02.10 16:58:58.161753 [uid:0,gid:0,pid:16062] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.033873>
2023.02.10 16:58:58.168013 [uid:0,gid:0,pid:16062] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.006077>
2023.02.10 16:58:58.170649 [uid:0,gid:0,pid:16062] lookup (707,f1): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.002382>
2023.02.10 16:58:58.172918 [uid:0,gid:0,pid:16062] getattr (1201): OK
```

```
(1201, [-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.002107>  
2023.02.10 16:58:58.173153 [uid:0,gid:0,pid:16062] readdir (707,4096,4): OK (0) <0.000003>  
2023.02.10 16:58:58.173576 [uid:0,gid:0,pid:0] releasedir (707): OK <0.000003>
```

测试场景 2: 打开 attr-cache, 关闭 entry-cache

```
| juicefs mount redis://10.221.103.160:16379/1 /mnt/jfs --debug --attr-cache 10 --entry-cache 0 --dir-entry-cache 0
```

```
real    0m0.065s
user    0m0.000s
sys     0m0.005s
```

```
lookup*12
releasedir*1
readdir*2
opendir*1
```

```
2023.02.10 16:13:26.780662 [uid:0,gid:0,pid:15831] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003177>
2023.02.10 16:13:26.784088 [uid:0,gid:0,pid:15831] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003253>
2023.02.10 16:13:26.784325 [uid:0,gid:0,pid:15831] opendir (707): OK [fh:2] <0.000006>
2023.02.10 16:13:26.794165 [uid:0,gid:0,pid:15831] readdir (707,4096,0): OK (4) <0.009573>
2023.02.10 16:13:26.797627 [uid:0,gid:0,pid:15831] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.003147>
2023.02.10 16:13:26.799943 [uid:0,gid:0,pid:15831] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002166>
2023.02.10 16:13:26.802100 [uid:0,gid:0,pid:15831] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.001883>
2023.02.10 16:13:26.806573 [uid:0,gid:0,pid:15831] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004246>
2023.02.10 16:13:26.809690 [uid:0,gid:0,pid:15831] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002903>
2023.02.10 16:13:26.813407 [uid:0,gid:0,pid:15831] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003491>
2023.02.10 16:13:26.819005 [uid:0,gid:0,pid:15831] lookup (707,d1): OK
(901,[drwxr-xr-x:0040755,2,0,0,1675936014,1675936014,1675936014,4096]) <0.005355>
2023.02.10 16:13:26.821380 [uid:0,gid:0,pid:15831] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002176>
2023.02.10 16:13:26.824193 [uid:0,gid:0,pid:15831] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002667>
2023.02.10 16:13:26.827883 [uid:0,gid:0,pid:15831] lookup (707,f1): OK
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.003514>
2023.02.10 16:13:26.828105 [uid:0,gid:0,pid:15831] readdir (707,4096,4): OK (0) <0.000003>
2023.02.10 16:13:26.828826 [uid:0,gid:0,pid:0] releasedir (707): OK <0.000004>
```

测试场景 3: 关闭 attr-cache, 打开 entry-cache

```
| juicefs mount redis://10.221.103.160:16379/1 /mnt/jfs --debug --attr-cache 0 --entry-cache 1 --dir-entry-cache 1
```

```
real    0m0.108s
user    0m0.000s
sys     0m0.005s
```

```
lookup*2
releasedir*1
getattr*20
readdir*2
opendir*1
```

```
2023.02.10 16:15:34.912702 [uid:0,gid:0,pid:15854] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.002386>
2023.02.10 16:15:34.918146 [uid:0,gid:0,pid:15854] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.005244>
2023.02.10 16:15:34.923999 [uid:0,gid:0,pid:15854] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.005693>
2023.02.10 16:15:34.927191 [uid:0,gid:0,pid:15854] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003011>
2023.02.10 16:15:34.931586 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004193>
2023.02.10 16:15:34.931780 [uid:0,gid:0,pid:15854] opendir (707): OK [fh:2] <0.000005>
2023.02.10 16:15:34.940551 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.008622>
2023.02.10 16:15:34.951830 [uid:0,gid:0,pid:15854] readdir (707,4096,0): OK (4) <0.011008>
2023.02.10 16:15:34.956179 [uid:0,gid:0,pid:15854] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.004114>
2023.02.10 16:15:34.958583 [uid:0,gid:0,pid:15854] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002347>
2023.02.10 16:15:34.963684 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.005025>
2023.02.10 16:15:34.968136 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004293>
2023.02.10 16:15:34.972328 [uid:0,gid:0,pid:15854] getattr (1): OK
```

```
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.003948>
2023.02.10 16:15:34.978388 [uid:0,gid:0,pid:15854] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.005872>
2023.02.10 16:15:34.993122 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.014554>
2023.02.10 16:15:35.002036 [uid:0,gid:0,pid:15854] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.008668>
2023.02.10 16:15:35.006060 [uid:0,gid:0,pid:15854] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.003722>
2023.02.10 16:15:35.008842 [uid:0,gid:0,pid:15854] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.002569>
2023.02.10 16:15:35.013101 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004056>
2023.02.10 16:15:35.020404 [uid:0,gid:0,pid:15854] getattr (901): OK
(901,[drwxr-xr-x:0040755,2,0,0,1675936014,1675936014,1675936014,4096]) <0.007083>
2023.02.10 16:15:35.027236 [uid:0,gid:0,pid:15854] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.006541>
2023.02.10 16:15:35.031766 [uid:0,gid:0,pid:15854] getattr (706): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.004358>
2023.02.10 16:15:35.036890 [uid:0,gid:0,pid:15854] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.004897>
2023.02.10 16:15:35.040737 [uid:0,gid:0,pid:15854] getattr (1201): OK
```

```
(1201,[-rw-r--r--:0100644,1,0,0,1675997697,1675997750,1675997750,2]) <0.003630>
2023.02.10 16:15:35.040960 [uid:0,gid:0,pid:15854] readdir (707,4096,4): OK (0) <0.000003>
2023.02.10 16:15:35.041316 [uid:0,gid:0,pid:0] releasedir (707): OK <0.000004>
```

测试场景 4: 打开 attr-cache, 打开 entry-cache

```
| juicefs mount redis://10.221.103.160:16379/1 /mnt/jfs --debug --attr-cache 1 --entry-cache 1 --dir-entry-cache 1
```

```
real    0m0.030s
user    0m0.001s
sys     0m0.005s
```

```
lookup*2
releasedir*1
getattr*2
readdir*2
opendir*1
```

```
2023.02.10 16:33:33.496629 [uid:0,gid:0,pid:15979] getattr (1): OK
(1,[drwxrwxrwx:0040777,3,0,0,1675827287,1675934138,1675934138,4096]) <0.002276>
2023.02.10 16:33:33.501514 [uid:0,gid:0,pid:15979] lookup (1,dir1): OK
(706,[drwxr-xr-x:0040755,3,0,0,1675934138,1675934138,1675934138,4096]) <0.004714>
2023.02.10 16:33:33.504255 [uid:0,gid:0,pid:15979] lookup (706,dir2): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.002559>
2023.02.10 16:33:33.504474 [uid:0,gid:0,pid:15979] opendir (707): OK [fh:2] <0.000006>
2023.02.10 16:33:33.514882 [uid:0,gid:0,pid:15979] readdir (707,4096,0): OK (4) <0.010097>
2023.02.10 16:33:33.519149 [uid:0,gid:0,pid:15979] getattr (707): OK
(707,[drwxr-xr-x:0040755,3,0,0,1675934138,1675997697,1675997697,4096]) <0.003968>
2023.02.10 16:33:33.519485 [uid:0,gid:0,pid:15979] readdir (707,4096,4): OK (0) <0.000003>
2023.02.10 16:33:33.519855 [uid:0,gid:0,pid:0] releasedir (707): OK <0.000004>
```

Fuse 对应 redis 请求

```
|
```

```
1676019538.011883 [1 10.246.159.29:60296] "get" "i1"
1676019538.017947 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d1" "dir1"
1676019538.018012 [1 lua] "HGET" "d1" "dir1"
1676019538.018048 [1 lua] "GET" "i706"
1676019538.029964 [1 10.246.159.29:60296] "get" "i706"
1676019538.034265 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d706" "dir2"
1676019538.034313 [1 lua] "HGET" "d706" "dir2"
1676019538.034335 [1 lua] "GET" "i707"
1676019538.036745 [1 10.246.159.29:60296] "get" "i707"
1676019538.040165 [1 10.246.159.29:60296] "get" "i707"
1676019538.045220 [1 10.246.159.29:60296] "get" "i707"
1676019538.047807 [1 10.246.159.29:60296] "hscan" "d707" "0" "match" "*" "count" "10000"
1676019538.050223 [1 10.246.159.29:60296] "mget" "i707" "i706" "i901" "i1201"
1676019538.053286 [1 10.246.159.29:60296] "get" "i1"
1676019538.056267 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d1" "dir1"
1676019538.056313 [1 lua] "HGET" "d1" "dir1"
1676019538.056339 [1 lua] "GET" "i706"
1676019538.059872 [1 10.246.159.29:60296] "get" "i706"
```

1676019538.065449 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d706" "dir2"
1676019538.065489 [1 lua] "HGET" "d706" "dir2"
1676019538.065512 [1 lua] "GET" "i707"
1676019538.069867 [1 10.246.159.29:60296] "get" "i707"
1676019538.073482 [1 10.246.159.29:60296] "get" "i707"
1676019538.076007 [1 10.246.159.29:60296] "get" "i1"
1676019538.079373 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d1" "dir1"
1676019538.079413 [1 lua] "HGET" "d1" "dir1"
1676019538.079434 [1 lua] "GET" "i706"
1676019538.082987 [1 10.246.159.29:60296] "get" "i706"
1676019538.086160 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d706" "dir2"
1676019538.086197 [1 lua] "HGET" "d706" "dir2"
1676019538.086217 [1 lua] "GET" "i707"
1676019538.089508 [1 10.246.159.29:60296] "get" "i707"
1676019538.092526 [1 10.246.159.29:60296] "get" "i706"
1676019538.095444 [1 10.246.159.29:60296] "get" "i1"
1676019538.098569 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d1" "dir1"
1676019538.098614 [1 lua] "HGET" "d1" "dir1"
1676019538.098639 [1 lua] "GET" "i706"
1676019538.101419 [1 10.246.159.29:60296] "get" "i706"
1676019538.104749 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d706" "dir2"
1676019538.104778 [1 lua] "HGET" "d706" "dir2"
1676019538.104794 [1 lua] "GET" "i707"
1676019538.115599 [1 10.246.159.29:60296] "get" "i707"
1676019538.119629 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d707" "d1"
1676019538.119666 [1 lua] "HGET" "d707" "d1"
1676019538.119688 [1 lua] "GET" "i901"
1676019538.122961 [1 10.246.159.29:60296] "get" "i901"
1676019538.130116 [1 10.246.159.29:60296] "get" "i1"
1676019538.134428 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d1" "dir1"

```
1676019538.134467 [1 lua] "HGET" "d1" "dir1"  
1676019538.134487 [1 lua] "GET" "i706"  
1676019538.140695 [1 10.246.159.29:60296] "get" "i706"  
1676019538.145637 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d706" "dir2"  
1676019538.145674 [1 lua] "HGET" "d706" "dir2"  
1676019538.145695 [1 lua] "GET" "i707"  
1676019538.180196 [1 10.246.159.29:60296] "get" "i707"  
1676019538.183599 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d707" "f1"
```

```
1676019538.183648 [1 lua] "HGET" "d707" "f1"  
1676019538.183676 [1 lua] "GET" "i1201"  
1676019538.186052 [1 10.246.159.29:60296] "get" "i1201"
```

getattr

| 对应 redis 一个 get 请求

```
1676019538.029964 [1 10.246.159.29:60296] "get" "i706"
```

lookup

| 向 redis 发送一个执行 lua script 请求

```
1676019538.056267 [1 10.246.159.29:60296] "evalsha" "8526c64984678fc27f67f44755c8ba7dcdf7ea60" "2" "d1" "dir1"  
  
1676019538.056313 [1 lua] "HGET" "d1" "dir1"  
  
1676019538.056339 [1 lua] "GET" "i706"
```

opendir

| 对应 redis 一个 get 请求

```
1676019538.040165 [1 10.246.159.29:60296] "get" "i707"
```

readdir

| 一个 hscan 请求获取目录下所有的 dentry
| 一个 mget 请求获取所有 dentry 对应的 inode

```
1676019538.047807 [1 10.246.159.29:60296] "hscan" "d707" "0" "match" "*" "count" "10000"
```

```
1676019538.050223 [1 10.246.159.29:60296] "mget" "i707" "i706" "i901" "i1201"
```

总结

- JuiceFS 的性能优势主要得利于内核缓存，但是开启内核缓存在极端情况下不能满足 close-to-open 一致性
- readdir 的优化：其在 readdir() 时就通过一个 batch 请求 (mget) 将该目录下所有的 inode attr 获取过来，并将 attribute/entry 缓存在内核中所以不会有 lookup/getattr 请求下发到 JuiceFS 中
- 关掉其缓存，JuiceFS(160 ms) vs CurveFS(220 ms)，主要体现在其内存 KV 上的优势，这部分优势目前可能不用考虑
- JuiceFS 提供 attr-cache, entry-cache, dir-entry-cache 3 个超时时间供用户配置，用户可能找到一种超时时间的组合，既能使用缓存，又能保证一致性