

版本信息

- OB版本：4.2.1.4

问题背景

- OCP平台持续告警租户长时间没合并成功，需要帮忙分析原因解除告警

问题诊断

OCP平台报错

- 以下报错存在关联性，租户合并失败告警以及OBServer产生了ERROR级别报错

| | | | | | | | |
|-------------------|---|-------|----------------------|---|-----------|----|-----|
| OBServer 程序日志 | 对象: ShiYuXian_OB / 172.32.1.3 类型: 日志/OceanBase 日志 | 日志 | 2024年10月28日 17:37:01 | - | 告警中: 22 天 | 严重 | 未配置 |
| OceanBase租户长时间未合并 | 对象: ShiYuXian_OB:1718595677 / shtpcinst 类型: OceanBase/租户 | 数据库内核 | 2024年10月28日 17:13:10 | - | 告警中: 22 天 | 严重 | 未配置 |
| OceanBase租户长时间未冻结 | 对象: ShiYuXian_OB:1718595677 / shtpcinst 类型: OceanBase/租户 | 数据库内核 | 2024年10月28日 17:13:10 | - | 告警中: 22 天 | 严重 | 未配置 |
| OBServer 程序日志 | 对象: ShiYuXian_OB / 172.32.1.1 类型: 日志/OceanBase 日志 | 日志 | 2024年10月28日 17:12:01 | - | 告警中: 22 天 | 严重 | 未配置 |

- OBServer ERROR级别报错的告警信息

告警详情 已恢复 3 小时 严重 告警分析数据 屏蔽告警

基本信息

ID: 11002223 应用: 日志 对象类型: OceanBase 日志 告警对象: ShiYuXian_OB / 172.32.1.3

告警规则: observer_all_error_log_alarm 产生时间: 2024年12月5日 10:29:01 恢复时间: 2024年12月5日 10:54:07 告警分类: 日志

最后一次告警时间: 2024年12月5日 10:49:01

告警事件详情

告警概述: alarm_template_id=0;ob_cluster=ShiYuXian_OB-1718595677;host=172.32.1.3;server_type=observer;error_code=4103;keyword= OBCServer 程序日志

告警详情: [OBCServer 程序日志] 集群: ShiYuXian_OB, 主机: 172.32.1.3, 日志类型: observer, 日志文件: /home/admin/oceanbase/log/observer.log, 日志级别: ERROR, 关键字=, 错误码=4103, 日志详情=[2024-12-05 10:48:10.937698] ERROR c heck_micro_block_checksum (ob_macro_block_writer.cpp:204) [82092][T1002_MAJOR_MER][T1002][YB42AC200103-00062585E72B2C7D-0-0] [lt=19][errcode=-4103] Data checksum error(msg="micro block checksum is not equal", new_checksum=703730528, checksum=860615231, ret=-4103, data_store_desc={ls_id:{id:1001}, tablet_id:{id:200929}, micro_block_size:16384, micro_block_size_limit:2095976, row_column_count:23, rowkey_column_count:3, schema_rowkey_col_cnt:1, full_stored_col_cnt:23, row_store_type:1, encoder_opt:{enable_bit_packing:true, store_sorted_var_len_numbers_dict:false, enable_raw:true, enable_dict:true, enable_int_diff:true, enable_str_diff:true, enable_hex_pack:true, enable_rle:true, enable_const:true}, compressor_type:6, schema_version:1729673192807072, merge_info:0x7f95eb54378, merge_type:4, snapshot_version:1729706403700683175, end_scn:{val:0, v:0}, need_prebuild_bloomfilter:false, bloomfilter_rowkey_prefix:0, encrypt_id:0, master_key_id:0, encrypt_key:"data_size:16, data:00000000000000000000000000000000", major_working_cluster_version:17180000516, sstable_index_builder:0x7f971dad0060, is_ddl:false, col_desc_array:[column_id=17 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=7 {type:"BIGINT", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=8 {type:"BIGINT", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=16 {type:"NUMBER", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=18 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=19 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=20 {type:"TIMESTAMP_WITH_TIME_ZONE", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=21 {type:"TIMESTAMP_WITH_TIME_ZONE", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=22 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=23 {type:"NUMBER", collation:"binary", coercibility:"NUMERIC"} order=...

[查看告警处理建议](#)

- OBCServer ERROR日志文本，显示微块校验失败，失败的表的为1002租户下tablet_id为200929的表

1 [OBCServer 程序日志] 集群: ShiYuXian_OB, 主机: 172.32.1.3, 日志类型: observer, 日志文件: /home/admin/oceanbase/log/observer.log, 日志级别: ERROR, 关键字=, 错误码=4103, 日志详情=[2024-11-19 12:17:11.173554] **ERROR check_micro_block_checksum** (ob_macro_block_writer.cpp:204) [82088][T1002_MAJOR_MER][**T1002**][YB42AC200103-00062585E72AEC8F-0-0] [lt=19][errcode=-4103] Data checksum error(msg="micro block checksum is not equal", new_checksum=703730528, checksum=860615231, ret=-4103, data_store_desc={ls_id:{id:1001}, tablet_id:{id:200929}, micro_block_size:16384, micro_block_size_limit:2095976, row_column_count:23, rowkey_column_count:3, schema_rowkey_col_cnt:1, full_stored_col_cnt:23, row_store_type:1, encoder_opt:{enable_bit_packing:true, store_sorted_var_len_numbers_dict:false, enable_raw:true, enable_dict:true, enable_int_diff:true, enable_str_diff:true, enable_hex_pack:true, enable_rle:true, enable_const:true}, compressor_type:6, schema_version:1729673192807072, merge_info:0x7f9a2d5f8378, merge_type:4, snapshot_version:1729706403700683175, end_scn:{val:0, v:0}, need_prebuild_bloomfilter:false, bloomfilter_rowkey_prefix:0, encrypt_id:0, master_key_id:0, encrypt_key:"data_size:16, data:00000000000000000000000000000000", major_working_cluster_version:17180000516, sstable_index_builder:0x7f9a2d556060, is_ddl:false, col_desc_array: [column_id=17 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=7 {type:"BIGINT", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=8 {type:"BIGINT", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=16 {type:"NUMBER", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=18 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=19 {type:"VARCHAR2", collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=20 {type:"TIMESTAMP_WITH_TIME_ZONE", collation:"binary", coercibility:"NUMERIC"} order=0, column_id=21 {type:"TIMESTAMP_WITH_TIME_ZONE", collation:"binary",

```
coercibility:"NUMERIC"} order=0, column_id=22 {type:"VARCHAR2",
collation:"utf8mb4_bin", coercibility:"INVALID"} order=0, column_id=23
{type:"NUMBER", collation:"binary", coercibility:"NUMERIC"}
```

查看当前各个租户的合并状态

- 在系统租户下，通过查看 `CDB_OB_MAJOR_COMPACTION` 视图，可以查看各个租户的合并状态。从下图可以看到，1002租户的合并状态处于COMPACTING，说明合并执行还未结束，另外也可以关注下 IS_ERROR 列，看是否存在合并错误。

在 RS Leader 所在的 OObserver 上去查看日志，基于 replica not merged 关键字

```
[2024-12-05 15:26:23.354540] INFO [RS] check_tablet_compaction_scn (ob_major_merge_progress_checker.cpp:517) [27364][T1002_MergeSche][T1002][YB42AC200101-0006287F0F15C5C0-0-0] [lt=57] replica not merged to target version or status not match (tenant_id=1002, current_version=1733374368375077283, global_broadcast_scn={val:1733378784178695206, v:0}, current_status=0, compaction_replica={tenant_id:1002, tablet_id:{id:201369}, ls_id:{id:1001}, server:"172.32.1.1:2882", snapshot_version:1733374368375077283, data_size:2841, required_size:2097152, report_scn:1733374368375077283, status:0})
[2024-12-05 15:26:23.354560] INFO [RS] check_tablet_compaction_scn (ob_major_merge_progress_checker.cpp:517) [27364][T1002_MergeSche][T1002][YB42AC200101-0006287F0F15C5C0-0-0] [lt=19] replica not merged to target version or status not match (tenant_id=1002, current_version=1733374368375077283, global_broadcast_scn={val:1733378784178695206, v:0}, current_status=0, compaction_replica={tenant_id:1002, tablet_id:{id:201369}, ls_id:{id:1001}, server:"172.32.1.2:2882", snapshot_version:1733374368375077283, data_size:2841, required_size:2097152, report_scn:1733374368375077283, status:0})
[2024-12-05 15:26:23.354569] INFO [RS] check_tablet_compaction_scn (ob_major_merge_progress_checker.cpp:517) [27364][T1002_MergeSche][T1002][YB42AC200101-0006287F0F15C5C0-0-0] [lt=8] replica not merged to target version or status not match (tenant_id=1002, current_version=1733374368375077283, global_broadcast_scn={val:1733378784178695206, v:0}, current_status=0, compaction_replica={tenant_id:1002, tablet_id:{id:201369}, ls_id:{id:1001}, server:"172.32.1.3:2882", snapshot_version:1733374368375077283, data_size:2841, required_size:2097152, report_scn:1733374368375077283, status:0})
```

确认该tablet_id对应的表

- 表为TBL_SYS_USER

```
1 MySQL [oceanbase]> select tenant_id,tablet_id,table_name from
  __all_virtual_table where tablet_id=200929
2 +-----+-----+-----+
3 | tenant_id | tablet_id | table_name |
4 +-----+-----+-----+
5 |      1002 |    200929 | TBL_SYS_USER |
6 +-----+-----+-----+
7
```

确认存储层 tablet 合并调度是否正常

- 查看合并诊断信息表，确认为FAILED状态，且附加信息中，记录的确为tablet_id=200929的表，也就是TBL_SYS_USER表合并失败导致整体合并进程失败。

```
1 MySQL [oceanbase]> select * from __all_virtual_compaction_diagnose_info;
```


- 调整过该参数为0后，合并正常

```
1 alter system set micro_block_merge_verify_level=0; -- 合并时关闭微块校验，跳过了合并微块校验，成功
2 alter system set micro_block_merge_verify_level=2; -- 合并过一次后再开启校验，依旧成功
```

初步结论

- 大概率是BUG，但没法分析到其中原由，目前只是采用临时方案绕过

参考链接

- <https://www.oceanbase.com/knowledge-base/oceanbase-database-100000000008209?back=kb>
- https://alidocs.dingtalk.com/i/nodes/a9E05BDRVQRkezKGC3E6rvw2J63zgzYA?iframeQuery=utm_source%3Dportal%26utm_medium%3Dportal_recent