

Problemes de recherche et indexation



This page is no longer being updated. From BlueMind 4.8, please refer to the [new BlueMind documentation](#)

- [The search bar doesn't show the option to search all folders](#)
- [Some search results are missing](#)
- [Error message during searches](#)
- [An error message appears when trying to access a message found by search](#)
- [ElasticSearch goes into "read-only" mode](#)
- [Logs show esQuota and imapQuota errors](#)
- [Global malfunction](#)

The search bar doesn't show the option to search all folders

For all users

If all users are affected, the issue probably occurs following an email migration at the folder system level. This means that the Elasticsearch search index doesn't exist: you can run `ConsolidateMailSpoolIndexJob` to create the index and index all the messages on the server:

- go to the admin console > System Management > Planning > run `ConsolidateMailSpoolIndexJob`

or

- using our [command line tool](#), run the `consolidateIndex` maintenance operation:

```
bm-cli maintenance consolidateIndex domain.net
```

This operation is resource hungry. From BlueMind 3.5.12, you can run it by groups of users using the option `--match`:

```
bm-cli maintenance consolidateIndex --match "a.*" domain.net
bm-cli maintenance consolidateIndex --match "[b-c].*" domain.net
```

For a few users only

If the issue affects a few users only, this means that their Elasticsearch index doesn't exist or is corrupt, you therefore have to create it:

- either by going into each user's admin page and running "**Validate and repair user data**" next "**Consolidate mailbox index**".
- or using our [command line tool](#):

```
bm-cli maintenance repair user@domain.net
bm-cli maintenance consolidateIndex user@domain.net
```

Some search results are missing

In case of temporary issues with the indexing service, some messages sent or received during that period may not have been indexed. Simply run `ConsolidateMailSpoolIndexJob` which computes the difference between messages within IMAP and the index, and then indexes the missing messages only.

Error message during searches

This may be caused by an inconsistency between the list of IMAP folders and the database. You can use the "validate and repair" (check and repair) maintenance operation -- which can be accessed from the Maintenance tab in the user's admin page – to rebuild this list. Then re-indexing the mailbox should fix the issue.

If this isn't the case, `/var/log/bm-webmail.errors` logs may be the cause.

An error message appears when trying to access a message found by search

This is probably due to an indexing fault when the message was moved. You can use the "Consolidate mailbox index" maintenance operation – which can be accessed from the Maintenance tab in the user's admin page – to update the search index.

ElasticSearch goes into "read-only" mode

Issue: ES puts itself in "read-only" mode, the cluster is green but a restart doesn't fix the issue.

Symptoms: Overall, BlueMind doesn't seem to work any more. Besides search and display issues, users can't create events, reply to emails, etc.

Confirmation: In ElasticSearch logs, you will find traces such as:

```
2018-12-27 08:15:58,273 [BM-Core21] n.b.c.r.b.RestServiceMethodHandler ERROR - Error during restcall
RestRequest [path=/internal-api/db_message_bodies/bm-master__blue-mind_loc
/8b904e61a5d73ceb8b7e02048986a1740c94b3bb, method=PUT, User-Agent=null, params=, remoteAddresses=
[192.168.132.14], origin=null]: class net.bluemind.core.api.fault.ServerFault: java.util.concurrent.
ExecutionException: ClusterBlockException[blocked by: [FORBIDDEN/12/index read-only / allow delete (api)];]
  at net.bluemind.backend.mail.replica.service.internal.DbMessageBodiesService.create
(DbMessageBodiesService.java:101)
  ... 18 common frames omitted
Caused by:
  class java.util.concurrent.ExecutionException: ClusterBlockException[blocked by: [FORBIDDEN/12/index read-
only / allow delete (api)];]
  at java.util.concurrent.CompletableFuture.reportGet(CompletableFuture.java:357)
  at java.util.concurrent.CompletableFuture.get(CompletableFuture.java:1915)
  at net.bluemind.backend.mail.replica.service.internal.DbMessageBodiesService.create
(DbMessageBodiesService.java:98)
  ... 18 common frames omitted

2018-12-27 08:15:58,274 [vert.x-eventloop-thread-1] n.b.b.c.r.s.s.ReplicationState ERROR - addMessage.
create: java.util.concurrent.ExecutionException: ClusterBlockException[blocked by: [FORBIDDEN/12/index read-
only / allow delete (api)];]
net.bluemind.core.api.fault.ServerFault: java.util.concurrent.ExecutionException: ClusterBlockException
[blocked by: [FORBIDDEN/12/index read-only / allow delete (api)];]
  at net.bluemind.core.rest.base.codec.JsonObjectCodec.parseFault(JsonObjectCodec.java:111)
  at net.bluemind.core.rest.base.codec.DefaultResponseCodecs$VoidCodec.decode(DefaultResponseCodecs.java:
288)
  at net.bluemind.core.rest.base.codec.DefaultResponseCodecs$VoidCodec.decode(DefaultResponseCodecs.java:1)
  at net.bluemind.core.rest.base.ClientProxyGenerator$MethodCallBuilder.parseResponse(ClientProxyGenerator.
java:417)
  at net.bluemind.core.rest.base.ClientProxyGenerator$EventMethodInvoker$1.success(ClientProxyGenerator.
java:215)
  at net.bluemind.core.rest.base.ClientProxyGenerator$EventMethodInvoker$1.success(ClientProxyGenerator.
java:1)
  at net.bluemind.core.rest.log.CallLogger$1.success(CallLogger.java:48)
  at net.bluemind.core.rest.log.CallLogger$1.success(CallLogger.java:1)
  at net.bluemind.core.rest.base.RestRootHandler$1.success(RestRootHandler.java:130)
  at net.bluemind.core.rest.base.RestRootHandler$1.success(RestRootHandler.java:1)
  at net.bluemind.core.rest.base.RestRootHandler$VertxAwareAsyncHandler$1.handle(RestRootHandler.java:172)
  at net.bluemind.core.rest.base.RestRootHandler$VertxAwareAsyncHandler$1.handle(RestRootHandler.java:1)
  at org.vertx.java.core.impl.DefaultContext$2.run(DefaultContext.java:144)
  at org.vertx.java.core.impl.DefaultContext$3.run(DefaultContext.java:175)
  at io.netty.util.concurrent.SingleThreadEventExecutor.runAllTasks(SingleThreadEventExecutor.java:399)
  at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:464)
  at io.netty.util.concurrent.SingleThreadEventExecutor$2.run(SingleThreadEventExecutor.java:131)
  at io.netty.util.concurrent.FastThreadLocalRunnable.run(FastThreadLocalRunnable.java:30)
  at java.lang.Thread.run(Thread.java:748)
2018-12-27 08:15:58,274 [vert.x-eventloop-thread-1] n.b.b.c.r.s.c.ApplyMessage INFO - Finished ApplyMessage.
2018-12-27 08:15:58,274 [vert.x-eventloop-thread-1] n.b.b.c.r.s.ReplicationSession INFO - REPL S: [frame-
00001297]: OK success
```

Cause: The server disk space is too tight

Solution:

1. add disk space (increase the partition, change the disk, etc.)

2. once this is done, play the following command:

```
curl -X PUT "localhost:9200/_all/_settings" -H 'Content-Type: application/json' -d' { "index.blocks.read_only_allow_delete" : false } }'
```

Logs show esQuota and imapQuota errors

You find messages such as the one below in `/var/log/bm-webmail/errors`:

```
10-Nov-2019 17:37:38 UTC] [jdoe@bluemind.loc] esQuota < (imapQuota * 0.8). disable es search. esQuota: 4199171, imapQuota: 6123568
```

This means that for the account shown, less than 80% of the mailbox is indexed (esQuota = elasticsearch quota), elasticsearch search (== advanced search engine) is therefore disabled because inefficient.

To fix this, you have to consolidate or reindex the account.

If only a few identified users are affected

If the issue only affects one or a few users, this means that their ElasticSearch index doesn't exist or is corrupt – it has to be created again:

- either by going into each user's admin page and executing "**Validate and repair user data**" then "**Consolidate mailbox index**" then, if there's no improvement, "**Reconstruct mailbox index**"
- or by using our [command line tool](#):

```
bm-cli maintenance repair user@domain.net
bm-cli maintenance consolidateIndex user@domain.net
```

If all users are affected

To repair all accounts, you can:

1. find the accounts by running a grep on the log file:

```
grep "disable es search. esQuota:" /var/log/bm-webmail/errors
```

2. copy the logins found into a text file (e.g. `/tmp/accountWithoutEsSearch.txt`)
3. use the following command combination to start the consolidation of the index for each file login:

```
while read account; do bm-cli maintenance consolidatedIndex $account;done < /tmp/accountWithoutEsSearch.txt
```

Global malfunction

Analyse

If you detect a **search malfunction in BlueMind**, you can see the status of the ElasticSearch cluster using the command:

```
curl -XGET --silent 'http://localhost:9200/_cluster/health'
```

If the cluster's status is 'green' then everything is fine, if it is 'red', this means there is an issue with Elasticsearch. This information is also fed into the [monitoring console](#).

Using our [command line admin tool](#), you can check a user's indexing status:

```
$ bm-cli index info admin@local.lan
{"email": "jdoe@domain.loc", "ESQuota": 3056, "IMAPQuota": 3058, "ratio": 95}
```

Solution

Several issues may stop ElasticSearch from working:

- **an insufficiently indexed mailbox:** if, when you check the index (see above), you find a **ratio below 80**, this means that fewer than 80% of the user's emails are indexed → you must reindex the mailbox:
 - go to the BlueMind admin console
 - go to the user's admin page > Maintenance tab
 - run **"Reconstruct mailbox index"**
- **corrupt index:** mainly due to low disk space. **You need at least 10% of free disk space.** If the disk containing ES data (`/var/spool/bm-elasticsearch`) has been short on space, search indices may have become corrupt. In ES logs, this translates as an error on service start up:

```
[2017-01-26 20:06:54,764][WARN ][cluster.action.shard] [Bill Foster] [mailspool][0] received shard failed for [mailspool][0], node[PcC6eICxRAajmWioKlmhDA], [P], s[INITIALIZING], indexUUID [IEJHQkOnTtOcdY0bMMIFRA], reason [master [Bill Foster][PcC6eICxRAajmWioKlmhDA][bluemind.exa.net.uk] [inet[/82.219.13.101:9300]] marked shard as initializing, but shard is marked as failed, resend shard failure]
[2016-01-26 20:06:55,828][WARN ][indices.cluster] [Bill Foster] [mailspool][0] failed to start shard org.elasticsearch.index.gateway.IndexShardGatewayRecoveryException: [mailspool][0] failed to recover shard
    at org.elasticsearch.index.gateway.local.LocalIndexShardGateway.recover
(LocalIndexShardGateway.java:287)
    at org.elasticsearch.index.gateway.IndexShardGatewayService$1.run(IndexShardGatewayService.java:132)
    at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145)
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615)
    at java.lang.Thread.run(Thread.java:745)
```

You must therefore:

1. Remove the index files and restart ElasticSearch:

```
service bm-elasticsearch stop
rm -rf /var/spool/bm-elasticsearch/data/nodes/0/indices/*
service bm-elasticsearch start
```

2. Reset the indices:

```
bm-cli index reset mailspool
bm-cli index reset mailspool_pending
bm-cli index reset event
bm-cli index reset contact
bm-cli index reset todo
bm-cli index reset im
```

3. Then start indexing again from scheduled jobs: Admin Console > System Management > Planning > select the domain "global.virt" and run all the `*IndexJob` jobs:

- CalendarIndexJob
- ContactIndexJob
- ConsolidateMailSpoolIndexJob



Email indexing is IO hungry, so **we recommend that you run it in the evening or at the weekend.** You can also run the indexing by lot using `bm-cli`.

- TodoListIndexJob
- HSMIndexJob

- **corrupt translog:** this can happen if **the server has crashed or because of low memory.** In this case, the general index is not corrupt and only the indexing of documents not written to the disk yet will be lost. In ES logs, this translates as the following error during service restart:

```
[2017-09-04 19:24:38,340][WARN ][indices.cluster          ] [Hebe] [mailspool][1] failed to start
shard
org.elasticsearch.index.gateway.IndexShardGatewayRecoveryException: [mailspool][1] failed to recover
shard
    at org.elasticsearch.index.gateway.local.LocalIndexShardGateway.recover
(LocalIndexShardGateway.java:287)
    at org.elasticsearch.index.gateway.IndexShardGatewayService$1.run(IndexShardGatewayService.
java:132)
    at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145)
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615)
    at java.lang.Thread.run(Thread.java:745)
Caused by: org.elasticsearch.index.translog.TranslogCorruptedException: translog corruption while
reading from stream
    at org.elasticsearch.index.translog.ChecksummedTranslogStream.read(ChecksummedTranslogStream.
java:70)
    at org.elasticsearch.index.gateway.local.LocalIndexShardGateway.recover
(LocalIndexShardGateway.java:257)
    ... 4 more
```

To remove the corrupt translogs:

```
service bm-elasticsearch stop
rm -rf /var/spool/bm-elasticsearch/data/nodes/0/indices/mailspool/*/translog
service bm-elasticsearch start
```

Running ConsolidateMailSpoolIndexJob reindexes the missing messages.