

HOSTED SAAS ELASTICSEARCH ISSUES WITH DOC.TIKI.ORG AND DEV.TIKI.ORG

Status

✖ Closed

Priority

0 Emergency

Task

Hosted SaaS Elasticsearch issues with doc.tiki.org and dev.tiki.org

Add a timesheet

[Add a timesheet for this task](#)

Description

It has been happening on and off in last few days: <https://dev.tiki.org/item8056-Wishlist-broken>



Could be related to recent upgrade to Tiki 24.x

[item73591-Tiki-Association-Elasticsearch-cloud-deployment-account](#)

[item18781-tiki-suite-wiki-root-and-MariaDB-root](#)

Files

1.  [Photo 2022 02 03 01 04 40](#) 

2.  [2022 02 05 005640](#) 

Emails

Assigned to

ricardo.melo, victor.emanouilov

Assigned by

bernard.sfez, marclaporte

Keep informed

bernard.sfez, brendan.ferguson, fabio.montefusco, gary.cunningham-lee, habiba.shahbaz, jean-

marc.libs, jonny.bradley, lukas.masek, oliver.hertel, roberto.kirschbaum, xavier.depedro

Also visible to

1- Visible to Anonymous (Public content)

Related

Created

Thursday 03 of February, 2022 01:07:27 EST

by Marc Laporte

LastModif

Wednesday 27 of March, 2024 05:26:36 EDT

Comments



Findings

victor.emanouilov 2022-02-03 03:56

Managed service runs just fine. It has some errors in the log but they are not connected with this downtime - channel closed errors might mean remote connection closed (by Tiki, for example) and ES service didn't properly handle that. I notice other users experiencing the same and it seems like a minor bug in the ES itself.

Noticed that the last rebuild didn't correctly assign the current index preference - it was pointing to an index that did not exist. Thus the search errors we see. It seems a rebuild isn't completing successfully for some reason.

I tested a couple of rebuilds from the console - it works just fine - current index pref is updated at the end and searching works.

I tested rebuilding from admin UI - php-fpm cut the process at the 31st second - config that disables long running scripts from the web. Index was partially built but current index pref was still pointing to the old index and searching was working.

Subsequent rebuild from the console correctly assigned the current index pref and index is still working.

So, I am unable to reproduce the issue. It might be someone trying to rebuild the index from the UI several times or something similar. Let's warn admins to rebuild index only via console. I see a couple of rebuild attempts in the past 24 hours from the ES managed service logs.



@Victor: I found the root cause
[Marc Laporte](#) 2022-02-05 01:04

I had seen this before: [item68719-After-a-Tiki-Manager-clone-with-Elasticsearch-index-prefix-in-tiki-ini-there-is-a-mismatch-with-Elasticsearch-current-index-that-can-cause-2-Tikis-to-write-to-one-index](#) but we never got to it. Our sites use our own Elasticsearch instances so we can rebuild at will.



Yes, this does seem to be the issue
[victor.emanouilov](#) 2022-02-07 11:01

I manually rebuilt the index on nextdoc which created successfully an index with the forced prefix. It seems when you clone a Tiki instance, not only the forced prefix must be set but also the current index prefix must be cleared and then the index manually rebuilt on the cloned instance. Otherwise, Elastic ends up mixing indexes as it uses index aliases.

One odd thing I noticed is that nextdoc rebuilt for 2 mins while doc is rebuilding more than an hour now... strange. Can check this more tomorrow.

Re: nextdoc - it failed to load in the browser due to a hanging connection to sentry error reporting lib on evoludata - I had to turn off error reporting on nextdoc to load it. Another thing to debug...



ok, GlitchTip is an Open Source alternative to sentry
[Marc Laporte](#) 2022-02-07
11:14

Here is our server: [item61498-GlitchTip-](#)

[self-hosted](#)

You can deactivate GlitchTip if it is causing issues.

<https://doc.tiki.org/GlitchTip>



yes, this is nice but

victor.emanouilov

2022-02-08 09:01

it causes some problems. I think the slow reindexing of doc.tiki.org (took 2 hours yesterday) is due to some php warnings that end up sent to glitchtip rather slowly and during rebuild this piles up to an enourmous amount of time. Not sure how to solve this besides fixing those errors which was my plan anyway - it is a good thing to clean up indexing errors before release. Do you want me to proceed here?

nextdoc.tiki.org still hangs trying to send error reports to the designated endpoint. I checked the self-hosted glitchtip setup and it seems correct. I have never used it, so it might take time to debug and see exactly why. Is this a priority and should I proceed or you will ask whoever installed that to debug more?



Additional information about the error (doc) Bernard Sfez 2022-02-07 05:28

This morning error after saving a page while connected as "bsfez" not "Medizkit".



```
x
  Error
Malformed search query: Parsing search query failed: "Field
searchable does not exist in the current index. If this is a
tracker field, the proper syntax is tracker_field_searchable."
x
  Error
Elasticsearch flush error: no such index (for object user-Medizkit)
x
  Error
The search index could not be updated. The site is misconfigured.
Contact an administrator.
no such index (for object user-Medizkit)
```



Results of further debugging and fix victor.emanouilov 2022-02-09 10:39

It turns out that the server is using Tiki Manager to clone and upgrade nightly all next.* sites. This copies original database version to next sites including the current ES index. When rebuilding over the same ES index from 2 sites, one of them deletes the index but the other one still uses it, thus the conflicts.

I have updated Tiki Manager on the server as a hot fix to delete `pref unified_elastic_index_current` when cloning. If there are no cloning errors, it will reindex the new site, so the new index prefix should be used. Tested this manually and it works. (another error is that currently it fails with `cyph` asset files standing on the way but I have fixed that in latest 24.x and master, so no sure why it fails again for nextdoc - maybe it starts from an older version of the branch? - please ask Ricardo for help here).

I am pasting a suggested patch of Tiki Manager here as I ran out of time to submit a PR - maybe Ricardo and team can take a look and apply it?



```
[root@tiki app]# git diff
diff --git a/src/Application/Instance.php
b/src/Application/Instance.php
index 304ad03..c6fe621 100644
--- a/src/Application/Instance.php
+++ b/src/Application/Instance.php
@@ -747,7 +747,7 @@ SQL;
     if ($databaseConfig) {
         try {
             $app = $this->getApplication();
-            $app->restoreDatabase($databaseConfig,
+            $app->restoreDatabase($databaseConfig,
 $database_dump);
+            $app->restoreDatabase($databaseConfig,
 $database_dump, $clone);
             $app->setPref('tmpDir', $this->tempdir);
         } catch (\Exception $e) {
             $restore->unlock();
diff --git a/src/Application/Tiki.php b/src/Application/Tiki.php
index 9216d4a..20319e3 100644
--- a/src/Application/Tiki.php
+++ b/src/Application/Tiki.php
@@ -582,7 +582,7 @@ class Tiki extends Application
     );
 }

-    public function restoreDatabase(Database $database,
+    public function restoreDatabase(Database $database,
 $remoteFile)
+    public function restoreDatabase(Database $database,
 $remoteFile, $clone = false)
    {
        $tmp = tempnam($_ENV['TEMP_FOLDER'], 'dblocal');

@@ -632,6 +632,10 @@ class Tiki extends Application
        if (!empty($errors)) {
            throw new \RuntimeException(implode(PHP_EOL,
 $errors));
        }

+        if ($clone) {
+            $database->query("DELETE FROM
 {$database->dbname}.tiki_preferences WHERE name =
 'unified_elastic_index_current'");
+        }
    }
}
```

//-----



Victor found it. It's because of Elasticsearch index collision with <https://nextdev.dogfood2.evoludata.com> and <https://nextdoc.dogfood2.evoludata.com>

[Marc Laporte](#) 2022-02-17 13:11

We have:

https://dev.tiki.org/Pre-dogfood-servers-for-Tiki-24-release-process#On_a_second_server



```
Victor Emanouilov, [2022-02-17 5:43]
```

```
Marc, tracked the index collisions down to this: this morning, live doc.t.o index was deleted and recreated around 6:29-6:32 am UTC according to elastic.co logs:
```

```
Victor Emanouilov, [2022-02-17 5:43]
```

```
Feb 17, 2022, 6:29:23 AM UTC
```

```
INFO
```

```
i3@eu-central-1c
```

```
[instance-0000000003]
```

```
[doc_tiki_org_main_620deb41be437/CZCzjkU6RhKwacxAJ02UQA]
```

```
create_mapping [_doc]
```

```
Victor Emanouilov, [2022-02-17 5:44]
```

```
Feb 17, 2022, 6:34:03 AM UTC
```

```
INFO
```

```
i3@eu-central-1c
```

```
[instance-0000000003] [doc_tiki_org_pref_en/_ES56rQMR0is1mfQDwyKg]
```

```
deleting index
```

```
Victor Emanouilov, [2022-02-17 5:44]
```

```
however, the live doc.t.o site's current index still points to the deleted one here
```

```
Victor Emanouilov, [2022-02-17 5:44]
```

```
something has rebuilt the index and deleted the old one but not
```

running from the live site

Victor Emanouilov, [2022-02-17 5:45]

I checked all the logs I found on the server and nothing has run 6:29-6:32 am on that server

Victor Emanouilov, [2022-02-17 5:45]

is it possible that doc.tiki.org has an instance on another server which runs index rebuild aroyund 6:30am?

Victor Emanouilov, [2022-02-17 5:46]

nextdoc rebuilds fine, instance update and clone and upgrade commands work fine, I double checked index prefixes, tried manual commands and it always succeeds with rebuild without disrupting other site indexes

Victor Emanouilov, [2022-02-17 5:46]

so it must be a rebuild processes running on another server but using the same ES managed instance

Victor Emanouilov, [2022-02-17 5:47]

the server I am checking on is this:

Victor Emanouilov, [2022-02-17 5:47]

IP address of tiki.suite.wiki is 172.104.155.141

Marc Laporte, WikiSuite.org, [2022-02-17 8:37]

[In reply to Victor Emanouilov]

Yes, there is another clone site

Marc Laporte, WikiSuite.org, [2022-02-17 8:38]

https://dev.tiki.org/Pre-dogfood-servers-for-Tiki-24-release-process#0n_a_second_server

Victor Emanouilov, [2022-02-17 8:38]

btw, I am rebuilding index on trunk.avan.tech site now, just FYI

Marc Laporte, WikiSuite.org, [2022-02-17 8:39]

<https://avan.tech/dogfood2-server>

Victor Emanouilov, [2022-02-17 9:03]

that is the problem, I see cron jobs running in the morning which clone production doc, dev and themes to local copies and then clone and upgrade to nextdoc/dev/themes

Victor Emanouilov, [2022-02-17 9:04]

when the ES server was local, it was not a problem as there were 2 ES servers

Victor Emanouilov, [2022-02-17 9:04]
however, since it is now an external service, the cloned version on
dogfood2 server reindexes, which deletes the original index on the
external ES server

Victor Emanouilov, [2022-02-17 9:04]
this leaves the live site without an index

Victor Emanouilov, [2022-02-17 9:04]
we need tiki.ini overrides for all dogfood2 sites

Marc Laporte, WikiSuite.org, [2022-02-17 9:05]
Thank you for finding.



I set tiki.ini.php for 3 sites

Marc Laporte 2022-02-17 13:48

<https://nextdev.dogfood2.evoludata.com>

<https://nextdoc.dogfood2.evoludata.com>

<https://nextthemes.dogfood2.evoludata.com>



Related

Marc Laporte 2022-02-18 10:44

https://gitlab.com/tikiwiki/tiki-manager/-/merge_requests/267