
Envie de booster ta carrière ?
Open source-toi !
DevFest
Dijon
25


ORACLE
DB
Oracle
bea
BEA
WebLogic
Java
J2EE
EJB
JMS
Eclipse IDE
PostgreSQL
JBoss
Spring
Framework
JSP
HTML
JS
CSS


DOUANIERS
DOUANIERS
DOUANIERS


Maitriser GIT - Version Control for Everyone
Commands
Wens
Paris JUG
Paris JUG
Paris JUG
Paris JUG

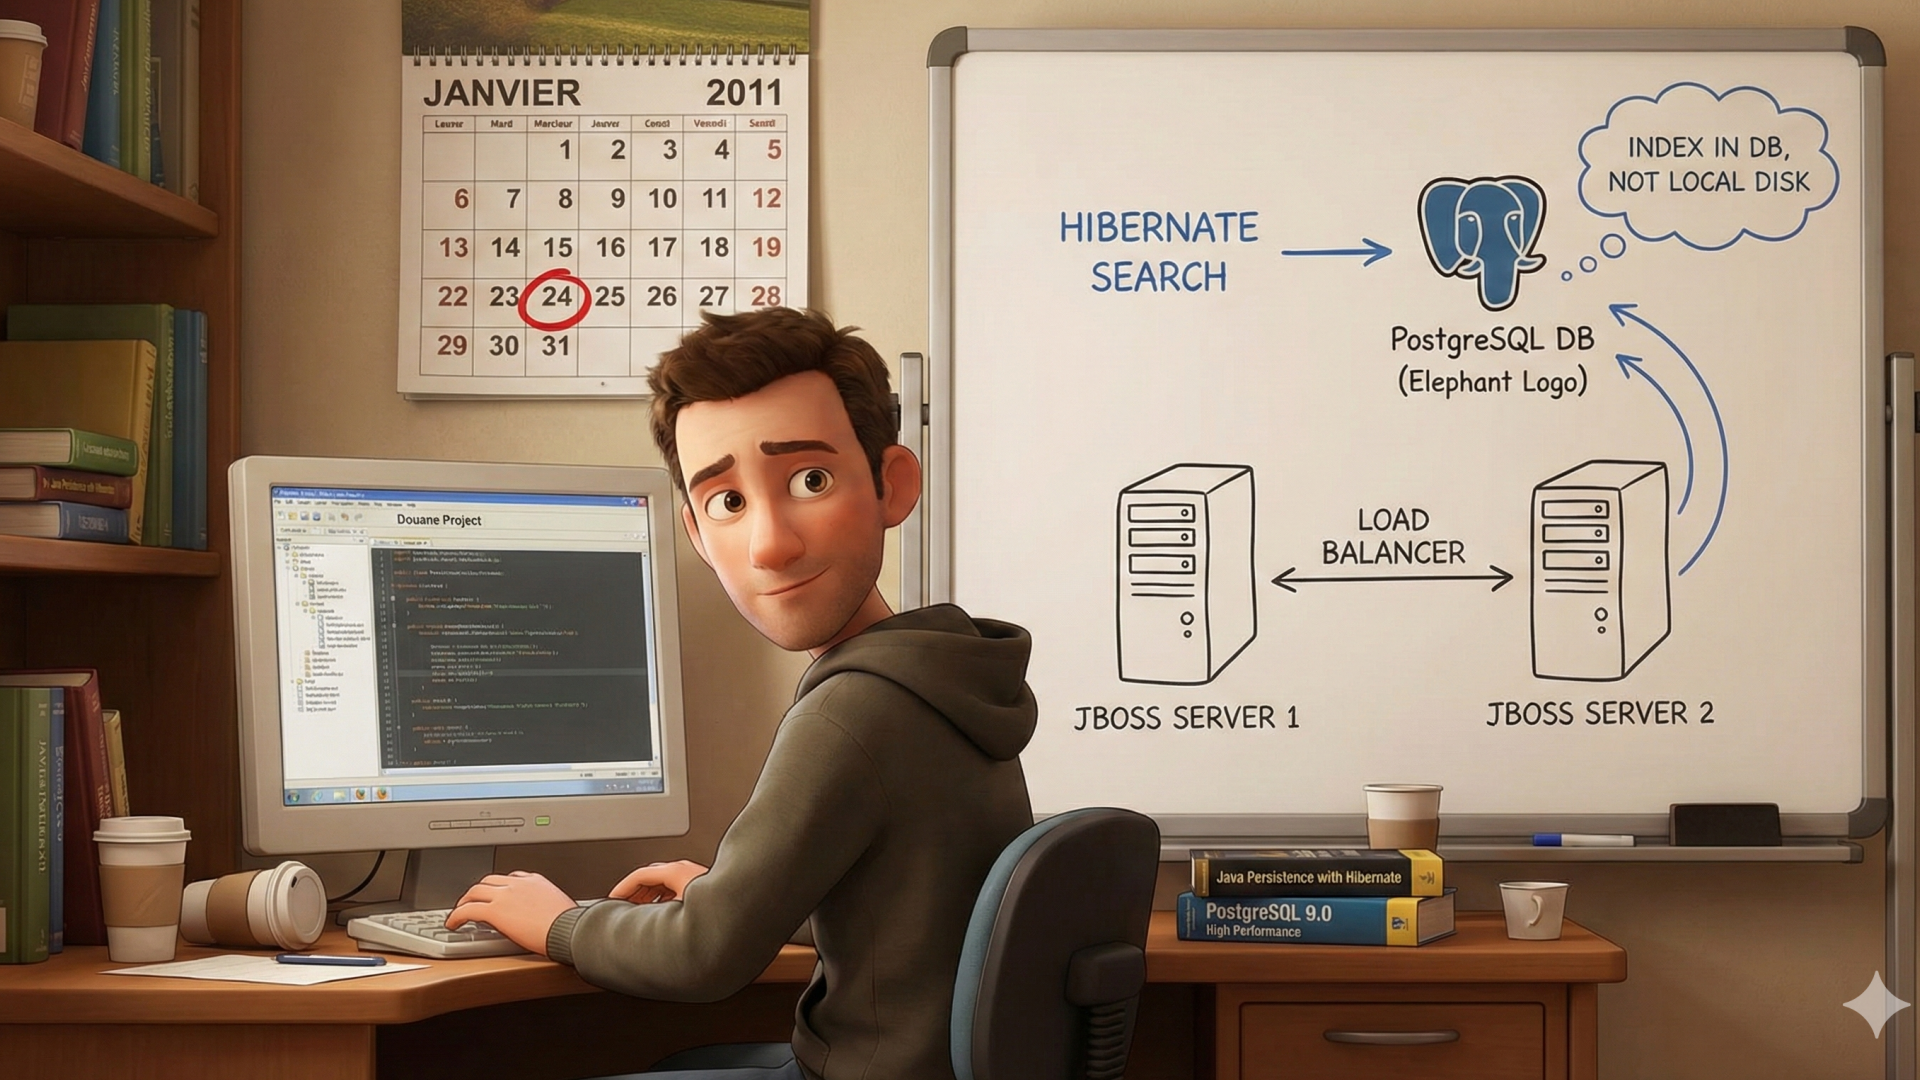
JANVIER 2011
HIBERNATE SEARCH
INDEX IN DB, NOT LOCAL DISK
PostgreSQL DB (Elephant Logo)
LOAD BALANCER
JBOSS SERVER 1
JBOSS SERVER 2
Douane Project
Java Persistence with Hibernate
PostgreSQL 9.0 High Performance

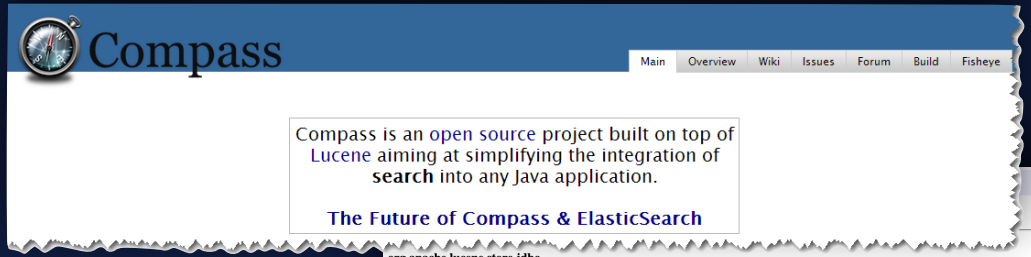

org.apache.lucene.store.jdbc

Class JdbcDirectory

java.lang.Object
  extended by org.apache.lucene.store.Directory
    extended by **org.apache.lucene.store.jdbc.JdbcDirectory**

All Implemented Interfaces:

MultiDeleteDirectory

```
public class JdbcDirectory
extends Directory
implements MultiDeleteDirectory
```

A Jdbc based implementation of a Lucene `Directory` allowing the storage of a Lucene index within a database. Uses a jdbc `DataSource`, `Dialect` specific for the database used, and an optional `JdbcDirectorySettings` and `JdbcTable` for configuration.

The directory works against a single table, where the binary data is stored in `Blob`. Each "file" has an entry in the database, and different `FileEntryHandler` can be defines for different files (or files groups).

Most of the files will not be deleted from the database when the directory delete method is called, but will only be marked to be deleted (see `MarkDeleteFileEntryHandler`. It is done since other readers or searchers might be working with the database, and still use the files. The ability to purge mark deleted files based on a "delta" is acheived using `deleteMarkDeleted()` and `deleteMarkDeleted(long)`. Note, the purging process is not called by the directory code, so it will have to be managed by the application using the jdbc directory.

For transaction management, all the operations performed against the database do not call `commit` or `rollback`. They simply open a connection (using `DataSourceUtils.getConnection(javax.sql.DataSource)`), and close it using `DataSourceUtils.releaseConnection(java.sql.Connection)`). This results in the fact that transcation management is simple and wraps the directory operations, allowing it to span as many operations as needed.

For none managed applications (i.e. applications that do not use JTA or Spring transaction manager), the jdbc directory implementation comes with `TransactionAwareDataSourceProxy` which wraps a `DataSource` (should be a pooled one, like Jakartat DBCP). Using it with the `DataSourceUtils`, or the provided `DirectoryTemplate` should make integrating or using jdbc directory simple.

Also, for none managed applications, there is an option working with autoCommit=true mode. The system will work much slower, and it is only supported on a portion of the databases, but any existing code that uses Lucene with any other `Directory` implemenation should work as is.

If working within managed environments, an external transaction management should be performed (using JTA for example). Simple solutions can be using CMT or Spring Framework abstraction of transaction managers. Currently, the jdbc directory implementation does not implement a transaction management abstraction, since there is a very good solution out there already (Spring and JTA). Note, when using Spring and the `DataSourceTransactionManager`, to provide the jdbc directory with a Spring's `TransactionAwareDataSourceProxy`.

Author:

kimchy

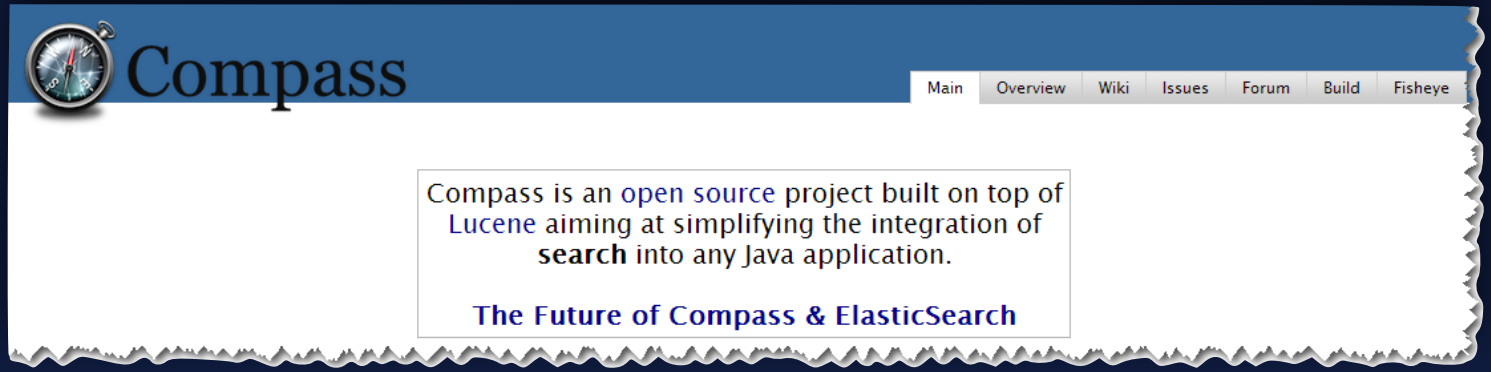
Compass
Main
Overview
Wiki
Issues
Forum
Build
Fisheye
Compass is an open source project built on top of Lucene aiming at simplifying the integration of search into any Java application.
The Future of Compass & ElasticSearch

Search

Fork me on GitHub

elasticsearch.

home
download
guide
blog
community
tutorials
videos

You know, for Search

So, we build a web site or an application and want to add search to it, and then it hits us: **getting search working is hard**. We want our search solution to be **fast**, we want a **painless setup** and a completely **free search schema**, we want to be able to index data simply using **JSON over HTTP**, we want our search server to be **always available**, we want to be able to start with one machine and **scale to hundreds**, we

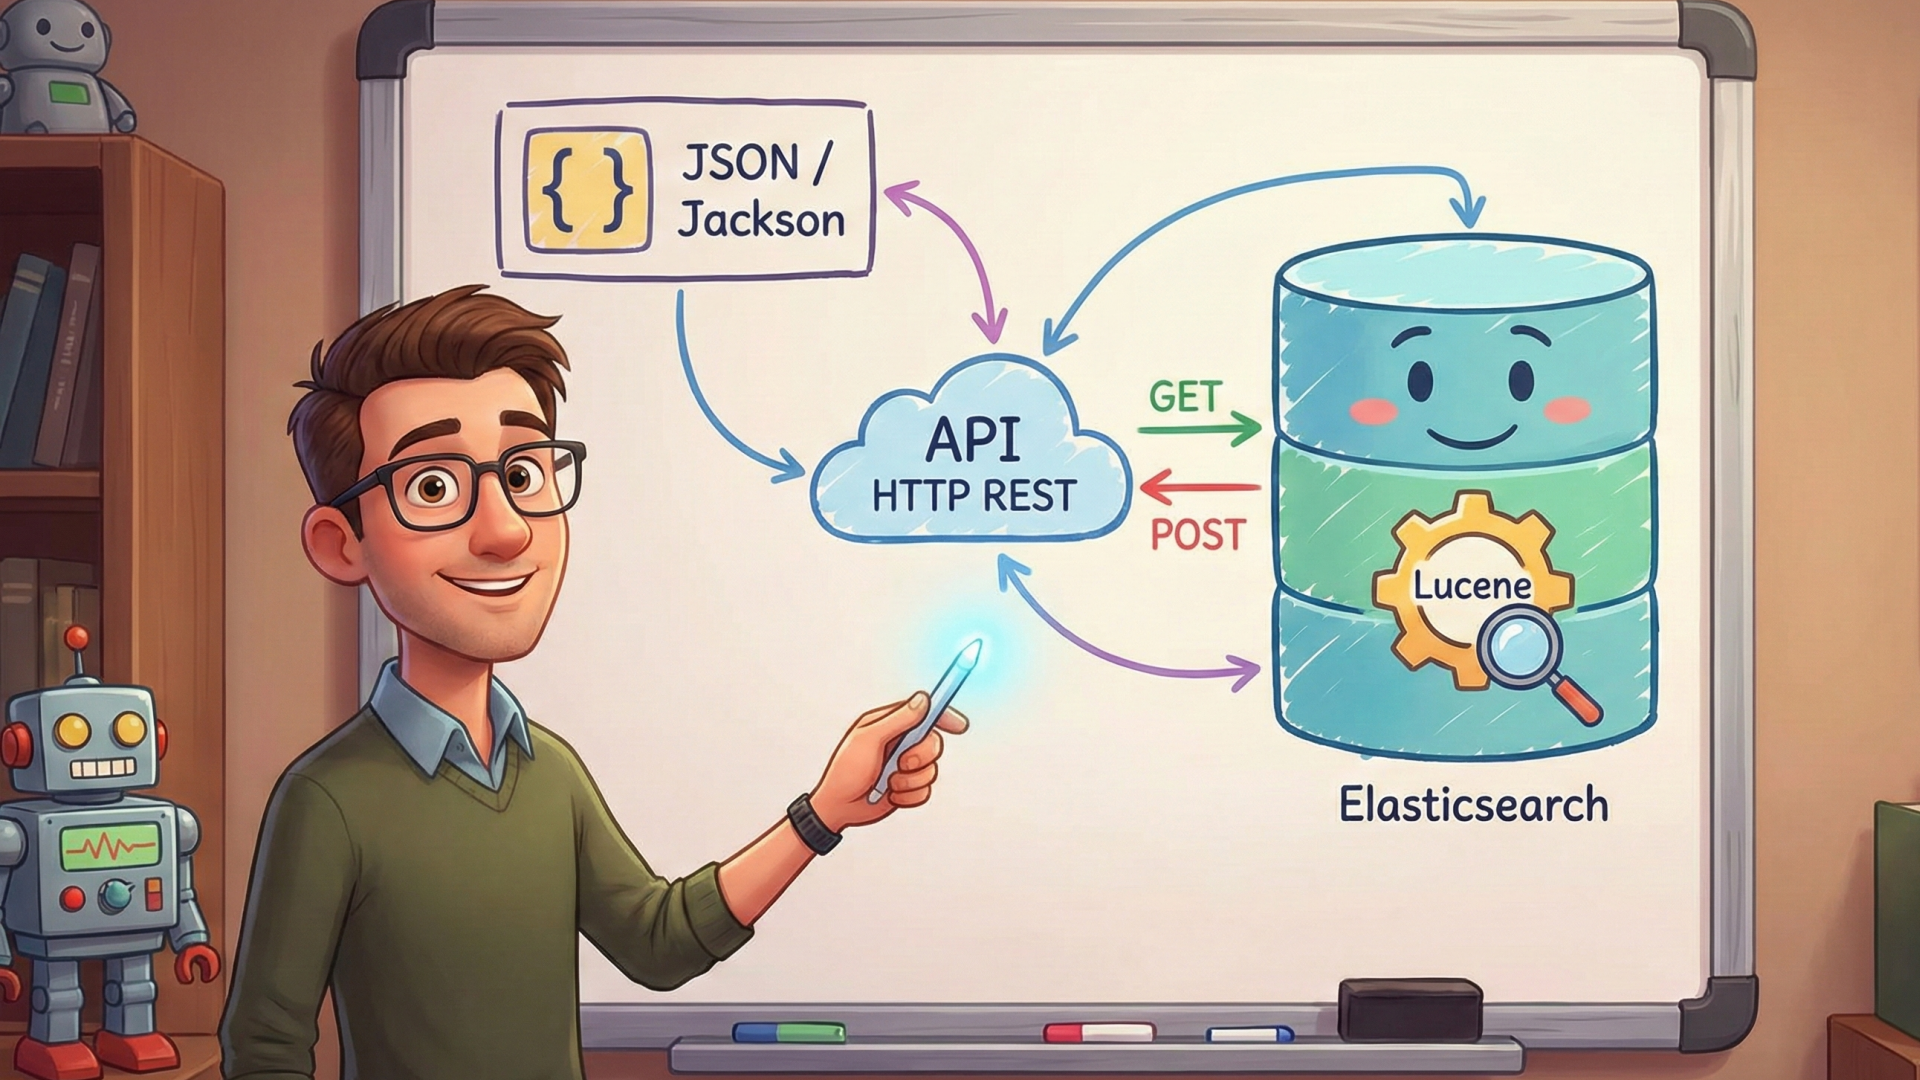
{ }
JSON / Jackson
API
HTTP REST
GET
POST
Lucene
Elasticsearch


COMMUNAUTÉ D'ENTRAIDE


Flux RSS
Elasticsearch

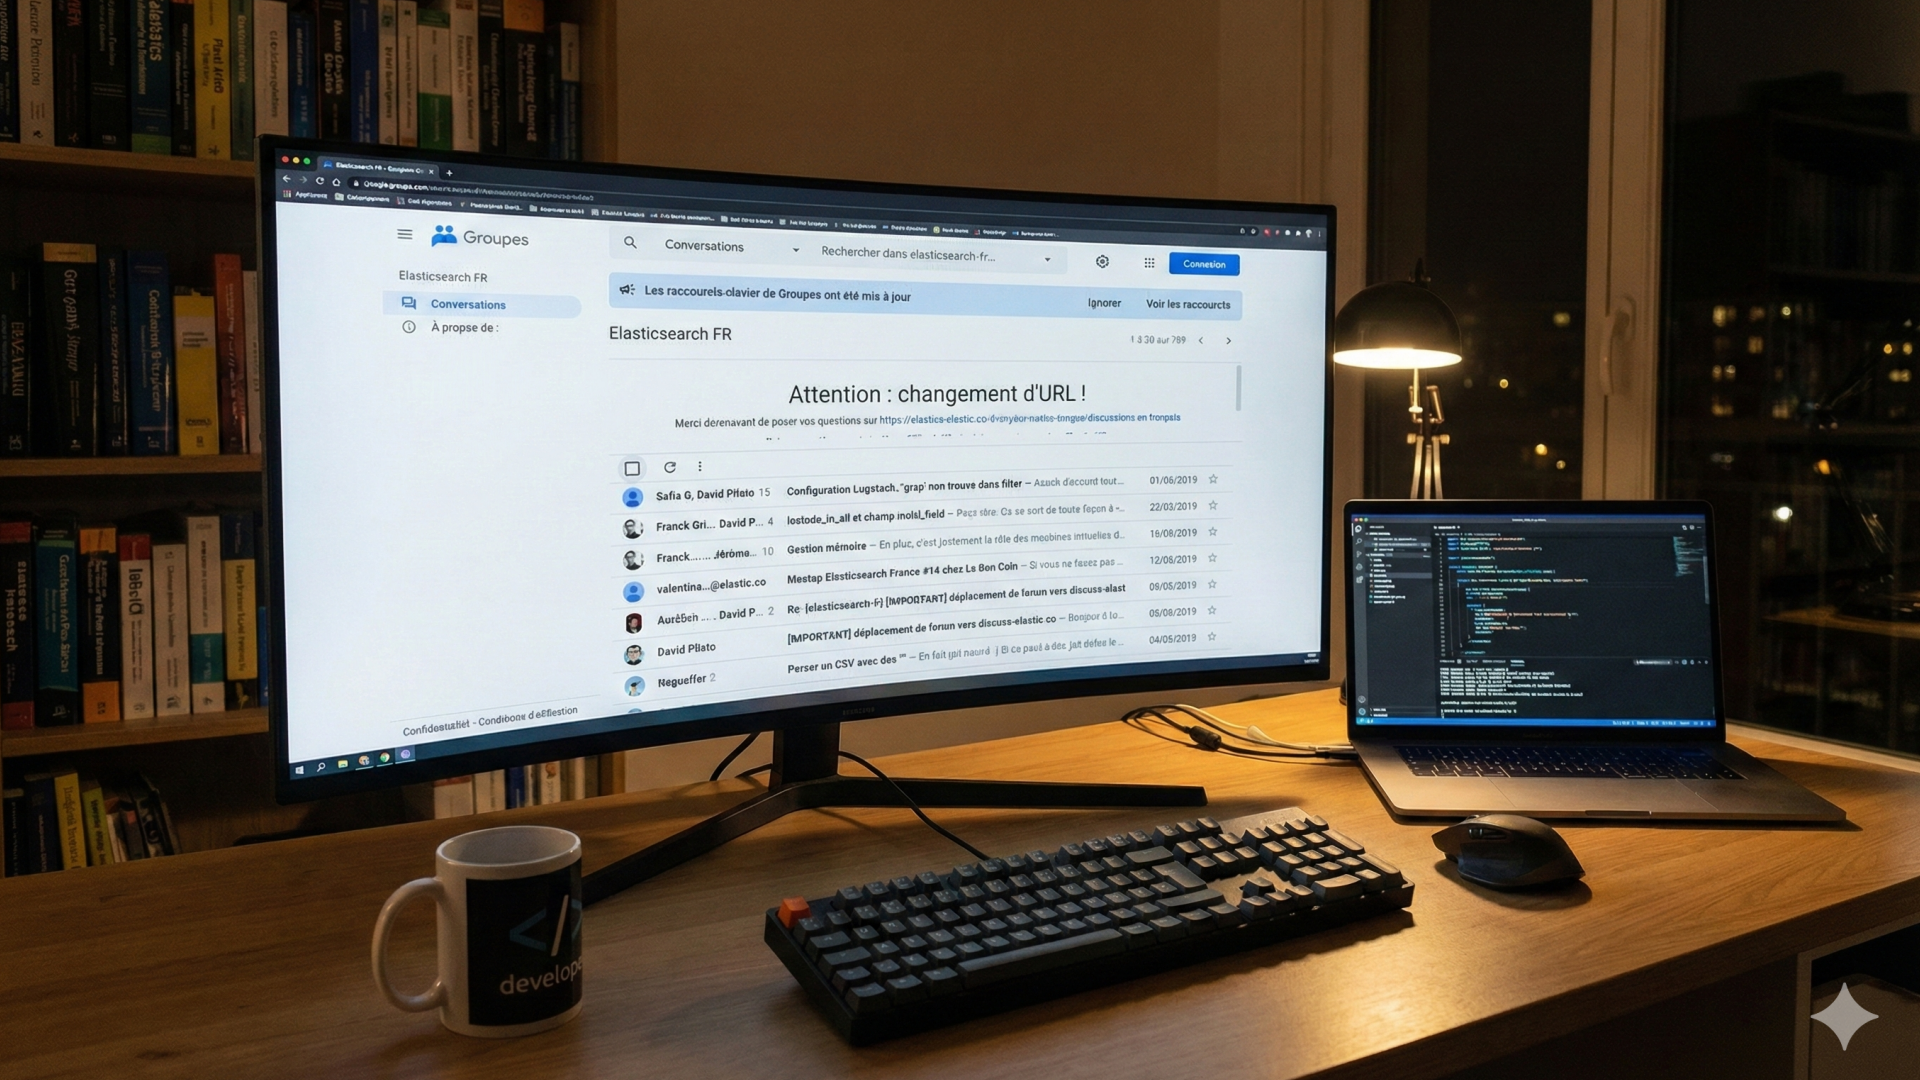
Groupes
Conversations
Rechercher dans elasticsearch-fr...
Connexion
Elasticsearch FR
Conversations
À propos de :
Les raccourcis clavier de Groupes ont été mis à jour
Ignorer
Voir les raccourcis
Elasticsearch FR
Attention : changement d'URL !
Merci dorénavant de poser vos questions sur https://elastics-elestic.co-dvsnyèor-natlss-tongue/discussions en tronpals
Safia G, David Pilato 15
Configuration Lugstach.."grap' non trouve dans filter — Azuck d'accord tout...
01/06/2019
Franck Gri... David P... 4
lostode_in_all et champ inolsl_field — Pacs sbre: Ca se sort de toute fepon à "...
22/03/2019
Franck..... Jérôme... 10
Gestion mémoire — En plus, c'est jostement la rôle des meobines intluelles d...
16/08/2019
valentina...@elastic.co
Mestap Elssticsearch France #14 chez Ls Bon Coin — Si vous ne faicez pas ...
12/08/2019
Aurélien ... David P... 2
Re: [elasticsearch-fr] [IMPORTANT] déplacement de forun vers discuss-alast
09/05/2019
David Pilato
[IMPORTANT] déplacement de forun vers discuss-elastic co — Bonjour à to...
05/08/2019
Negueffer 2
Perser un CSV avec des "" — En fait ipit naurd ; Et ce pauit à dec Jait défee le...
04/05/2019
develop


ANNONCE EXCLUSIVE :
Devoxx France 2012 –
Bientôt à Paris !
DEVOXX
PARIS JUG
PARIS JUG


DEVOXX
France
2012
Twitter River Demo


DOUANE
DOUANE
CUSTOMS
elasticsearch.
Elasticsearch Discussion Forum
Indexation issue
Cluster health
Query performance

elastic


DOUANE
CUSTOMS
elastic / elasticsearch
Docs: Fix typo in documentation about index settings #22
Write
Preview
This PR corrects a typo in the index settings documentation.
Fixes #12345.
Create pull request

elastic


PROBLÈME CRITIQUE : DÉPLOIEMENT ÉCHOUÉ – URGENT !!!
3:45 AM


Envie de booster ta carrière ? Open source-toi !
* Soyez visible
* N'hésitez pas
* Imposteur
* Contribuez:
- documentation
- tests
- code
- articles de blogs
* Aidez les autres

David Pilato
@dadoonet
@david.pilato.fr
elastic

Envie de booster ta carrière ?

Open source-toi !

DevFest Dijon 25