

EndecaCommunity.com

Everything you want to know about Endeca

Search

or

Ask a Question

Endeca Community Solutions

MDEX



Dear Endeca Community user,

On July 31, 2015 we closed EndecaCommunity.com, and asked users to instead take advantage of the collaboration tools and communities now provided by Oracle.

Why we did this? Because EndecaCommunity.com has fulfilled its purpose.

We created EndecaCommunity.com to fill a temporary void, to support teams working with Endeca technology in both Commerce and Data Insight capacities. The intention was to provide continuity through the transition period that followed Oracle's acquisition of Endeca, until Oracle had ramped up its own Endeca support community.

Now that Oracle's Endeca communities and formal user groups are well established, it's better for you as an Endeca user to find all the support you need from OAUG Special Interest Group and Oracle sources, rather than split your time having to cross-post to multiple disparate sites to find the answers you need.

We recommend joining the [Endeca Community on LinkedIn](#) or the [OAUG Endeca Special Interest Group](#). More support for Endeca can be found on the Oracle community website at [Technical Questions](#) and [Endeca Information Discovery](#).

On the following pages you will find the most frequently asked questions about MDEX on EndecaCommunity.com and the best solutions as determined by members of the Community.

Thank you for taking advantage of EndecaCommunity.com and for trusting RealDecoy as your Endeca authority.

Sincerely,

Richard Isaac
President & CEO, RealDecoy

Table of Contents

- **LiveDgraph** 4
- **How can I secure the MDEX/DGraph admin page?** 4
- **MDEX hosts out of sync** 5
- **My MDEX is not performing well what can I do to make it work faster?** 5
- **How do I perform performance testing on the MDEX?** 7

MDEX

LiveDgraph

Q: I have upgraded MDEX from 6.3.0 to 6.4.0 and after that everything is working fine and baseline indexing has also completed successfully. Now the endeca_jspref page is showing a weird issue. I am getting correct data on port 15002(authoring) in endeca_jspref but on port 15000(live) I'm getting the data which was indexed when MDEX was 6.3.0. I am totally stuck. Any guidance will be helpful.

A: Endeca_jsp ref web application still has 6.3 related navigation jar files. Update the jar files and restart Platformservices.

Validate jar versions

- 1) PlatformServices\<version>\lib\java
- 2) PlatformServices\<version>\tools\server\webapps\endeca_jspref\WEB-INF\lib
- 3)ToolsAndFrameworks_1\<version>\reference\endeca_jspref\WEB-INF\lib

How can I secure the MDEX/DGraph admin page?

Q: I'd like to restrict the DGraph admin page (available at [host]:[port]/admin) to only allow the read-only operations for monitoring purposes while disallowing the potentially service-impacting operations.

I would like to allow: audit, ping, stats, etc.

And disallow: exit, reload-services, restart, flush, etc.

How can this be achieved?

A: MDEX's should be on a secure network and not exposed.

For reference sake here's the Security Guide:

http://docs.oracle.com/cd/E28911_01/PlatformServices.612/pdf/SecurityGuide.pdf . It shows how to enable ssl for your MDEX and other Oracle Endeca components. You need to define trust. Either you trust someone to not bring down your network or you don't. Without any access to the admin functions there are DOS attacks--through issuing a lot of bad performing queries.

MDEX hosts out of sync

Q: I have 2 MDEX 6.1.1 hosts and they are currently out of sync with the indexes. For some reason one of them failed to update while the other was successful. Any idea how I can synch them up without a baseline update?

A: Are they both in the same dgraph group? Do you know why they failed? It's probably a good idea to check the logs and determine the issue before proceeding but if you don't have time to investigate:

QUICK FIX = If you just want to re-apply to all indexes without running a baseline:
<project_dir>\control\runcommand.sh DistributeIndexAndApply

You will dump your cache on the mdexs but it's really the fastest way to re-apply the indexes. If you want you could get more complicated, backup your config, remove the definitions for the working dgraphs, then distribute and apply, and then restore your configuration.

My MDEX is not performing well what can I do to make it work faster?

Q: My application is not performing up to where it needs to be. We've identified that the app isn't the problem, it's the length of time it takes to get a response from Endeca. What are some things I can do to improve performance?

A: Performance of a solution is the result of decisions throughout the solution lifecycle. The following describes some principles, steps, and techniques that can be used to help ensure performance is top-of-mind throughout the development cycle.

DGraph Server Statistics

Did you know that Endeca keeps tracks of some key information to help diagnose what is going on?

Try pointing your web browser at your application's dgraph and prepare to be amazed.

http://dgraph_host:dgraph_port/admin?op=stats

This will give you access to information such as:

- Most Expensive Queries
- Hot Spots
- Status of the Cache
- Summary statistics on queries and throughput
- Command line arguments passed to the dgraph process
- Next week's lottery numbers :)

Threads

The default configuration is not suitable for large scale production deployments. Attention needs to be given to the number of threads available to the Dgraph process. More isn't always better, and the following guidelines should be used as part of your design, testing, and deployment planning

Endeca Threads

When to increase the number of threads?

Endeca recommends using a higher setting for threads than in previous releases. Increasing the number of threads allows the MDEX Engine to handle more queries simultaneously. Use the following recommendations:

-If you are using an application with a low throughput without long-running queries, this implementation can run in a singlethreaded mode in which one thread is used to process all query requests to the MDEX Engine. The same thread is used for other query-related processes of the MDEX Engine.

If you are using a single MDEX Engine server with one thread, it is worth increasing the number of threads to improve performance. A simple recommendation is to configure at least one thread per core. Higher ratios may generate more throughput, but due to the potential impact on latencies, Endeca recommends running further testing to find the thread count most beneficial to the needs of a specific application. If increasing the number of threads stops improving query performance, this is an inflection point at which you can start considering the need to switch to a configuration with more Dgraphs.

A typical estimate that you can use to start testing with the increased number of threads is about 1 thread per core.

For example:

On a standard processor, enable 1 thread per processor.

On a dual-core processor, enable 2 threads per processor.

On a quad-core processor, enable 4 threads per processor.

Queries

Search query processing order

Understanding how Endeca will process the query goes a long way in understanding the likely impact of the query and resulting performance. From a conceptual perspective, the first thing Endeca tries to do is reduce the set of records through filtering. Then, it looks for path to expand the set through auto correction, stemming. Note that range filtering, business rules, analytics, and relevance ranking are at the end, and will be applied to each result set, to they constitute performance hits.

This section summarizes how the MDEX Engine processes record search queries. While this summary is not exhaustive, it covers the processing steps likely to occur in most application contexts. The process outlined here assumes that other features (such as spelling correction and thesaurus) are being used. The MDEX Engine uses the following high-level steps to process record search queries:

- Record filtering
- Endeca Query Language (EQL) filtering
- Tokenization
- Auto correction (spelling correction and automatic phrasing)
- Thesaurus expansion
- Stemming
- Primitive term and phrase lookup
- Did you mean
- Range filtering
- Navigation filtering
- Business rules and keyword redirects
- Analytics
- Relevance ranking

How do I perform performance testing on the MDEX?

Q: My application is running considerably slow based on the hardware and the data size. How do I go about running performance tests on the MDEX?

A: You can run ENERPerf on your MDEX to verify the performance of your MDEX engine, and further information on individual queries can be found using the Request Log Analyzer. Documentation on this can be found in the Oracle Endeca Performance Testing Guide. Additionally, the Cheetah tool can be used to provide further information on an MDEX Engine by analyzing the DGraph log files.