

Agenda - OpenSearch Search Meeting

OpenSearch Search Meetup - Standing Agenda

<https://www.meetup.com/opensearch/> >> "Search Backlog & Triage Meeting"

- Introduction (5 mins)
 - Host's Intro
 - What do we cover here?
 - New members Intro
 - Select a volunteer to write meeting notes below
- Triage New Search Issues -
 - **Label:untriaged on the Search Project Board:**
 - <https://github.com/orgs/opensearch-project/projects/45/views/21>
 - Recent Issues to discuss:
 - [Search Project Board](#) -> New -> Scroll at the end (move bottom up)
 - <https://github.com/orgs/opensearch-project/projects/45/views/21?filterQuery=label%3A%22good+first+issue%22>
 - Comment on an issue you want to work on and ask to be assigned.
 - Browse through Search component issues not tagged/labeled correctly.
- New Contributors Can Pick Up Issues -
 - **label:"good first issue"**
 - <https://github.com/orgs/opensearch-project/projects/45/views/34>
- Flaky Tests
 - <https://github.com/opensearch-project/OpenSearch/issues?q=is%3Aissue+is%3Aopen+flaky+test+label%3ASearch>
 - Aim to review this once in a while
 - 10/15/2024 - 23 count
 - 11/27/2024 - 11 count (closed up 12 issues manually opened since either they haven't been flaky in a while or already been tracked by autocut issues)
 - 12/23/2024 - 5 count (closed up flaky test issues which haven't been reported flaky in last 5+ months)
 - 03/26/2025 - 4 count
- Project Roadmap Review (one per week - host can help pick), ex -
 - Search Vertical Application : <https://github.com/orgs/opensearch-project/projects/45>
 - Performance : <https://github.com/orgs/opensearch-project/projects/153>

- OpenSearch Roadmap:
<https://github.com/orgs/opensearch-project/projects/206/views/17>
- Core Search & Query Capabilities
- Resiliency
- Varying Topics
 - Attendees put their name next to a topic below or add their own topic
 - Specific topics which were raised of interest from the previous meeting

Generally, in this meeting we cover all aspects of “Search”, including core search functionalities, query capabilities, resiliency, and relevance. We can also use the time to review the roadmap of various streams, and gather broader feedback on planning and releases.

- We welcome community discussion and contribution.
- We are all learning together, even the organizers of this meeting.
- Agendas may look full; it's a community meeting, so we can take a tangent.
- Please feel free to introduce yourself to the group, especially if you're speaking, but if you're just observing, no pressure.

Meeting Discussion Topics

Do you want to discuss a topic? Add yours here. We may start voting on topics as well so if you have something you really, really want to discuss, but a meeting looks full, add the topic anyway.

Date	Topic	Presenter	Presentation/Discussion Time	Recording OK
2026-05-27	OpenSearch API Spec: What's the policy regarding fixing api-spec bugs and downstream client library impact?	Sean Chittenden		
2026-05-13	OpenSearch API Spec: 1. api-specification#1103 2. api-specification#1102 3. api-specification#1101 4. api-specification#1099 5. api-specification#1098 6. api-specification#1107 7. api-specification#1108	Sean Chittenden		
2026-05-13	OpenSearch Deadlock: 4. OpenSearch#21404	Sean Chittenden		

	2. OpenSearch#21570 3. OpenSearch#21572			
2026-03-11	Conclude on https://github.com/opensearch-project/dashboards-search-relevance/issues/780 https://github.com/opensearch-project/search-relevance/issues/402	Prithvi, Venkateshwar		
2025-11-11	Migrating Index Schema for Plugin Between Versions	Diogo Pontual and Eric Pugh	15 minutes	yes
2025-10-08	Please have a look at https://github.com/opensearch-project/OpenSearch/pull/19350 Just a reminder, PR is obvious.	Mikhail Khudnev	I'm hardly able to speak during my commute.	yes
	Reducing build times and reducing flaky tests with the goal to reduce mean time to merge a PR	Sandesh Kumar	15 min	
?	HLD for Dynamic Weight Optimization for Hybrid Search	Martin?		
2025-08-06	Remote Query Capability RFC (for SRW)	Scott Stults	15 min	yes
2025-08-06	Discussion : IO optimizations in Search https://github.com/opensearch-project/OpenSearch/issues/1883	Saurabh Singh	15 min	yes
2025-07-30	Understand role and processes around app sec team when modifying API's	Eric Pugh (hopefully Martin Gaievski as well)	15 min maybe?	yes
2025-07-06	Discussion: Understand how Snapshots are built for OpenSearch and Plugins (search-relevance depends on user-behavior-insights for testing)	Eric Pugh	15 min	
TBD	https://github.com/opensearch-project/OpenSearch/issues/18608	Ankit Jain	10 min	
2025-06-25	Update: Team Draft Interleaving	David Fisher	10 min	Yes

	for A/B testing https://github.com/opensearch-project/OpenSearch/issues/18383			
2025-06-04	Eric Demo SRW Hybrid Search Optimizer Show n Tell	Eric Daniel/Martin	15 min 15 min	Yes
2025-06-04	System Indices, its limitations with SRW	Eric Pugh	15 min	
2025-05-28	Team Draft Interleaving for A/B testing https://github.com/opensearch-project/OpenSearch/issues/18383	David Fisher	15 min	Yes
	https://github.com/opensearch-project/opensearch-api-specification/issues/901	Scott Stults	5 min	
2025-05-21	Expanding the scope of this meeting			
2025-05-14				
2025-05-07	Learnings from OpenSearchCon 2025	Sandesh Kumar, Ankit Jain, Govind Kamat, etc	15 min	
2025-04-16	How do we trust results from the perspective of Search Relevance? (Predictability of Search Results)	Eric Pugh	15 min	
2025-03-12	How to effectively resolve relative time range when resolving using star-tree? https://github.com/opensearch-project/OpenSearch/issues/17443	Sandesh Kumar & Bharathwaj G	15 min	
2025-02-12	Discuss Lucene Changelog		15 min	
2024-11-20	Edge case of parsing terms over IP masks PR#16391	Mikhail Khludnev		
2024-11-13	Features to improve maintainability of LTR plugin	Johannes Peter	15 minutes	

	within managed services			
2024-11-06	IP field terms query: Max clauses exception .	Mikhal Khudnev	7.5 mins PR limit number of ranges only PR Doesn't support DV only fields Lucene PR SortedSet DV Query	Yes. Let me speak at the very end of the meeting, since I'm commuting.
2024-10-23	Hybrid Search Optimizer	Daniel Wrigley	15 minutes Slide deck: https://docs.google.com/presentation/d/1ehiXixvTtK0ByXYxGLKfdpoj6FyRZoUViA6Ooat0ral/edit?usp=sharing	yes
2024-10-23	IP field terms query: Max clauses exception .	<can't attend. Sorry>	PR limit number of ranges only PR Doesn't support DV only fields	Please leave a comment at those PRs
2024-09-04	Hello LTR! I'm rebooting OS LTR	Scott Stults	5 minutes	
2024-08-28	Search Quality Evaluation Framework RFC	Jeff Zemerick / Eric Pugh	20 minutes	yes
2024-08-07	Strategy for Implicit Judgements from UBI Tracking Data in OpenSearch	Daniel Wrigley	30 minutes	yes
2024-07-31	Where is UBI Headed.	Eric Pugh	10 minutes	yes
2024-07-10	https://github.com/opensearch-project/OpenSearch/issues/14637			
2024-07-03	Indexing and search separation	Shuyi / Philip	15 minutes	
2024-06-19	Discuss to get consensus on labeling support with headers	Kaushal / Chenyang	15 minutes	
2024-06-05	Labeling support for all types of queries - discuss approaches	Chenyang / Michael	10 minutes	
	Host Rotation for next week	Saurabh Singh	5 minutes	
2024-05-29	Progress on POC improvements in BKD Leaf block encodings https://github.com/opensearch-project/OpenSearch/issues/13686#issuecomment-2137643638	Aniketh Jain	10 minutes	

2024-05-22	Progress with POC improvements on Multi-term aggregation	Sandesh Kumar	10 minutes	
2024-05-01	Tenant labeling & POC	Chenyang / Michael	10 minutes	
2024-04-24	Concurrent Segment Search Performance Benchmarking	Ganesh Ramadurai	15 mins + QA	
2024-04-17	UBL no longer using Plugin Approach	Eric Pugh	10 minutes + 10 min Q&A	
2024-03-19	Maintainers in Training: https://github.com/opensearch-project/OpenSearch/issues/12779 https://github.com/opensearch-project/OpenSearch/pull/11650	Sandesh Kumar	15 minutes	
2024-03-13	Capture query level resource usage - discuss approaches	Chenyang Ji	10 minutes	
2024-03-06	User Behavior Insights/Query Insights Discussion	Chenyang/ Eric	20 minutes	
2024-02-21	Query sandboxing: https://github.com/opensearch-project/OpenSearch/issues/12342	Kiran Prakash	20 minutes presentation + 10 minutes Q&A	
2024-2-14	UBL Show n Tell	Eric Pugh	20 Minutes (walk through)	yes
2024-2-7	Query Insights and Top N Queries	Chenyang Ji	10 - 15 minutes	
2024-2-7	"UBL" Update (Eric) and RFC Feedback Discussion the (Jeff)	Eric Pugh / Jeff Zemerich	10 - 20 minutes?	
2023-1-31	Use Collector#setWeight to improve terms aggregation performance - discuss approach & take feedback	Sandesh Kumar	10 minutes	Don't think it's worth it
2024-1-24	Introduce User Behavior project	Eric Pugh	20 minutes	
2024-1-17	Brainstorm idea for expanding the Search Relevance meeting to cover all aspects of Search, including core functionalities, query capabilities, resiliency, and relevance.	Saurabh Singh	20 minutes	

2024-1-3	Good first issues 1. Just a reproducer for sorting by basic * missing dates	Mike Khl		Don't think it's worth it
2024-1-3	What about introducing join between indices like Solr ?	Mike Khl		Might be useful
2023-12-xx	Discussion: LTR in the Age of Vector Search/LLMs		45 minutes	Yes
2023-12-20	Phone Number Analyzer - https://github.com/opensearch-project/OpenSearch/issues/11326	?	?	
2023-12-13	Continue LTR Discussion - Connectors to External Models/Performance & Feature Logging Options; Any volunteers to work on a 2.10 open source release? Anyone want to volunteer to test the 2.9 build	Mark	?	
2023-12-13	Phone Number Analyzer: https://github.com/opensearch-project/OpenSearch/issues/11326	Mark		
2023-12-06	LTR plugin digest	Louis Chu	15 minutes	
2023-11-29	OSCI Dashboard Search Relevance Contributions	Nicholas Ung	10 minutes	
2023-11-08	BM25F - Add Support combined fields in OpenSearch	TBD	10-15 minutes	
2023-11-15	Hello LTR on OpenSearch Demo/Walkthrough	Mark Cohen	15-20 minutes	Yes
2023-11-15	How do people actually use LTR for shard-level reranking?			
2023-11-15	OSCI Dashboard Search Relevance Contributions	Nicholas Ung	10-15 minutes	

Meeting Notes

2024-02-14

- Action item: Discuss commit messages on #maintainers channel on OpenSearch Slack.
- Action item: Fix search links in the agenda

- Action item: Follow up on aggregation test failures
- Question: Do we want hooks for stale Good First Issues to unassign? Ankit J will look into it.

2024-02-07

- Eric Pugh discussed user behavior logging RFC (<https://github.com/opensearch-project/OpenSearch/issues/12084>). Will present the early version next week (IIRC). Discussed some options around client-side libraries for behavior logging, looking to reuse an existing standard if possible. Example client libraries include [search-collector](#), [Apache Flagon](#), Google Analytics gtags ([example usage](#) from OpenSearch documentation website).
- Chenyang Ji discussed [Query Insights Framework](#) and demoed the functionality coming in OpenSearch 2.12. The feature in 2.12 will be limited to in-memory storage and retrieval of top N queries by latency for the current time window only. Future work will add retention of windowed data, a dashboards UI, and ability to collect top N by other criteria (e.g. CPU utilization)
 - Andriy Redko asked why we need to implement this as a separate solution, when all this data will be available via the tracing framework. In theory, the two solutions can be complementary, as Query Insights can be modified to use/implement an OpenTelemetry collector (to watch traces as they're emitted), rather than an explicit request listener.

2024-01-10

I think we have a few topics for today's Search Relevance meeting: there are two issues I missed that Mike K added to the agenda for last week. We're making some progress on a 2.10 LTR plugin, thanks to Mike K and @Austin Lee for input. Implementing the `isFragment()` method. There's also this issue that I think is interesting to discuss and raise for Search Pipelines: [\[Feature Request\] Remote Search Processor](#).

- Triage
 - [\[BUG\] exists query for flat object ignores field paths](#)
 - Not sure if this is definitely a bug
 - [\[BUG\] Query Workbench: Semicolon in SQL strings terminates query](#)
 - Requested admin to move to dashboards-query-workbench repo
 - [Logic duplication in SearchRequestContext and TimeProvider](#)
 -
 - [\[Feature Request\] provide dedicated API to get all stored search templates](#)
 - Consider working on. Anyone
 - [\[Feature Request\] Conditional Search Processor](#)
 - Closed two other issues as duplicate
 - [\[Feature Request\] Remote Search Processor](#)
 - [\[BUG\] Exists query fails if object has "search as you type" field](#)

- Bug confirmed; labeled “good first issue”
 - [\[Feature Request\] we should print detail reason why the parent task cancelled](#)
 - Labeled other; no good label is available for this...
 - [\[Feature Request\] Tie breaker for search with search after pagination](#)
- LTR Plugin Build for 2.10 (Thanks to Mike K and Austin for feedback on
 - Mike K submitted a PR to override isFragment() which passes all the tests, but not yet sure why...
<https://github.com/macohen/opensearch-learning-to-rank-base/blob/52e17a6095db806be78c48234521e9b8ad89f4a8/src/main/java/com/o19s/es/ltr/feature/store/StoredLtrModel.java#L201>
 - Review: <https://github.com/macohen/opensearch-learning-to-rank-base/pull/5>
 -
 -

2023-12-13

- Triage
- Deepak Sharma, Mikhail Khludnev, Andriy Redko, Michael Froh
- Deepak asks about deploying a huggingface model into OpenSearch
 - Michael Froh suggests the OpenSearch Slack workspace
- OpenSearch LTR
 - Remote Inference
[\(https://opensearch.org/docs/latest/ml-commons-plugin/remote-models/index/\)](https://opensearch.org/docs/latest/ml-commons-plugin/remote-models/index/)
 will downgrade performance, but can have better results...
 - Integration is relatively straightforward.
 - LTR accesses API over https instead of
 - Or LTR can work at the shard-level and remote models can be accessed at a later stage
 - Is there a difference between shard-level and coordinator level reranking?
 - Good question...
 - Depends...
 - Proof of Concept to hook up a remote model to LTR
 - Reranking at the model happens during query phase
 -
- Continue Conversation on Learning to Rank
 -
 - Remote Models:
<https://opensearch.org/docs/latest/ml-commons-plugin/remote-models/index/>
 - Can we do a deeper dive on feature logging?
 - What would async feature logging look like?
 -

2023-12-06

- Triage: 1 untriaged issue
- Eric Pugh, Minal Shah, Louis Chu, Michael Froh, Andriy Redko
- Reviewed
 - <https://github.com/opensearch-project/opensearch-learning-to-rank-base/issues/26>
- Discussed the difference between LTR as a concept and the LTR plugin
 - Video:  Search Relevance: Learning-to-Rank Discussion (2023-1206)
 - LTR is about learning from user behavior (or judgment) how search should change, providing the most likely set of results that will satisfy that user.
 - LTR, the plugin, needs updating
 - Support larger models
 - Can be done with
 - Asynchronous feature logging
 - What else?
 -

2023-11-01


Agenda

- Introductions
- Triage New Issues - <https://github.com/orgs/opensearch-project/projects/45>
- New Contributors
 - <https://github.com/orgs/opensearch-project/projects/45/views/1?filterQuery=label%3A%22good+first+issue%22>
- BM25F - [Add Support combined fields in OpenSearch](#)
- LTR
 - Overview and what should we do with LTR
 - Aryn uses cross-encoding for reranking;
 - Judgement Lists can be useful outside of LTR as well; feature engineering
 - Search Pipelines?
 - Reranking on individual shards that LTR does now would be prohibitive

2023-10-25

Agenda

- Introductions
- Triage New Issues - <https://github.com/orgs/opensearch-project/projects/45>
- New Contributors

- <https://github.com/orgs/opensearch-project/projects/45/views/1?filterQuery=label%3A%22good+first+issue%22>
 - Discussion & Discovery
 - Measuring Relevance (precision, recall, accuracy)
 - [Ranking evaluation - OpenSearch documentation](#)
 - What else do we have in OpenSearch?
 - Louis: there are 5 evaluation metrics provided by OpenSearch
- 

A screenshot of a code editor showing five evaluation metrics from the OpenSearch modules. Each metric is listed with its fully qualified class name in parentheses and the module path. The metrics are: DiscountedCumulativeGain, ExpectedReciprocalRank, MeanReciprocalRank, PrecisionAtK, and RecallAtK. All of them are associated with the module path OpenSearch.modules.rank-eval.main.

Metric	Module Path
DiscountedCumulativeGain	OpenSearch.modules.rank-eval.main
ExpectedReciprocalRank	OpenSearch.modules.rank-eval.main
MeanReciprocalRank	OpenSearch.modules.rank-eval.main
PrecisionAtK	OpenSearch.modules.rank-eval.main
RecallAtK	OpenSearch.modules.rank-eval.main
- Other tools? RRE?
 - Learning-to-Rank: Current/Future State
 - Opensearch-learning-to-rank-base:
<https://github.com/opensearch-project/opensearch-learning-to-rank-base>
 - Getting it to build for latest versions as a self-install plugin
 - 2.6
 - Should LTR be baked into the OpenSearch Release?
 - Open Sourcing Private LTR Repo Features
 - Circuit breaker
 - Other
 - Divergence from Elasticsearch source of this fork
 - BM25F - [Add Support combined fields in OpenSearch](#)
 -
- Other Topics?
 -  OpenSearchCon 2023 - Opening Keynote
 - Other Interesting Search Relevance Presentations
 - <https://simonwillison.net/2023/Oct/23/embeddings/> (credit to Doug Turnbull for posting this on LinkedIn)
- Next Time

- LTR
 - Can we refactor the LTR plugin back into the elasticsearch plugin?
 - How do we keep the divergence of OpenSearch and Elasticsearch separate from the implementation of LTR?
 -
- Relevance Ranking Tools
 - <https://dl.acm.org/doi/10.1145/1645953.1646033>
 - [Ranking evaluation - OpenSearch documentation](#) - this is what we have today.
- NodeAnalyzers
 - Add optional section of node analyzers into NodeInfo
 - Cannot see what analyzers are available on the cluster.
 - This solves that problem

- <https://github.com/opensearch-project/OpenSearch/pull/10894>