Workstation Enterprise Search

Brief Overview

Workstation Enterprise Search enables users to discover applications and resources with a single unified search. It searches across WalkMe content, and content within 3rd-party apps that were connected in the Settings Integrations page.

It allows to:

- Provide fast and efficient knowledge discovery, eliminating the need to search across multiple disparate data sources
- Generate personalized, AI-powered results where users can filter by app and file type
- Preserve business security and end-user privacy, with zero indexing and respecting access permissions

Sorting search results

The Enterprise Search uses a **Sorter Service** that sorts the results according to its relevance. How the sorter service performs the sorting-

1. A user searches for a term in the Workstation app
2. Search service gets the term and calls for each connected app's service
3. Connected app's service search for this term, each with its own API call and implementation
4. Search service gets all the search results from all the connected apps
5. Search service calls the sorter service with the results
6. Sorter service sorts/orders the results and sends the sorted/ordered results back to the Search service
7. Search service returns the results back to the Workstation app

The Sorter service uses 8 different scoring methods, including stemming, fuzziness and NLP, in addition to the basic Levenshtein distance. Recently Viewed items get a higher prioritization in search results as well, since they're most likely to be more relevant to the user during the search process.

Indexing

Workstation does not index results. Some data is stored on the client side. For example, the data which appears on Recent Search Results or in the Recently Viewed widget.

Rest of the data displayed to the user on the Search flow is brought in real-time from the 3rd-party apps API's.

Integrations

The Enterprise Search uses 3rd-party integrations to implement a "federal search". Searches within Workstation are backed by an NLP engine, and a graph database that supports a great user experience. Workstation Enterprise Search doesn't index 3rd-party data on an independently searchable database. See below sequence diagram describes the searching algorithm:

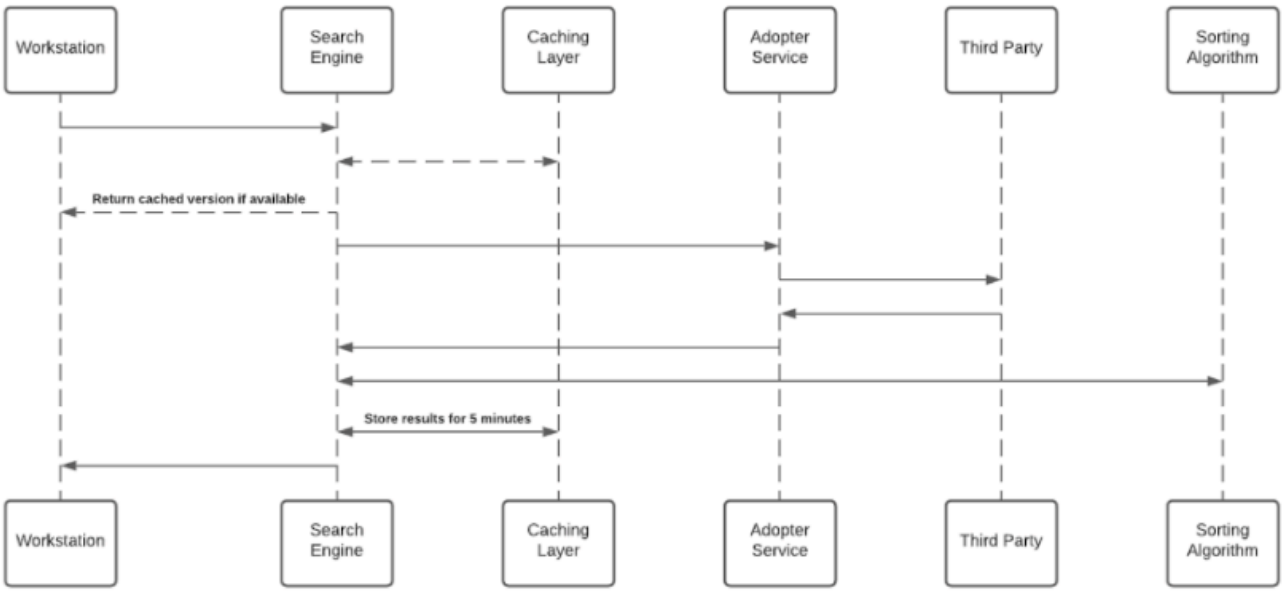

1. Cache layer saves results for a period of five minutes
2. Each Adopter Service creates a unique identifier for the results which is meaningless without access to the 3rd-party and stores it in the graph database

3rd-Party Access and Refresh Tokens

To activate the Enterprise Search (and the Personalized Workspace widgets), each employee is required to grant Workstation permission to access the 3rd-party. The granting process is using the OAuth2.0 protocol. Each time a new access token is granted to the Workstation, the application will encrypt the access and refresh tokens and store it in a remote database.

The encryption process includes a unique private key ("salt") that is generated for each individual at the very first bootstrap and stored in the local machine Keychain. The salt is irreplaceable and not restorable – losing it causes the access tokens to be voided. This security measure is being taken to eliminate identity spoofing when accessing high-sensitive data.

See the diagram below to review the salt generation and storage flow.

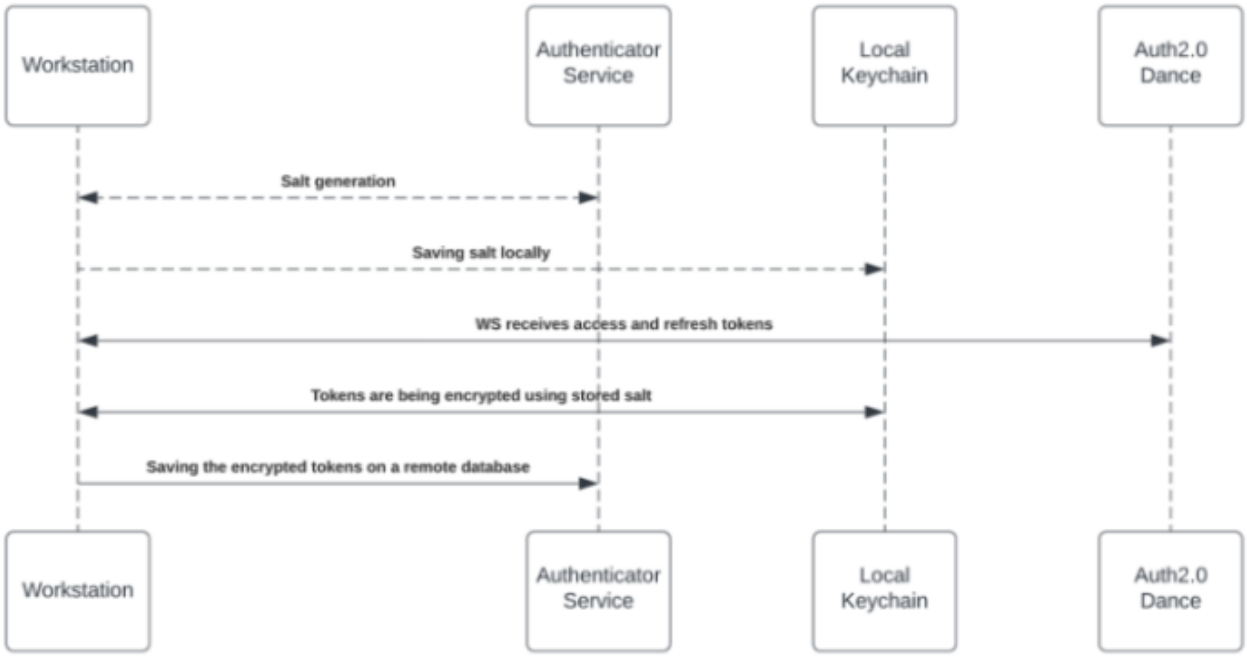

Accessing 3rd-Party Content

Accessing 3rd-party content requires user consent, and in some cases, mostly on Microsoft products, an organization admin consent. Users grant Workstation the necessary permission by approving an OAuth2.0 consent screen that is being triggered by them from the Workstation application (“Third-party apps”).

The third-party apps are being approved and verified by third-parties products. By the end of the granting process, the third-party apps provide access and refresh tokens that are used by the search engine to establish the requests.

See 3rd-Party Access and Refresh Tokens section above for more information about the storing mechanism.

While searching, the search engine forwards the request, before hitting the Adopter Service, through the Token Injector; a service that injects the relevant tokens to accomplish the request. The local private key is being handed off over the search HTTPS request for runtime decryption.

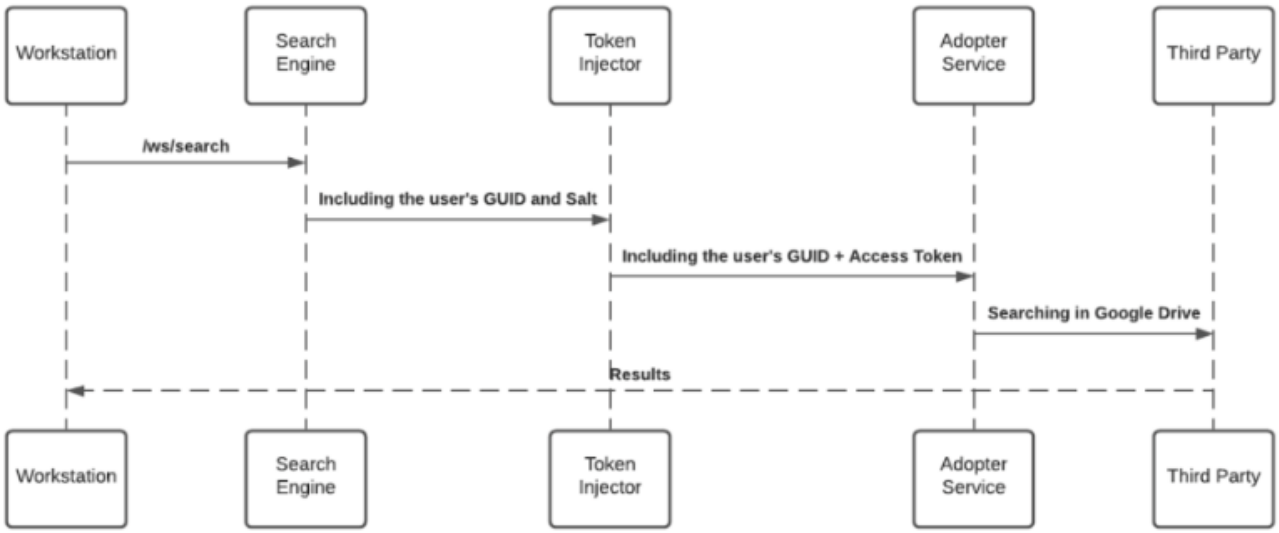

JWT Protection

When a user initiates a search query – the WalkMe enterprise search starts a search flow that is being protected by a JWT assigned by WalkMe IdP integration, as part of the user signing flow:

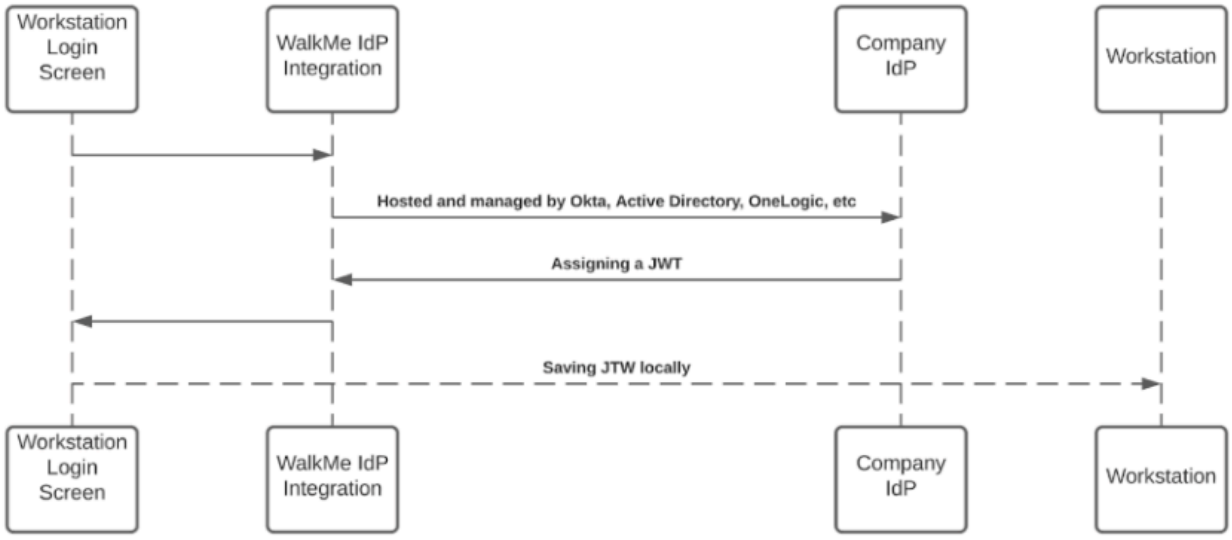

The JWT is proxying the user identity and keeping any HTTPS request secured and individual.

All Workstation requests are protected by a JWT validation.