



Red Hat OpenShift Lightspeed 1.0tp1

About

Introduction to OpenShift Lightspeed

Red Hat OpenShift Lightspeed 1.0tp1 About

Introduction to OpenShift Lightspeed

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides an overview of OpenShift Lightspeed features.

Table of Contents

CHAPTER 1. ABOUT OPENSIFT LIGHTSPEED	3
1.1. OPENSIFT LIGHTSPEED OVERVIEW	3
1.2. OPENSIFT REQUIREMENTS	3
1.3. LARGE LANGUAGE MODEL (LLM) REQUIREMENTS	3
1.3.1. About OpenAI	3
1.3.2. About Azure OpenAI	3
1.3.3. About WatsonX	4
1.4. ABOUT DATA USE	4
1.5. ABOUT DATA, TELEMETRY, TRANSCRIPT, AND FEEDBACK COLLECTION	4
1.6. ABOUT REMOTE HEALTH MONITORING	4
1.6.1. Transcript collection overview	5
1.6.2. Feedback collection overview	5

CHAPTER 1. ABOUT OPENSIFT LIGHTSPEED

The following topics provide an overview of Red Hat OpenShift Lightspeed and discuss functional requirements.

1.1. OPENSIFT LIGHTSPEED OVERVIEW

Red Hat OpenShift Lightspeed is a generative AI-powered virtual assistant for OpenShift Container Platform. Lightspeed functionality uses a natural-language interface in the OpenShift web console.

This early access program exists so that customers can provide feedback on the user experience, features and capabilities, issues encountered, and any other aspects of the product so that Lightspeed can become more aligned with your needs when it is released and made generally available.

1.2. OPENSIFT REQUIREMENTS

Openshift Lightspeed requires OpenShift Container Platform 4.15 or later running on x86 hardware. Any installation type or deployment architecture is supported so long as the cluster is 4.15+ and x86-based.

For the Openshift Lightspeed Developer Preview release, the cluster you use must be connected to the Internet and it must have telemetry enabled. Telemetry is enabled by default. If you are using a standard installation process for OpenShift confirm that it does not disable telemetry.

1.3. LARGE LANGUAGE MODEL (LLM) REQUIREMENTS

As part of the Developer Preview release, Openshift Lightspeed relies on Software as a Service (SaaS) LLM providers. You will need to have either a trial or paid subscription that allows for API access to completions and inferences with one of the following providers:

- OpenAI
- Azure OpenAI
- WatsonX



NOTE

Many self-hosted or self-managed model servers claim API compatibility with OpenAI. It is possible to configure the OpenShift Lightspeed OpenAI provider to point to an API-compatible model server. If the model server is truly API-compatible, especially with respect to authentication, then it may work. These configurations have not been tested by Red Hat, and issues related to their use are outside the scope of Developer Preview support.

1.3.1. About OpenAI

To use OpenAI with Red Hat OpenShift Lightspeed, you will need access to the OpenAI [API platform](#).

1.3.2. About Azure OpenAI

To use Azure OpenAI with Red Hat OpenShift Lightspeed, you must have access to Azure OpenAI. If you have already applied for and received access to Azure OpenAI, you can ignore the rest of the information in this topic.

Azure OpenAI access currently requires an approval process initiated from within the Microsoft Azure portal. When you create an Azure OpenAI instance, you may see a warning that contains a link. The link in the warning will take you to the [signup form](#). It can take several days for Microsoft to approve your access. When your access to Azure OpenAI is approved, make sure you know which Azure subscription your Azure OpenAI access is associated with.

1.3.3. About WatsonX

To use WatsonX with Red Hat OpenShift Lightspeed, you will need an account with [IBM Cloud's WatsonX](#).

1.4. ABOUT DATA USE

Red Hat OpenShift Lightspeed is a virtual assistant you interact with using natural language. Using the Openshift Lightspeed interface, you send chat messages that Openshift Lightspeed transforms and sends to the Large Language Model (LLM) provider you have configured for your environment. These messages can contain information about your cluster, cluster resources, or other aspects of your environment.

The Openshift Lightspeed Developer Preview release has limited capabilities to filter or redact the information you provide to the LLM. Do not enter information into the Openshift Lightspeed interface that you do not want to send to the LLM provider.

By using the Openshift Lightspeed as part of the Developer Preview release, you agree that Red Hat may use all of the messages that you exchange with the LLM provider for any purpose. The transcript recording data uses the Red Hat Insights system's back-end, and is subject to the same access restrictions and other security policies.

You may email [Red Hat](#) and request that your data be deleted at the end of the Developer Preview release period.

1.5. ABOUT DATA, TELEMETRY, TRANSCRIPT, AND FEEDBACK COLLECTION

Openshift Lightspeed is a virtual assistant that you interact with using natural language. Communicating with Openshift Lightspeed involves sending chat messages, which may include information about your cluster, your cluster resources, or other aspects of your environment. These messages are sent to Openshift Lightspeed, potentially with some content filtered or redacted, and then sent to the LLM provider that you have configured.

Do not enter any information into the Openshift Lightspeed user interface that you do not want sent to the LLM provider.

The transcript recording data uses the Red Hat Insights system back-end and is subject to the same access restrictions and other security policies described in [Red Hat Insights data and application security](#).

1.6. ABOUT REMOTE HEALTH MONITORING

Red Hat records basic information using the Telemeter Client and the Insights Operator, which is generally referred to as Remote Health Monitoring in OpenShift clusters. The OpenShift documentation for remote health monitoring explains data collection and includes instructions for opting out. If you wish

to disable transcript or feedback collection, you must follow the procedure for opting out of remote health monitoring. For more information, see "About remote health monitoring" in the OpenShift Container Platform documentation.

1.6.1. Transcript collection overview

Openshift Lightspeed temporarily logs and stores complete transcripts of conversations that users have with the virtual assistant. This includes:

- Queries from the user.
- The complete message sent to the configured LLM provider, which includes system instructions, referenced documentation, and the user question.
- The complete response from the LLM provider.

Transcripts are associated with the cluster from which they originated, and Red Hat can attribute specific clusters to specific customer accounts. Transcripts do not contain any information about which user submitted a question, and cannot be tied to any individual user.

Transcripts are sent to Red Hat every two hours, by default. If you are using the filtering and redaction functionality, the filtered or redacted content is sent to Red Hat. Red Hat will not see the original non-redacted content, and the redaction takes place before any content is captured in logs.

1.6.2. Feedback collection overview

Openshift Lightspeed collects feedback from users who engage with the feedback feature in the virtual assistant interface. If a user submits feedback, the feedback score (thumbs up or down), text feedback (if entered), the user query, and the LLM provider response are stored and sent to Red Hat on the same schedule as transcript collection. If you are using the filtering and redaction functionality, the filtered or redacted content is sent to Red Hat. Red Hat will not see the original non-redacted content, and the redaction takes place before any content is captured in logs.

Feedback is associated with the cluster from which it originated, and Red Hat can attribute specific clusters to specific customer accounts. Feedback does not contain any information about which user submitted the feedback, and feedback cannot be tied to any individual user.