

Kafka

Configure XpoLog to collect Kafka topics:

1. Background -

In case Kafka listener does not appear in the listeners management console. Follow these steps:

- a. Stop XpoLog service.
 - b. Enter the file system with XpoLog user and enter the path `/INSTALL_DIR/plugins/templates/` then copy the directory **Kafka** and its content and place it at the path `/SHARED_CONF/plugins/user/`
 - c. Restart XpoLog.
2. Go to Manager > left panel Data > Listen to Data. The Listeners management console opens. Kafka listener should be visible.
3. Add Kafka account, for each account the following should be configured:
- a. **Name**: the name of the Listener account
 - b. **Description**: the description of the Listener account
 - c. **Listening Node**: the XpoLog node that will perform the subscription (display if XpoLog runs in cluster mode)
 - d. **Host**: the Kafka server(s):port(s) that the listener receives topics from (semicolon separated). For example: **SRV1:PORT;SRV2:PORT**
 - e. **Topics**: the list of topics Kafka Listener will subscribe to (comma separated). For example: **TP1,TP2,TP3**
You can also use a regular expression that represents all topics that the Kafka listener will subscribe to

f. Advanced Settings:

General Information:

Enabled: determine whether this account is enabled or not

Security Details (optional if Kafka requires authentication):

User Name: the user name to use to subscribe to the Kafka topic(s)

Password: the user name's password

Auth Mechanism: PLAIN

Security Protocol: SASL_PLAINTEXT

Security Module: org.apache.kafka.common.security.plain.PlainLoginModule

Advanced Kafka Details:

Consumer Group Id: use xplg

Dynamic Log Creation Configuration:

Parent Folder: the parent folder which all logs from this listener will be place under in XpoLog Folders and Logs tree

Collection Policy: the [collection policy](#) which will be associated to all logs from this listener (used mainly for storage location and data retention)

AppTags: the [AppTags](#) which will be associated to all the logs from this listener (used mainly for data enrichment)

Log Name Prefix: a prefix which will be added to any of the the logs from this listener (used to easily distinguish between multiple listener accounts logs)

Split by Source Device:

- i. Do not split - by default, XpoLog will not split the incoming data. All data will be stored under a single log in XpoLog.
- ii. Create log by unique IP / host name - XpoLog will split the incoming data based on the source that sends it to different logs, the log name structure will be "Log_Name_Prefix Source_IP/Name"
- iii. Create log by IP mask - XpoLog will split the incoming data based on matched source to the configured IP mask that sends it to different logs, the log name structure will be "Log_Name_Prefix IP_Mask"
- iv. Create log with Regular expression - XpoLog will split the incoming data to different logs based on a regular expression that will be applied on the Message field. The part in the regular expression that will be used to determine the split should be in (round parentheses).

Records that the regular expression does not return a value will be directed to the global Kafka log.

Listener Data:

Listener Data Location: the location which data will be stored to, by default XpoLog stores it in its data directory

Indexing Interval: the frequency in which received Kafka topics are indexed

4. Save the account.
5. Data received to the Kafka listener account will be created under the configured parent folder and will be available for searching, reporting and alerting.

Note: multiple listeners accounts may be configured. However, listeners which run on the same machine must listen on a different network port.