

---

# **IBM DataPower Add-on for Splunk Documentation**

***Release 0.1.0***

**Diogo Silva**

**Dec 23, 2019**



---

## Contents:

---

|          |                           |           |
|----------|---------------------------|-----------|
| <b>1</b> | <b>Introduction</b>       | <b>1</b>  |
| <b>2</b> | <b>Requirements</b>       | <b>3</b>  |
| <b>3</b> | <b>Installation</b>       | <b>5</b>  |
| <b>4</b> | <b>Configuration</b>      | <b>7</b>  |
| <b>5</b> | <b>Troubleshooting</b>    | <b>9</b>  |
| <b>6</b> | <b>Support</b>            | <b>11</b> |
| <b>7</b> | <b>Indices and tables</b> | <b>13</b> |



# CHAPTER 1

---

## Introduction

---



## CHAPTER 2

---

### Requirements

---

- Splunk 7.0 or newer
- IBM DataPower 7.5 or newer



## 3.1 Install the IBM DataPower Add-on for Splunk

- Get the IBM DataPower Add-on for Splunk by downloading it from [Splunkbase](#) or browsing to it using the app browser within Splunk Web.
- Determine where and how to install this add-on in your deployment, using the tables on this page.
- Perform any prerequisite steps before installing, if required and specified in the tables below.
- Complete your installation.

### 3.1.1 Distributed deployments

Reference the tables below to determine where and how to install this add-on in a distributed deployment of Splunk Enterprise or any deployment for which you are using forwarders to get your data in. Depending on your environment, your preferences, and the requirements of the add-on, you may need to install the add-on in multiple places.

#### Where to install this add-on

Unless otherwise noted, all supported add-ons can be safely installed to all tiers of a distributed Splunk platform deployment. See [Where to install Splunk add-ons](#) in Splunk Add-ons for more information.

This table provides a reference for installing this specific add-on to a distributed deployment of Splunk Enterprise.

| Splunk platform component | Supported | Required | Comments  |
|---------------------------|-----------|----------|---|
| Search Heads              | Yes       | Yes      | Install this add-on to all search heads.  |
| Indexers                  | Yes       | Optional | Required for the parsing operations (sourcetype renaming) if the data is not coming from a heavy forwarder. |
| Heavy Forwarders          | Yes       | Yes      | Required for the parsing operations (sourcetype renaming).  |
| Universal Forwarders      | No        | No       | This add-on requires heavy forwarders.  |

### Distributed deployment compatibility

This table provides a quick reference for the compatibility of this add-on with Splunk distributed deployment features.

| Distributed deployment feature | Supported | Comments  |
|--------------------------------|-----------|---|
| Search Head Clusters           | Yes       | You can install this add-on on a search head cluster for all search-time functionality. |
| Indexer Clusters               | Yes       |   |
| Deployment Server              | Yes       | Supported for deploying via Deployment server   |

### 3.1.2 Installation walkthroughs

The Splunk Add-Ons manual includes an [Installing add-ons](#) guide that helps you successfully install any add-on to your Splunk platform. For a walkthrough of the installation procedure, follow the link that matches your deployment scenario:

- [Single-instance Splunk Enterprise](#)
- [Distributed Splunk Enterprise](#)
- [Splunk Cloud](#)

### 4.1 Splunk

- Configure a new index (e.g. storage) for the new logs

The IBM DataPower Add-on contains two base sourcetypes: - `ibm:datapower:syslog` - this should be used if you are sending data via UDP - `ibm:datapower:syslog:tcp` - this should be used if you are sending data via TCP

The reason behind having multiple base sourcetypes is due to the fact that DataPower logs different timestamp formats depending on how you are sending the logs. - Sending data via UDP doesn't allow for much configuration and the timestamp will look something like "Jul 10 10:45:32". - Sending data via TCP allows for extra time granularity since you can choose to include the microseconds and time zone. It will look something like "2019-07-10T10:45:32.123415+01:00".

#### 4.1.1 Receiving syslogs on Splunk

NOTE: Its recommended to use a separate and dedicated syslog solution (e.g. rsyslog, syslog-ng, etc) - Configure new TCP port (e.g. 514) pointing to the new index using the "`ibm:datapower:syslog:tcp`" sourcetype

#### 4.1.2 Monitoring log files

- Configure a new file monitor input pointing to the new index using the "`ibm:datapower:syslog:tcp`" sourcetype

### 4.2 IBM DataPower

- Configure syslog outputs

For more information please refer to the [IBM DataPower documentation](#).



## CHAPTER 5

---

### Troubleshooting

---



### 6.1 Bugs & Support Issues

You can file bug reports on our [GitHub issue tracker](#) and they will be addressed as soon as possible. **Support is a volunteer effort** and there is no guaranteed response time.



## CHAPTER 7

---

### Indices and tables

---

- `genindex`
- `modindex`
- `search`