



X-PACK: MACHINE LEARNING

Overview

Elastic Machine Learning lets you get the most out of your data. Hunting down irregularities in your data is difficult and this on-demand course will teach you how to make this easy. In this course you will learn what to look for, how to look for it and why it is important. With 2.5 hours of video content and multiple labs, you will explore different datasets with the Machine Learning user interface to reveal the anomalies in them. You will also learn how Machine Learning integrates with Elasticsearch and how you can automate the processing of discovered anomalies so you can take action quickly. Elastic Machine Learning helps you discover irregularities before they become problems and this course teaches you how.

This course is based on Elastic Stack 5.x.

Audience

Anybody who would want to use X-Pack Machine Learning to discover anomalies in their data and create automation of Machine learning jobs.

Duration

With nearly 2.5 hours of instructional video, 4 labs and 30 quizzes we expect participants to allocate between 6-8 hours to complete this course.

Language

English

Prerequisites

Some experience with Elasticsearch and Kibana is helpful and we recommend you complete at least one of the following:

- Elasticsearch Engineer I
- Elastic Stack Data Administration I

This course uses Logstash to load datasets, but all the configurations are written for you.

Requirements

- Participants must use their own desktop or laptop system
- Internet connection capable of streaming video
- Mac, Linux OS, or Windows 7 or later
- Java version 1.8u20 or later installed
Note: Logstash 5.x does not support Java 9
- A modern web browser
- At least 20% free disk space

Additional Information

All training materials are provided via the Elastic training portal at <https://training.elastic.co>.

X-PACK: MACHINE LEARNING

Modules

X-Pack: ML Introduction

- X-Pack: ML Overview
- Statistical Learning Models & Anomalies
- Bucket Spans & Analysis Functions Overview
- Machine Learning Jobs & Data Feeds
- Creating A Single Metric Job
- Creating A Distinct Count Job
- Installing X-Pack

X-Pack: ML Concepts

- Investigating Anomalies
- Influencers Overview
- Creating A Multi-Metric Job
- Analysis with The Explorer
- Splitting Analysis
- Using Partition Fields to Identify Anomalies

Functions and Detectors

- Individual VS Population Analysis
- By Fields VS Over Fields
- Partition Fields VS By Fields
- Creating An Advanced Job
- Analyzing Script Attacks
- Automating Machine Learning

Architecture and Operations

- Production Architecture
- Distributed Persistent Tasks
- Controlling Data Feeds
- Machine Learning Failover