Elastic Stack 8. x Cookbook : Over 80 Recipes to Perform Ingestion, Search, Visualization, and Monitoring for Actionable Insights /

Chen, Huage,
author

Monografía

Unlock the full potential of Elastic Stack for search, analytics, security, and observability and manage substantial data workloads in both on-premise and cloud environments Key Features Explore the diverse capabilities of the Elastic Stack through a comprehensive set of recipes Build search applications, analyze your data, and observe cloud-native applications Harness powerful machine learning and AI features to create data science and search applications Purchase of the print or Kindle book includes a free PDF eBook Book Description Learn how to make the most of the Elastic Stack (ELK Stack) products--including Elasticsearch, Kibana, Elastic Agent, and Logstash--to take data reliably and securely from any source, in any format, and then search, analyze, and visualize it in real-time. This cookbook takes a practical approach to unlocking the full potential of Elastic Stack through detailed recipes step by step. Starting with installing and ingesting data using Elastic Agent and Beats, this book guides you through data transformation and enrichment with various Elastic components and explores the latest advancements in search applications, including semantic search and Generative AI. You'll then visualize and explore your data and create dashboards using Kibana. As you progress, you'll advance your skills with machine learning for data science, get to grips with natural language processing, and discover the power of vector search. The book covers Elastic Observability use cases for log, infrastructure, and synthetics monitoring, along with essential strategies for securing the Elastic Stack. Finally, you'll gain expertise in Elastic Stack operations to effectively monitor and manage your system. What you will learn Discover techniques for collecting data from diverse sources Visualize data and create dashboards using Kibana to extract business insights Explore machine learning, vector search, and AI capabilities of Elastic Stack Handle data transformation and data formatting Build search solutions from the ingested data Leverage data science tools for in-depth data exploration Monitor and manage your system with Elastic Stack Who this book is for This book is for Elastic Stack users, developers, observability practitioners, and data professionals ranging from beginner to expert level. If you're a developer, you'll benefit from the easy-to-follow recipes for using APIs and features to build powerful applications, and if you're an observability practitioner, this book will help you with use cases covering APM, Kubernetes, and cloud monitoring. For data engineers and AI enthusiasts, the book covers dedicated recipes on vector search and machine learning. No prior knowledge of the Elastic Stack is required

https://rebiunoda.pro.baratznet.cloud:28443/OpacDiscovery/public/catalog/detail/b2FpOmNlbGVicmF0aW9uOmVzLmJhcmF0ei5yZW4vMzYxNTMwNTc

Título: Elastic Stack 8. x Cookbook Over 80 Recipes to Perform Ingestion, Search, Visualization, and Monitoring for Actionable Insights Huage Chen, Yazid Akadiri, and Shay Banon

Edición: First edition

Editorial: Birmingham, England Packt Publishing [2024] 2024

Descripción física: 1 online resource (688 pages)

Contenido: Cover -- Title Page -- Copyright and Credits -- Dedication -- Foreword -- Contributors -- Acknowledgments -- Table of Contents -- Preface -- Chapter 1: Getting Started - Installing the Elastic Stack -- Deploying the Elastic Stack on Elastic Cloud -- How to do it... -- How it works... -- There's more... -- Installing the Elastic Stack with ECK -- Technical requirements -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Installing a self-managed Elastic Stack -- Getting ready -- How to do it... -- How it works... -- There's more... -- Creating and setting up data tiering -- Getting ready -- How to do it on your local machine... -- How it works (on self-managed)... -- How to do it on Elastic Cloud... -- How to do it on ECK... -- There's more... -- See also -- Creating and setting up additional Elasticsearch nodes -- Getting ready -- How to do it... -- How it works... -- How to do it on Elastic Cloud... -- How to do it on ECK... -- There's more... -- See also -- Creating and setting up Fleet Server -- Getting ready -- How to do it on a self-managed Elastic Stack... -- How it works... -- Setting up on Elastic Cloud -- See also -- Setting up snapshot repository -- Getting ready -- How to do it... -- How it works... -- There's more... -- Chapter 2: Ingesting General Content Data -- Introducing the Wikipedia Movie Plots dataset -- Technical requirements -- Adding data from the Elasticsearch client -- Getting ready -- How to do it... -- How it works... -- There's more... -- Updating data in Elasticsearch -- Getting ready -- How to do it... -- How it works... -- There's more... -- Deleting data in Elasticsearch -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Using an analyzer -- Getting ready -- How to do it... -- How it works... -- There's more... -- Defining index mapping -- Getting ready -- How to do it... -- How it works There's more... -- See also -- Using dynamic templates in document mapping -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Creating an index template -- Getting ready -- How to do it... -- How it works... -- There's more... -- Indexing multiple documents using Bulk API -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Chapter 3: Building Search Applications -- Technical requirements -- Searching with Query DSL -- Getting ready -- How to do it... -- How it works... -- There's more... -- Building advanced search queries with Query DSL -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Using search templates to pre-render search requests -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Getting started with Search Applications for your Elasticsearch index -- Getting ready -- How to do it... -- How it works... -- Building a search experience with the Search Application client -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Measuring the performance of your Search Applications with Behavioral Analytics -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Chapter 4: Timestamped Data Ingestion -- Technical requirements -- Deploying Elastic Agent with Fleet -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Monitoring Apache HTTP logs and metrics using the Apache integration -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Deploying standalone Elastic Agent -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Adding data using Beats -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also Setting up a data stream manually -- Dataset -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Setting up a time series data stream manually -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Chapter 5: Transform Data -- Technical requirements -- Creating an ingest pipeline -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Enriching data with a custom ingest pipeline for an existing Elastic Agent integration -- Getting ready -- How to do it... -- How it works... -- There's more... -- Using a processor to enrich your data before ingesting with Elastic Agent -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Installing self-managed Logstash -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Creating a Logstash pipeline -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Setting up pivot data transform -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Setting up the latest data transform -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Downsampling your time series data -- Getting ready -- How to do it... -- How it works... -- There's more... -- See also -- Chapter 6:

ISBN: 1-83763-350-9

Materia: Client/server computing Open source software

Autores: Akadiri, Yazid, author Banon, Shay, author

Enlace a formato físico adicional: 1-83763-429-7