PROVOKE™
Training | Consulting | OutSourcing

SPLUNK COURSE CONTENT

Module 1: Splunk Fundamentals

- Introduction to SIEM & Role of Splunk
- Evolution and History of Splunk
- Splunk as a Platform: ITSI, SOAR, ES Overview
- Splunk Cloud vs Splunk Enterprise
- Splunk Pricing & Licensing (Per GB vs Ingest)
- Splunk UI & Navigation Tour
- REST API Overview (Introduction level)

Module 2: Splunk Architecture – Deep Dive

- Forwarders, Indexers, Search Heads, and Deployment Server
- Role of Cluster Master, License Master, Deployer, DMC
- Splunk Processing Pipeline – Event Lifecycle
- How Data Moves: Parsing, Indexing, Searching
- Index Time vs Search Time knowledge object separation
- Indexer Clustering Concepts: RF/SF, Multisite vs Single Site
- Search Head Clustering – Captain, KV Store, Deployer

Module 3: Installation & Deployment

➤ **Windows/Linux Deployment**

- Installation on Windows (GUI)
- Installation on Linux (RPM, TAR)
- Running Splunk as a Service
- Installation best practices (non-root, ports, permissions)

➤ **Cloud Server Setup**

- Creating Linux Server in AWS/GCP/Oracle Cloud
- IP, Firewall, User, and Storage Setup
- SSH Keys and Connectivity

Module 4: Splunk File System & Internals

- Splunk Directory Structure:
  - o $SPLUNK_HOME/etc/
  - o bin, var/lib/splunk, apps, users, system
- Configuration Layering:
- o System Default → App Default → App Local → User Local
  Configuration File Precedence & Merging
  App Packaging & Deployment Structure

Module 5: Core Configuration Files (Mastery)

- **inputs.conf** – Monitor, Scripted, TCP, UDP, HTTP
- **props.conf** – Line Breaking, Timestamping, Field Extraction
- **transforms.conf** – Routing, Filtering, Masking, Aliasing
- **outputs.conf** – Forwarder Outputs, Load Balancing
- **indexes.conf** – Index Definitions, Freezing, Thawing
  limits.conf, server.conf, web.conf,deploymentclient.conf,serverclass.conf
  authentication.conf, savedsearches.conf and many more
  Real-world scenarios with each file

Module 6: Forwarder Deployment

- Universal vs Heavy Forwarder (with use cases)
- Forwarder Installation & Configuration
- Outputs.conf and Forwarding Pipelines
- Managing Forwarders via Deployment Server
- Indexer Acknowledgement & Load Balancing
  Splunk Deployment Apps & Clients

Module 7: Buckets, Indexing & Storage

- Index Lifecycle: Hot → Warm → Cold → Frozen → Thawed
- Bucket Rotation & Retention
- Indexing Volumes & Storage Paths
- License Impact on Indexing
  Troubleshooting Index Issues

Module 8: Data Onboarding Techniques

- File & Directory Monitoring
- TCP/UDP Streams
- Windows Event Logs & Performance Counters
- Syslog Integration (UDP/TCP/Rsyslog/Chronyd)
- **Scripted Inputs (Python, Shell)**
- **HTTP Event Collector (HEC)** – Secure Token-based ingestion
- Modular Inputs (Advanced API-based ingestion)

Module 9: Add-ons and Technology Integrations

- Installing and Managing Add-ons (TA-*)
- **Splunk Add-on for AWS** (CloudTrail, CloudWatch, VPC Flow)
- **Splunk DB Connect** – SQL/Oracle/MySQL integrations
- **TA for Windows, TA for Linux**
- **Cloud Monitoring** – Azure Monitor, Google Cloud
- **Cisco, Palo Alto, Fortinet TA Setup**
- Integration with Sysmon, Winlogbeat, and NXLog

Module 10: Splunk Security & Access Control

- Role-based Access Control (RBAC)
- Authentication:
  - Splunk Native
  - LDAP/Active Directory
  - SAML Integration with Azure/Gsuite
- Secure Web Access (HTTPS setup, certs)
- KV Store & secrets.conf (Encrypted credentials)

Module 11: Administering Apps, Alerts & Dashboards

- Installing and Configuring Splunk Apps
- Managing Knowledge Objects as Admin
- Managing Scheduled Searches and Alerts
- Dashboard App Permissions
  Backup and Export of Dashboards & Reports

Module 12: Troubleshooting & Monitoring

- splunkd.log, metrics.log, btool command

- CLI Commands for Admins (splunk show, clean, rebuild)
- Splunk Diag Collection
- Troubleshooting:
  - Forwarding issues
  - Parsing and timestamp issues
  - Missing events or lag
  - License violation troubleshooting
- Using DMC and Monitoring Console
- Health Report and Disk Usage

Module 13: Cluster Implementation (Enterprise Level)

➤ **Indexer Cluster (Single & Multi-site)**

- Setting up CM, Peer Nodes, Searchable Indexes
- Validating RF/SF and Bucket Replication

➤ **Search Head Cluster**

- Deployer Configuration
- Captain Election
- App Deployment & Bundle Push
- KV Store Synchronization

Module 14: REST API & Automation (Introductory)

- Introduction to Splunk REST API for Admin Tasks
- Use curl & tokens for search/jobs management
- Introduction to Automating via Ansible & Bash

Module 15: Real-Time Projects, Labs & Interview Readiness

A. Project Environment Setup

- **Distributed Environment Simulation:**
  - 1 Search Head, 1 Indexer, 1 Deployment Server
  - 1 Universal Forwarder (Windows/Linux)
  - 1 Heavy Forwarder
- **Add-on Integration:**
  - TA-Windows, TA-Linux
  - AWS Add-on (CloudTrail/VPC)

- DB Connect (MySQL/Oracle)

B. Single-Site Indexer Clustering (Hands-On)

- Indexer Cluster Setup with:
  - 1 Cluster Master (CM)
  - 3 Peer Nodes (Indexers)
  - 1 Search Head
- Configure:
  - Replication Factor (RF) and Search Factor (SF)
  - App deployment to peer nodes using CM
- Validate:
  - Bucket Replication
  - Failover and Recovery
  - Search Head > Indexer communication

C. Multi-Site Indexer Clustering (Enterprise-Grade Practice)

- Multi-site Configuration:
  - 2 Sites with 2 Indexers each
  - CM with site awareness
  - SH to perform multisite searches
- Configure:
  - Site-specific RF/SF (e.g., origin:2, total:3)
  - Smart Data Placement
- Disaster Recovery (DR) simulation and testing

D. Search Head Clustering (Production Ready)

- Setup of 3 SH nodes + 1 Deployer
- Enable Captain Election and KV Store Sync
- Deploy dashboards, apps, and alerts via Deployer
- Validate Bundle Push and SHC health

E. Real-Time Use Cases

- Onboard Logs:
  - Windows AD & Security Logs
  - Linux Syslog & Secure Log

- AWS CloudTrail/VPC Flow Logs
- MySQL DB Logs via DB Connect
- Apache/Nginx Access Logs

- Alerts:
  - Failed Logins (Windows/Linux)
  - Database Connection Failures
  - Cloud Access Outside Office Hours
- Dashboards:
  - Linux Resource Monitoring
  - AWS Cloud Overview
  - Security Alert Summary

F. Interview and Certification Readiness

- 30+ Real-time Splunk Admin Interview Questions
- Sample Splunk Architect Scenario Q&A
- Resume Building with Role-Specific Keywords
- Preparation Guide for:
  - Splunk Core Certified Admin
  - Splunk Enterprise Certified Admin

Outcomes After Completion

Deploy, configure, and manage a secure and scalable Splunk environment
Confidently onboard data from any source
Implement Splunk enterprise clustering with full fault tolerance
Troubleshoot and optimize Splunk for performance and reliability

Wishing you the best in your Splunk Admin journey — go build secure, scalable, and high-performance environments!

---- P. Ravikumar

PROVOKE™
Training | Consulting | OutSourcing

SPLUNK DEVELOPMENT COURSE CONTENT

Module 1: Introduction to Splunk Development

- Role of a Splunk Developer vs Admin
- How Splunk fits into the IT & Security Ecosystem
- Introduction to Splunk Search Language (SPL)
- Overview of Dashboards, Reports, Alerts, and Apps
  Developer Tools: UI, REST API, Postman, CLI

Module 2: Mastering SPL (Search Processing Language)

- Basic Search Commands: search, table, sort, where
- Filtering, Wildcards, Boolean Expressions
- Statistical Commands: stats, eventstats, streamstats
- Transforming Commands: chart, timechart, top, rare
- Time Commands: earliest, latest, now, relative_time
  Advanced Commands:
  - rex, spath, eval, lookup, join, append
  - tstats, mstats, inputlookup, outputlookup
  - bin, bucket, xyseries, fillnull
- Creating Fields on the Fly: eval, fieldalias, rename
- SPL Optimization Techniques & Search Job Inspector

Module 3: Knowledge Objects (KO) Management

- Fields, Tags, Event Types
- Lookups:
  - Static CSV, External Lookup Scripts, KV Store
- Aliases, Calculated Fields, Field Extractions (rex vs props)
- Macros & Event Types
- Transaction vs Stats-based sessionization
  Creating and Managing Data Models
  CIM (Common Information Model) & Acceleration
  Tags and Event Tagging Strategies

Module 4: Reports, Dashboards & Visualization

- Creating Reports and Saving Searches
- Dashboard Studio vs Classic Simple XML Dashboards
- Dashboard Panels, Tokens, and Drilldowns
- Inputs: Time Pickers, Dropdowns, Text Boxes
- Dynamic Dashboards with Advanced Token Logic
- Reusable Dashboard Templates
- Scheduled Reports & Email Configuration
- Real-World Dashboard Examples:
    - Server Health
    - SOC Analyst Dashboard
    - IT Ticketing Overview (ServiceNow integration)

Module 5: Alerts & Scheduled Searches

- Creating Alerts (Real-time, Scheduled, Rolling Window)
- Trigger Conditions and Throttling
- Alert Actions:
    - Email
    - Webhook
    - Script
    - Indexing alerts (into summary index)
- Correlation Searches (Intro to ES-style logic)
- Alert Management Best Practices

Module 6: Splunk App Development & Packaging

- What is a Splunk App? App Structure Overview
- Building a Custom App with Dashboards, Lookups, and Permissions
- Creating and Organizing Your App in $SPLUNK_HOME/etc/apps/
- App.conf, default.meta and packaging
- User/Role-based App Permissions
- Deploying Apps to SHC and Search Peers
- Publishing Private/Internal Apps

Module 7: Summary Indexing & Accelerations

- What is Summary Indexing? Why and When to Use It
- Setting up Summary Index Searches
- Best Practices for Performance Optimization
- Report Acceleration vs Data Model Acceleration vs Summary Index
- Managing Accelerated Data Models

Module 8: REST API for Developers

- Introduction to Splunk REST API
- Exploring Endpoints: Search, Jobs, Users, Apps, Alerts
- Use Cases: Dashboard automation, Search triggers
- REST via Postman and Python
- Generating Tokens for Secure API Access
  Sample Use Case: Triggering Alert from External System

Module 9: Integration with External Systems

- Integration with:
  - o ServiceNow (via Webhooks or REST)
  - o Email & SMTP
  - o Slack/Teams Alerts
  - o SOAR Playbooks (intro)
- Data Onboarding via HEC (Developer view)
- Pushing Data into Splunk (scripted/API)

Module 10: CIM Compliance and Enterprise Security (Intro)

- What is CIM? Use in ES and Security Add-ons
- Mapping your fields to CIM models
- Importance of Tags & Event Types
- ES Correlation Searches (Basic Developer Role)
  Developer's role in Security Content Authoring

Module 11: Hands-on Projects & Use Cases

Use Case 1: Server Monitoring

- Build a dashboard to show CPU, Memory, Disk from Linux Logs
- Scheduled Alerts for critical resource usage

Use Case 2: Windows Security Monitoring

- Monitor AD logs for failed logins
- Build correlation logic for brute force detection

Use Case 3: AWS Cloud Monitoring

- Use AWS Add-on for CloudTrail & GuardDuty logs
- Build KPI dashboard for Cloud Assets

Use Case 4: DB Monitoring

- Use DB Connect to ingest MySQL queries
- Build dashboard to track long-running queries

Use Case 5: Application Logs

- Parse JSON logs using spath
- Visualize app performance using timechart and bar charts

Module 12: Certification & Interview Readiness

- Splunk Core Certified Power User (SPL Focus)
- Splunk Enterprise Certified Developer
- Resume Tips for Splunk Developer Role
- Real-world Developer Scenarios
  40+ Interview Q&A covering SPL, Dashboards, Lookups, Apps

Bonus Topics

- Git Versioning for Splunk App Code
- Working with SimpleXML and JavaScript for advanced dashboards
- Dynamic Forms and Token Chaining
- Creating Custom Visualizations (HTML/JS overview)
- Developer Best Practices: Naming, Tagging, Reusability

Good developers write queries. Great developers turn data into decisions.

---- P. Ravikumar