INFOSOFT IT SOLUTIONS

Training | Projects | Placements

Revathi Apartments, Ameerpet, 1st Floor, Opposite Annapurna Block,

Infosoft It solutions, Software Training & Development Institute, 9059683947|9182540872

Apache Storm

Introduction to Apache Storm

- Overview of Apache Storm: Features, benefits, and use cases
- History and evolution of real-time stream processing
- Comparison with other stream processing frameworks (e.g., Apache Kafka, Apache Flink)

Setting Up Apache Storm

- Installation and Configuration of Apache Storm
- Setting up Zookeeper for coordination and state management
- Configuring and deploying Storm clusters (single-node and multi-node setups)

Apache Storm Architecture

- Understanding Storm architecture components: Nimbus, Supervisors, Workers, Zookeeper
- Topologies and Tasks: Spouts, Bolts, and their roles in data processing
- Fault tolerance and reliability mechanisms in Storm

Developing Storm Topologies

- Writing Storm Topologies in Java: Creating Spouts and Bolts
- Implementing Real-time Data Processing: Data flow and transformations

- Configuring Parallelism: Strategies for task and worker parallelism

Data Model and Message Passing

- Tuples and Streams: Storm's data model for message passing
- Groupings and Stream Partitions: Controlling how data is distributed among tasks
- Using Fields Grouping, Shuffle Grouping, and Custom Grouping

Stream Processing Patterns

- Common Stream Processing Patterns: Filtering, Transformation, Aggregation
- Windowing and Time-based Operations: Sliding and tumbling windows
- Exactly-once Processing Semantics: Achieving message processing guarantees

Trident API for Stateful Stream Processing

- Introduction to Trident: High-level abstraction for stateful processing
- Trident Topologies: Using stateful operators and transactions
- Integrating Trident with existing Storm topologies

Integrating Storm with Data Sources

- Integration with Messaging Systems: Apache Kafka, RabbitMQ
- Reading and Writing to Databases: Using Storm with JDBC connectors
- Implementing Custom Spouts and Bolts for specific data sources

Performance Tuning and Optimization

- Monitoring Storm Clusters: Metrics and monitoring tools
- Tuning Storm Configurations: JVM settings, parallelism hints
- Scaling and managing large-scale Storm deployments

Security and Reliability in Storm

- Securing Storm Clusters: Authentication and authorization mechanisms
- Configuring SSL/TLS for secure communication within Storm
- Handling Failures and Recovery: Strategies for fault tolerance

Advanced Topics in Apache Storm

- Advanced Bolt Techniques: Implementing custom stateful bolts, complex event processing (CEP)
- Machine Learning with Storm: Real-time predictive analytics and model deployment
- IoT Use Cases: Handling Internet of Things (IoT) data streams with Storm

Real-world Use Cases and Projects

- Implementing Apache Storm in Production: Case studies across industries
- Project Work: Hands-on projects to apply learned concepts
- Designing and deploying Storm topologies for specific business requirements

Career Development and Job Preparation

- Building a career in Real-time Data Processing: Skills and certifications
- Interview Preparation: Apache Storm-related interview questions
- Freelancing and Consulting Opportunities in real-time stream processing