

Effective J2EE Server Performance Tuning

Course Summary

Description

This course teaches J2EE System Administrators and others how to tune a J2EE Server. A J2EE Server is designed to serve 100s of simultaneous users. The number of end-users that effectively connect to one server depends on how the Server is configured. The same software might support 500 users if the system is configured poorly and as many as 2000 simultaneous users if the system is configured properly. This course gives the practical skills needed to tune a server to serve the maximum number of simultaneous users while still maintaining adequate levels of service as measured in service response times. This course can be split into two 2 1/2 day classes. There are versions of this course for various J2EE Servers including WebLogic and Tomcat

Objectives

At the end of this course, students will be able to:

- Specify Service Level Agreements (SLAs)
- Tune the JVM memory usage
- Understand the Elements of Tuning
- Obtain Performance Measurements
- Create Simulated Workloads
- Understand Wait Based Tuning
- Use Fine-tuned Performance
- Understand the Principles of Tuning Clusters
- Understand Trending, Forecasting and Capacity Planning

Topics

- Specifying Service Level Agreements
- Tune the JVM Memory Usage
- The Elements of Tuning
- Obtaining Performance Measurements
- Create Simulated Workloads
- Wait Based Tuning
- Fine Tuning Performance
- Principles of Tuning Clusters
- Trending, Forecasting and Capacity Planning

Audience

This course would benefit Weblogic, Tomcat and other J2EE System Administrators who are responsible for the performance of the Server.

Prerequisites

The student should be familiar with the basics of deploying and monitoring J2EE Applications on a J2EE Server and with the basic structure of J2EE applications.

Duration

Five days

Effective J2EE Server Performance Tuning

Course Outline

- I. Overview**
 - A. The Performance Problem
 - B. How Fast is Fast Enough?
 - C. Hiding Bad Performance
 - D. The 6 main Performance costs: Memory, CPU Time, Network Traffic, I/O,
 - E. System Calls, Resource Blocking
- II. Specifying Service Level Agreements**
 - A. The Four Types of Performance Measures
 - B. End-User Response Time
 - C. Request Throughout
 - D. Resource Utilization
 - E. Application Availability
 - F. Specifying a Service Level Agreement
- III. Tune the JVM Memory Usage**
 - A. The JVM Memory and Garbage Collector Model
 - B. Configuring the Garbage Collector
 - C. Tracing the behavior of the Garbage Collector
 - D. Recommended initial memory configuration
- IV. The Elements of Tuning**
 - A. Tuning Memory
 - B. Tuning Thread Pools
 - C. Tuning JDBC Connection Pools
 - D. Tuning other Caches
 - E. Using a separate server for static content
- V. Obtaining Performance Measurements**
 - A. Using the Server Console
 - B. Using Connection Pools that provide Statistics
 - C. Polling Statistics with Java Programs
 - D. Using JMX Scripts
 - E. Using Other Tools
- VI. Create Simulated Workloads**
 - A. Workload Generation Tools
 - B. Using JMeter
 - C. Creating Graduated Loads
 - D. Test Load Strategies
- VII. Wait Based Tuning**
 - A. The J2EE Tiers
 - B. Wait Based Tuning
 - C. Tuning From the Back to the Front - Theory and Method
 - D. Tuning Example
 - E. Tuning the Database Tier
- VIII. Fine Tuning Performance**
 - A. Precompiling JSPs
 - B. JMS Tuning
 - C. Advanced JDBC Tuning
 - D. Other Items to Tune
- IX. Principles of Tuning Clusters**
 - A. Horizontal Clustering
 - B. Vertical Clustering
 - C. Minimizing Object Movement
 - D. Load Testing a Cluster
- X. Trending, Forecasting and Capacity Planning**
 - A. Predict Future Trends
 - B. Anticipating Server Overload
 - C. Performance Management Plans