

# Demystifying Performance Issues in Java based Products



Sudhir Kakkar, Senior Technology Architect, Accenture



GREAT INDIAN  
**DEVELOPER**  
**SUMMIT**

A circular graphic element for the Developer Summit, featuring a grid of small dots in the center, surrounded by concentric circles and a larger outer ring.

# Agenda

- Importance of Performance
- Product Architecture
- Key challenges
- Issues and resolution
- Tools used
- Other considerations - NFRs

# Why is Performance so important?

- Transaction volumes reaching far beyond expected limits
  - Flipkart sells 60,000 Xiaomi Redmi 1s in 5.2 s
  - Salesforce handles 1.3 billion transactions per day
  - Amazon handled 3 million transactions on 'Cyber Monday'
- Royal Wedding site witnessed 6 million unique visitors with 15 million page views on a single day processing more than 2000 requests/second
- Among applications with 100,000 concurrent users (typically very large Web-based systems), by 2016, Java will capture a 55% market share — up from 10% in 2006 and 40% in 2010.

© Gartner, Inc. Examining Java EE and Microsoft Software Platforms: Competitive Trends and Future Directions, By Mark Driver, November 22, 2011.

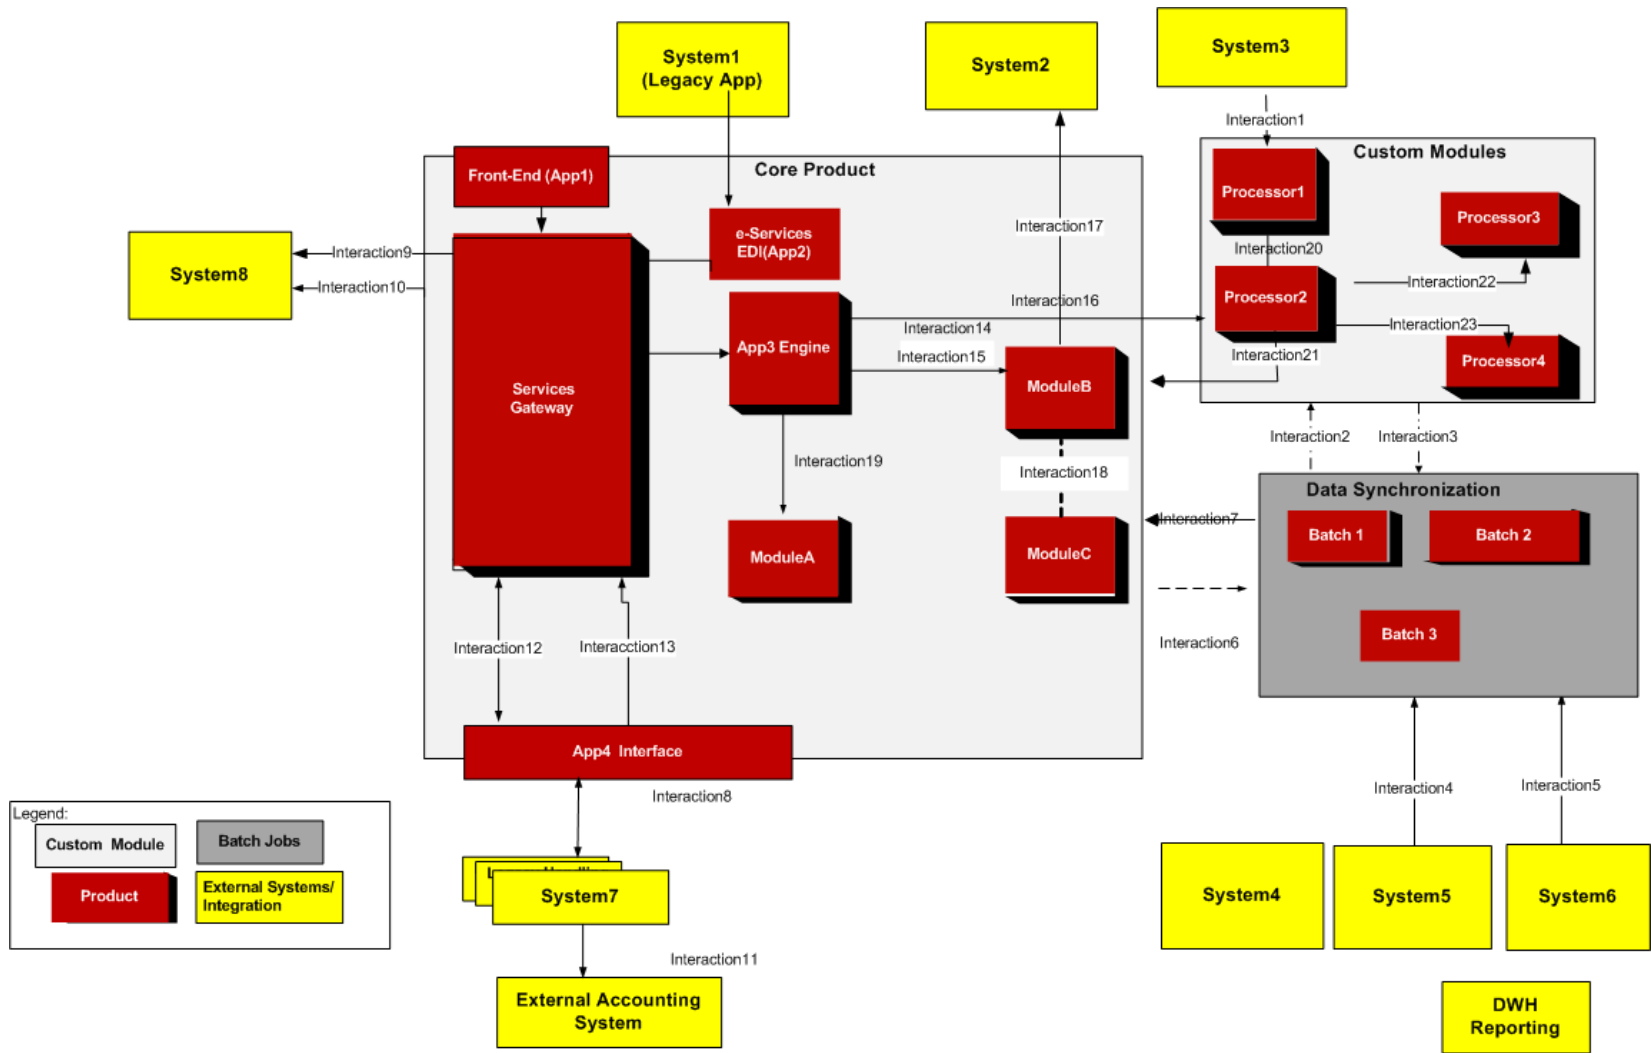
## Sources:

<http://gadgets.ndtv.com/mobiles/news/60000-redmi-1s-smartphones-go-out-of-stock-in-52-seconds-xiaomi-596605>

<https://developer.salesforce.com/blogs/engineering/2013/10/under-the-hood-how-the-salesforce-platform-handles-1-3-billion-transactions-per-day>

html <http://www.dailymail.co.uk/news/article-2072837/Amazons-Cyber-Monday-saw-3million-items-sold-Christmas-rush.html>

# Sample System Architecture of a Product



# Key Challenges

- Performance requirements differ from one client to another
- Gathering Non functional requirements are challenging (often vague)
  - The response time should meet or exceed that of current system
  - Batch jobs should complete processing in 4 hours
  - End to end response time not to exceed 3 seconds
  - System should be able to handle the load for the next 3 years
- Products need to work with disparate software, platforms and protocols
- Support for multiple platforms e.g. Linux, Websphere, Weblogic, Oracle, MySQL etc
- Different Tool set for each implementation
  - Dump Analyzers for IBM products different from Oracle products
  - CA Introscope by one client, Dynatrace for another and AppDynamics for another
- Difficult to perform load test with different legacy systems for individual client

# Issues and Resolution

Issue	Resolution
Queries taking very long time to execute in spite of proper indexes on columns	Size of redo log file too small. The frequency of log switches becomes very high hence slowing down the overall disk performance. The amount of I/O for the particular implementation was much higher than normal. Increasing the size of this file improved performance.
Poor response time during certain periods of time	Higher heap utilization due to heavy load. Insufficient heap size leads to excessive GC thereby degrading the performance. Amount of data cached large for certain implementation. Increasing the heap resolves the issue.
Search operations extremely slow at certain times	Number of rows returned too large causing multiple round trips between the Application and DB server. Setting the <code>setFetchsize</code> of prepared statement to a high value limits the number of round trips improving performance
Some pages take long time to load	Enabling gzip compression and mod deflate option on the http server to compress the pages results in quicker loading.
Performance issues with rules engine	Knowledge Base creation is an expensive operation and takes considerable amount of time. The system may move to unstable state within few minutes of operation. By caching the knowledge base, transaction time gets reduced considerably.
Insert and update queries taking very long for a few tables	Increasing the <code>IniTrans</code> property for a particular table increases its ability to handle more concurrent update transactions on that table.

# Tools used

- AppDynamics
- CA Introscope
- Riverbed
- Dynatrace
- YourKit
- JProfiler
- Visual VM
- Oracle JRockit Mission Control
- IBM Heap Analyzer
- IBM Thread and Monitor Dump Analyzer
- Oracle AWR
- WebLog Expert
- HP Load Runner

# Other considerations - NFRs in Products

## ▪ Scalability

- The product must support thousands of concurrent users
- Must be horizontally and vertically scalable

## ▪ Security

- Prevent un-authorized access
- Protect data at rest and data in motion via encryption and firewalls

## ▪ Integration

- Impact of Performance of legacy systems both
- Support multiple interaction paradigms (Req/Resp, Async Messaging etc)

## ▪ Availability, reliability, backup and recovery

- Availability requirements vary from one client to another and can be 99.99% or more
- Support for tools to monitor availability against the stated SLAs
- Ability to automatically trigger and monitor batch jobs
- Provision for automatic data backup
- Easy recovery from failures in different environments

## ▪ Usability

- Rich UI for better user experience
- Easy Configuration - by people with minimal technology knowledge



Thank you!