Web Sphere Server 9.0.5.1 Best Practices
Oracle FLEXCUBE Investor Servicing
Release 12.0.4.8.9
[June] [2021]

ORACLE
FINANCIAL SERVICES

Table of Contents

1. Introduction 1
   1.1 Background 1
   1.2 Basics of Websphere 1
      1.2.1 Profile 2
   1.3 References 3
2. JDBC Tuning 4
3. JVM Tuning 6
   3.1 JVM Heap Size 6
   3.2 Thread Pool Size 7
4. Logging 8
   4.1 Diagnostic Trace 8
   4.2 JVM Logs 9
   4.3 Process Logs 10
   4.4 IBM Service Logs 10
   4.5 Change Log Level Details 11
   4.6 NCSA Access and HTTP Error Logging 12
5. Session Management 13
6. Appendix A: Frequently Encountered Errors 15
7. Appendix A: Frequently Encountered Errors 16

1. Introduction

1.1 Background

IBM® WebSphere® Application Server 9.0.5.1 supports a range of applications, each with their own unique set of features, requirements, and services. Just as no two applications will use an application server in exactly the same way, no single set of tuning parameters will likely provide the best performance for any two different applications.

This document explains the best practices followed for Web Sphere application server tuning in the perspective of Oracle FLEXCUBE IS.

1.2 Basics of Websphere

IBM websphere application server cluster deployment contains the below key elements:

- Cell
- Nodes
  - Deployment Manager Node- “DMGR”
  - Node- “NodeXX”
  - Node Agent- “NAXX”
- Profiles
- Cluster
- Cluster Members
- Data Sources

1.2.1 Profile

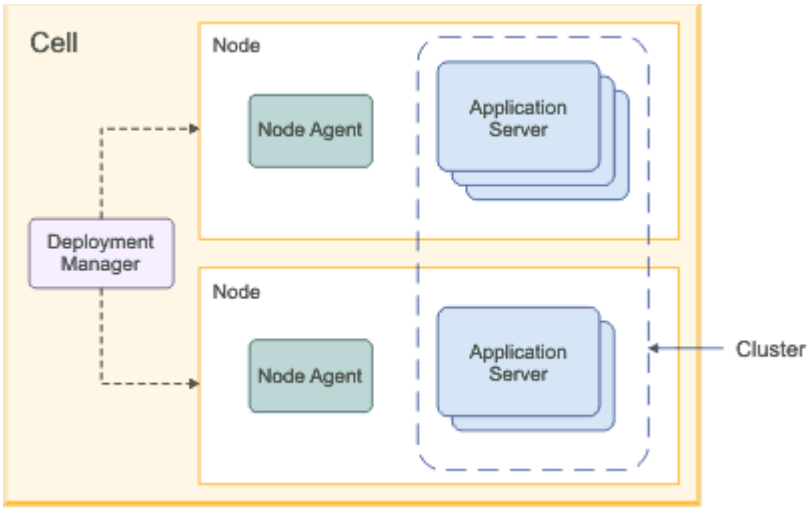

- Cell: A cell is a grouping of nodes into a single administrative domain. In a Network Deployment environment, a cell can consist of multiple nodes (and node groups), which are all administered from a single point, the deployment manager.
- Node: A node is an administrative grouping of application servers for configuration and operational management within one operating system instance
- Node Agent: In distributed server configurations, each node has a node agent that works with the deployment manager to manage administration processes. A node agent is created automatically when you add (federate) a stand-alone node to a cell.
- Cluster: A cluster is a logical collection of application server processes that provides workload balancing and high availability. Application servers that belong to a cluster are members of that cluster and must all have identical application components deployed on them.
- A profile is a Websphere runtime environment formed by collection of User data and Product files. Product Files are shared application binaries for Websphere. User data is set of user customizations for a specific runtime environment.

Prominent profile types are:

- Stand-alone Application Server: An application server environment runs Enterprise Application. Application server is managed from its own administrative console and functions independently from other application server.
- Deployment Manager: A Deployment Manager manages operations for a logical group or cell of other servers. It is the central administration point of a cell that consists of multiple nodes and node groups in a distributed server configuration. The deployment manager uses the node agent to manage the application servers within one node. A deployment manager provides management capability for multiple federated nodes and can manage nodes that span multiple systems and platforms. A node can only be managed by a single deployment manager and must be federated to the cell of that deployment manager.

Note: Deployment Manager is part of Network Deployment Edition of Websphere.

1.3 References

Before proceeding with the document, ensure the below documents are followed for setup

- For resource creation(JDBC Provider, Datasource, Queue connection factory, JMS queue), refer to Resource_Creation_WAS.doc
- For cluster configuration, refer to FCIS_Cluster_Creation_WAS.doc
- For application deployment, refer to FCIS_Application_WAS.doc
- For deployment of Gateway applications, refer to GATEWAY_Applications_WAS.doc
- For SSL configuration in Websphere, refer to SSL_Configuration_WAS.doc

2. JDBC Tuning

The JDBC provider object encapsulates the specific JDBC driver implementation class for access to the specific vendor database of your environment. The datasource object supplies your application with connections for accessing the database. In a typical production environment database and application needs to use the perfect combination of parameters to achieve higher throughput.

Navigation Path >> Resources->JDBC->Data Sources ->DATA_SOURCE_NAME->Connection Pools ->

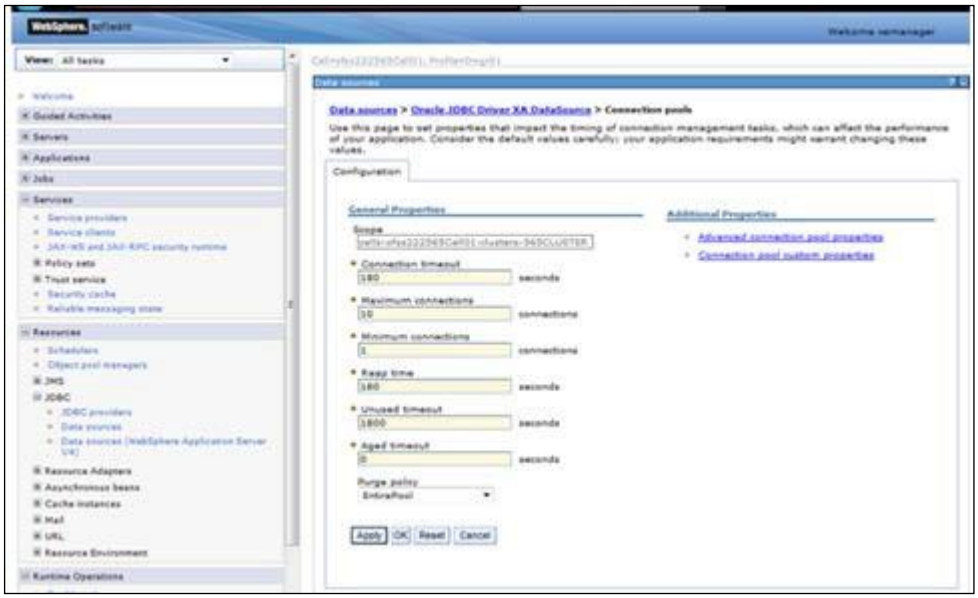

PARAMETER	VALUE
Connection Time out	180
Maximum Connection	10
Minimum Connections	1
Reap Time	180
Unused Time	180
Aged Timeout	0
Purge Policy	Entire Pool

Navigation Path >> Resources-> JDBC->Data Sources ->DATA_SOURCE_NAME->Connection Pools ->Advanced Connection Pool Properties

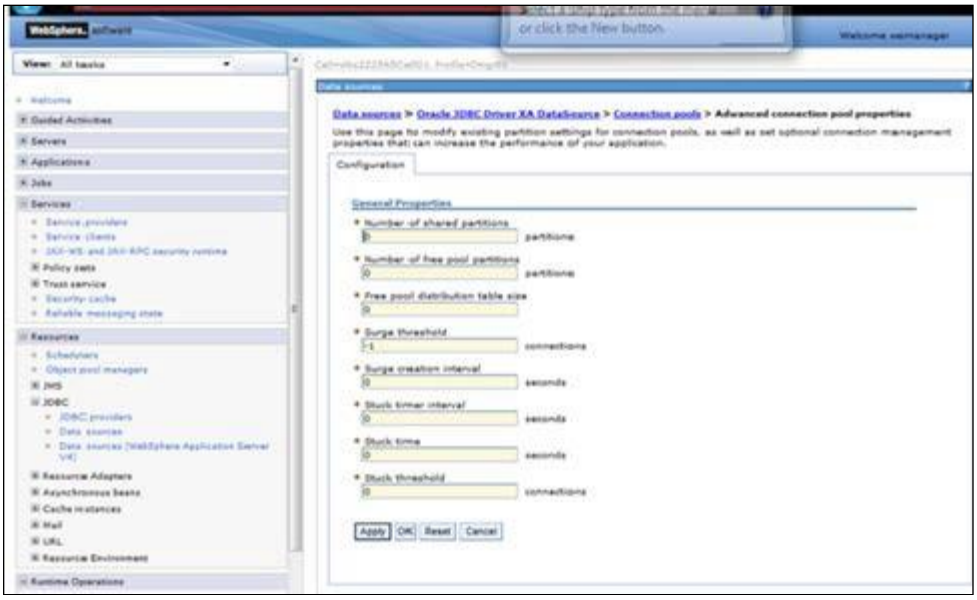

PARAMETER	VALUE
Number of shared pool partitions	0
Number of free pool partitions	0
Free pool distribution table size	0
Surge threshold	-1
Surge creation interval	0
Stuck timer interval	0
Stuck time interval	0
Stuck threshold	0

3. JVM Tuning

3.1 JVM Heap Size

When a Java Virtual Machine (JVM) is started, it obtains a large area of memory from the underlying operating system. This area is called the heap, and Java performs its own memory management by allocating areas of the heap as memory is needed by the process. Performance tuning generally starts with the Java Virtual Machine (JVM), which serves as the foundation for the application server. From that point forward, tuning is primarily driven by the application server components that are used by the application.

> Every 75 concurrent Oracle FLEXCUBE users require one websphere application server of size 4GB. I.e. for 300 concurrent Oracle FLEXCUBE IS users, it is recommended to have 4 websphere application servers.

Navigation Path >> Application Server->APPLICATION_SERVER_NAME-> Server

Infrastructure->Java And Process Management->Process Definition-> Java Virtual machine

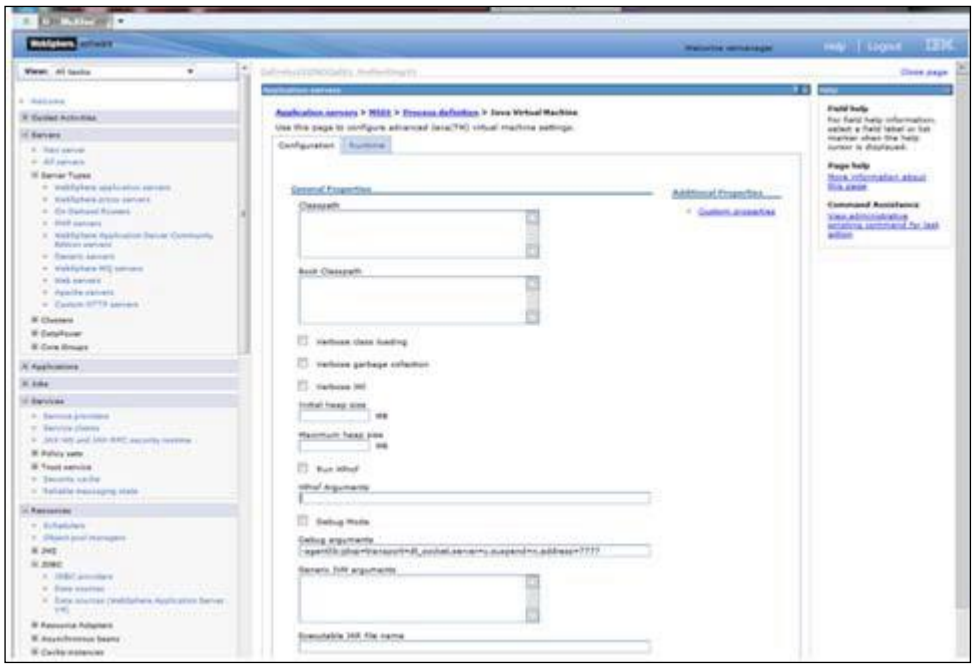

PARAMETER	VALUE
Initial Heap Size	1024
Maximum Heap Size	4096

3.2 Thread Pool Size

Each task performed by the server runs on a thread obtained from one of WebSphere Application Server's many thread pools. A thread pool enables components of the server to reuse threads, eliminating the need to create new threads at run time to service each new request. Three of the most commonly used (and tuned) thread pools within the application server are:

Web container: Used when requests come in over HTTP

Default: Used when requests come in for a message driven bean or if a particular transport chain has not been defined to a specific thread pool.

ORB: Used when remote requests come in over RMI/IIOP for an enterprise bean from an EJB application client, remote EJB interface, or another application server.

Navigation Path >> Application Server->APPLICATION_SERVER_NAME->Additional Properties->Thread Pools->

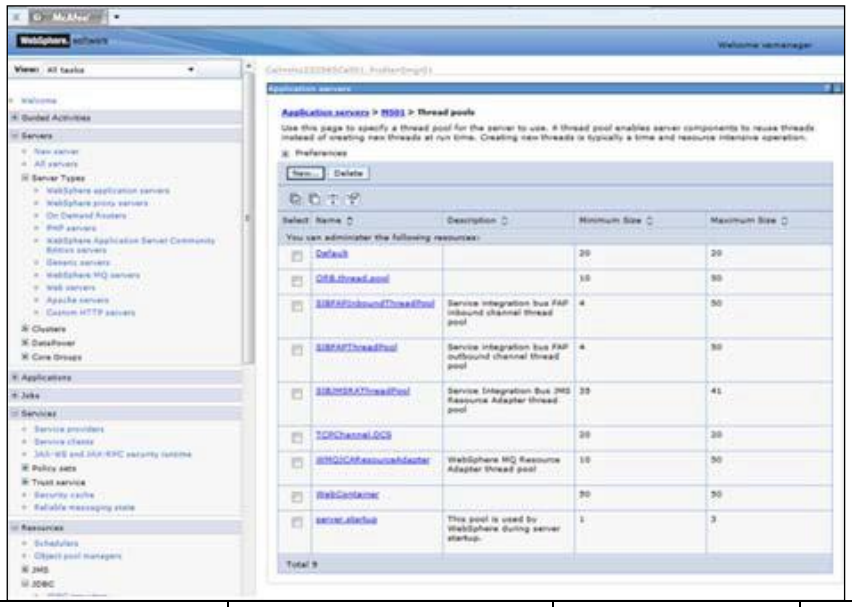

PARAMETER	Minimum Size	Maximum Size	Thread Inactive
Default	20	20	5000
ORB.thread.pool	10	50	3500
Web Container	50	50	60000

4. Logging

When working with IBM to debug request processing problems, there might be occasions where low-level tracing components must be enabled to capture details for how the request is processed. These low-level Application server trace components do not have knowledge of the request intent or the potential data within. Therefore, when enabled, it is possible that these tracing components might potentially include sensitive information, in plain text, in the trace file.

It is recommended whenever possible to not enable these types of tracing components on a production system and attempt to simulate the problem on a quality assurance environment to capture the appropriate information.

Navigation Path >> Application Server->APPLICATION_SERVER_NAME->Process Definition-> Logging and Tracing

4.1 Diagnostic Trace

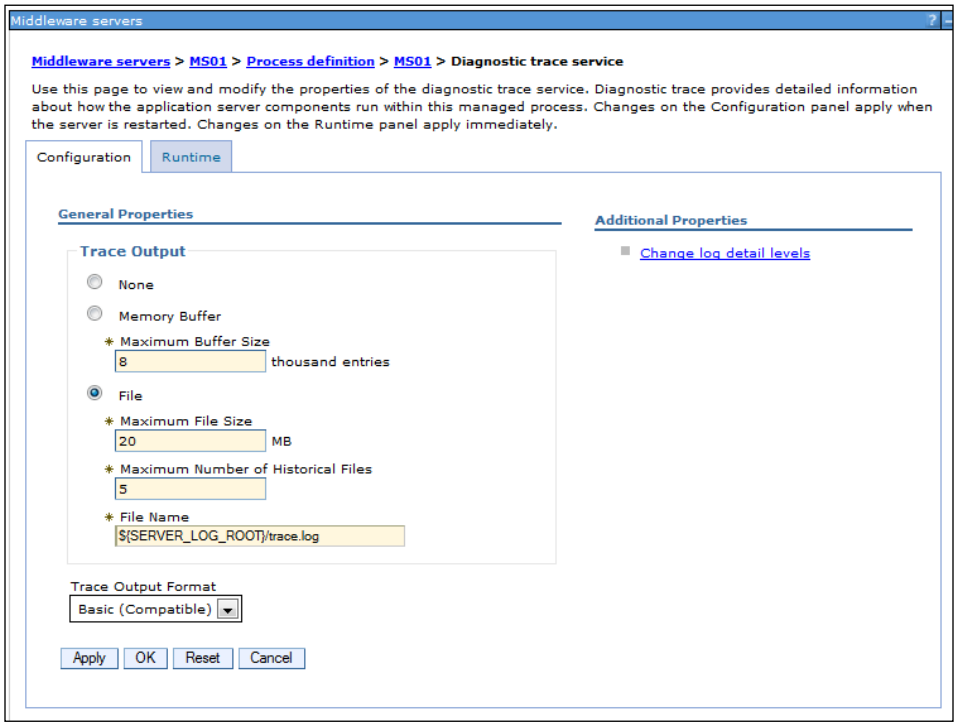

PARAMETER	VALUE
Trace Output	File

4.2 JVM Logs

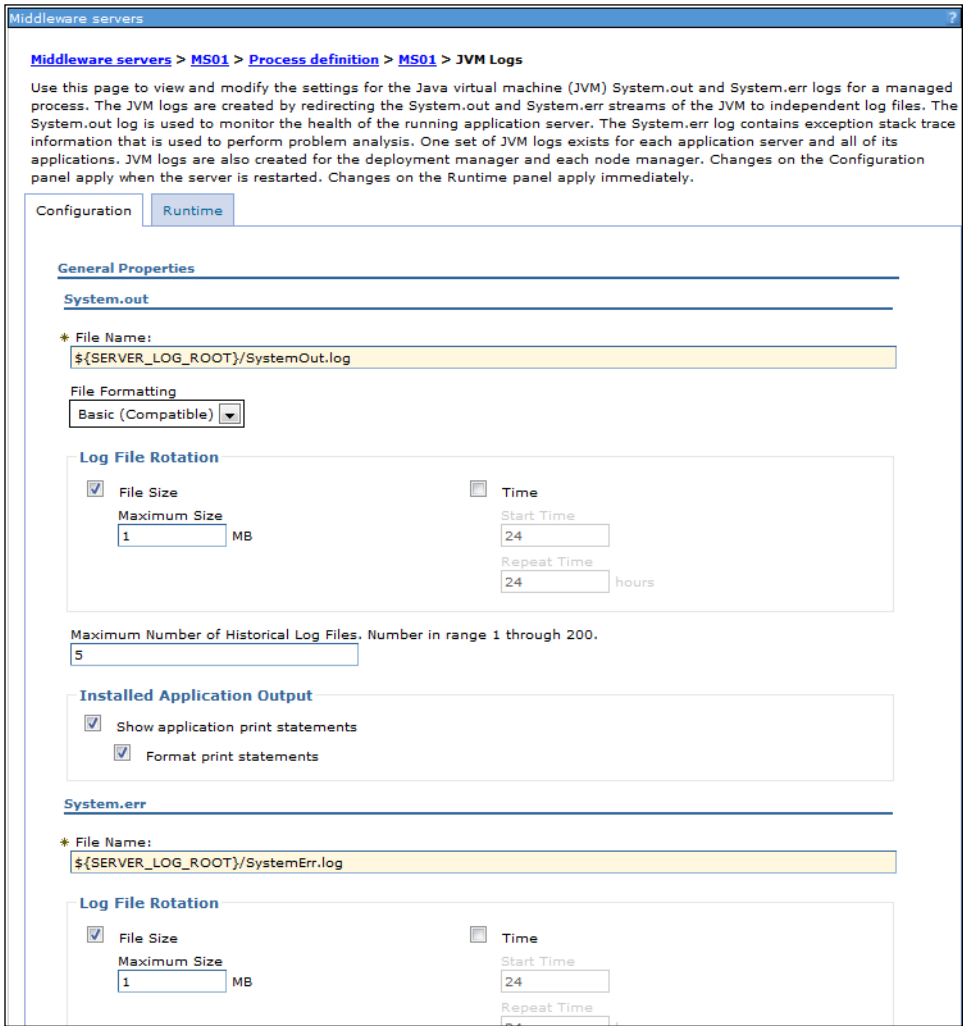

PARAMETER	VALUE	Remarks
System. out	${SERVER_LOG_ROOT}/SystemOut.log	File Name and File
Sysem.err	${SERVER_LOG_ROOT}/SystemErr.log	File Name and File

4.3 Process Logs

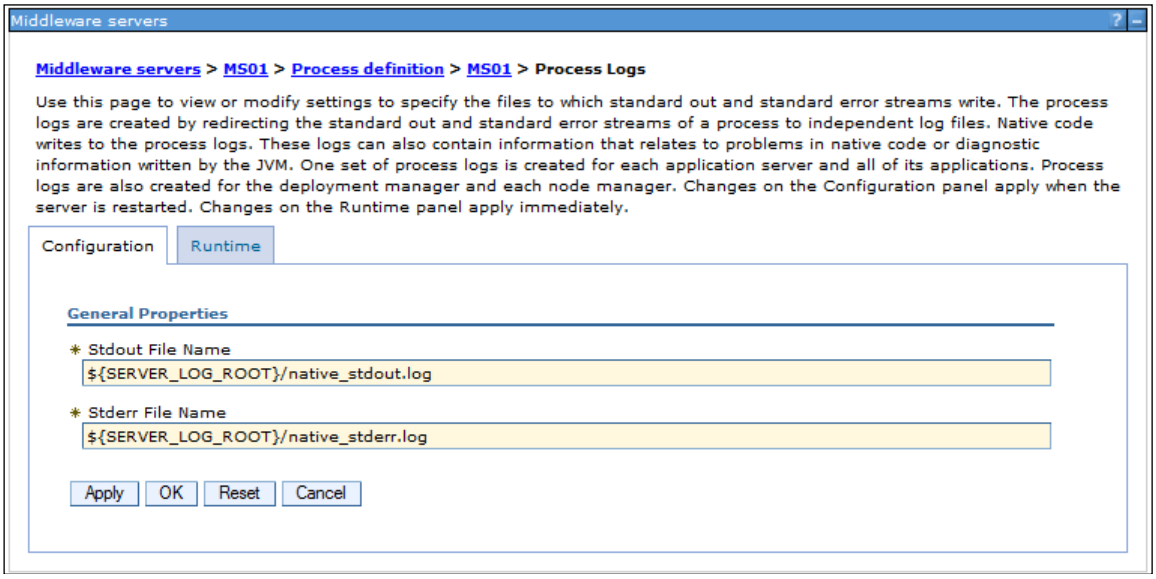

PARAMETER	VALUE
Stdout File	${SERVER_LOG_ROOT}/native_stdout.log
Stderr file name	${SERVER_LOG_ROOT}/native_stderr.log

4.4 IBM Service Logs

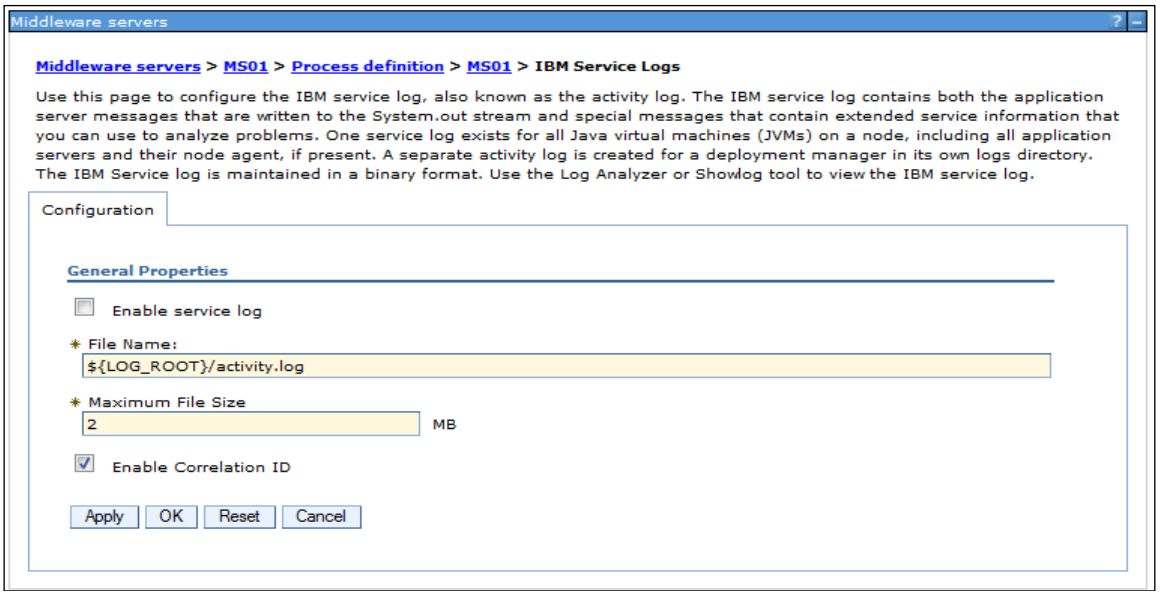

PARAMETER	VALUE
Enable Service Log	False
Maximum File Size	2
File Name	${LOG_ROOT}/activity.log
Enable Correlation Id	True

4.5 Change Log Level Details

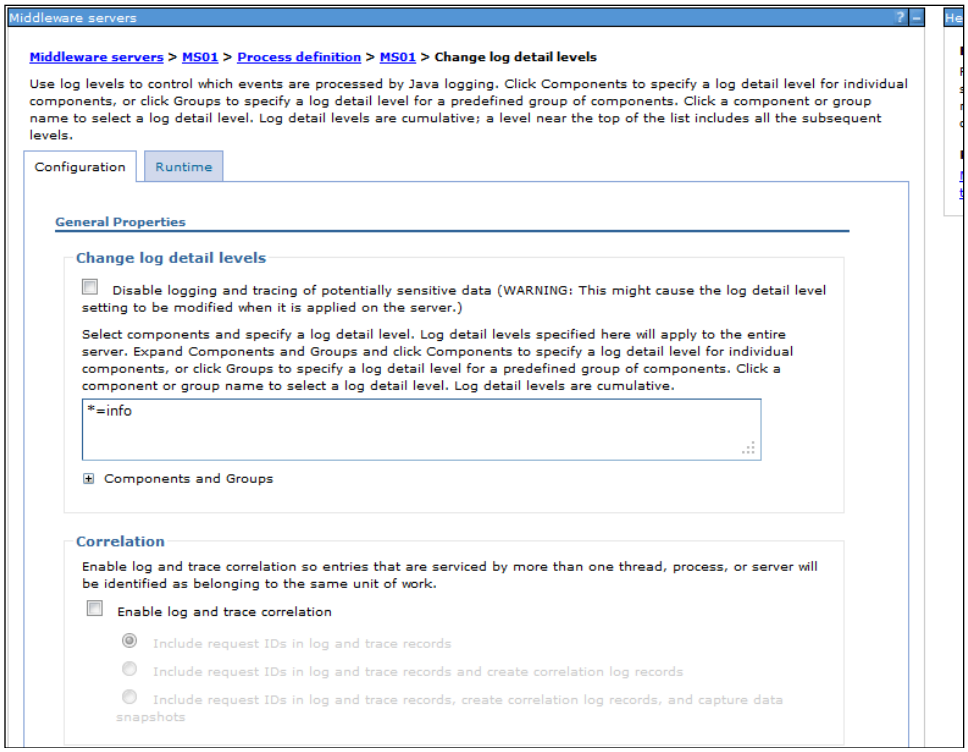

PARAMETER	VALUE
Disable logging	False
Enable Log and Trace	False

4.6 NCSA Access and HTTP Error Logging

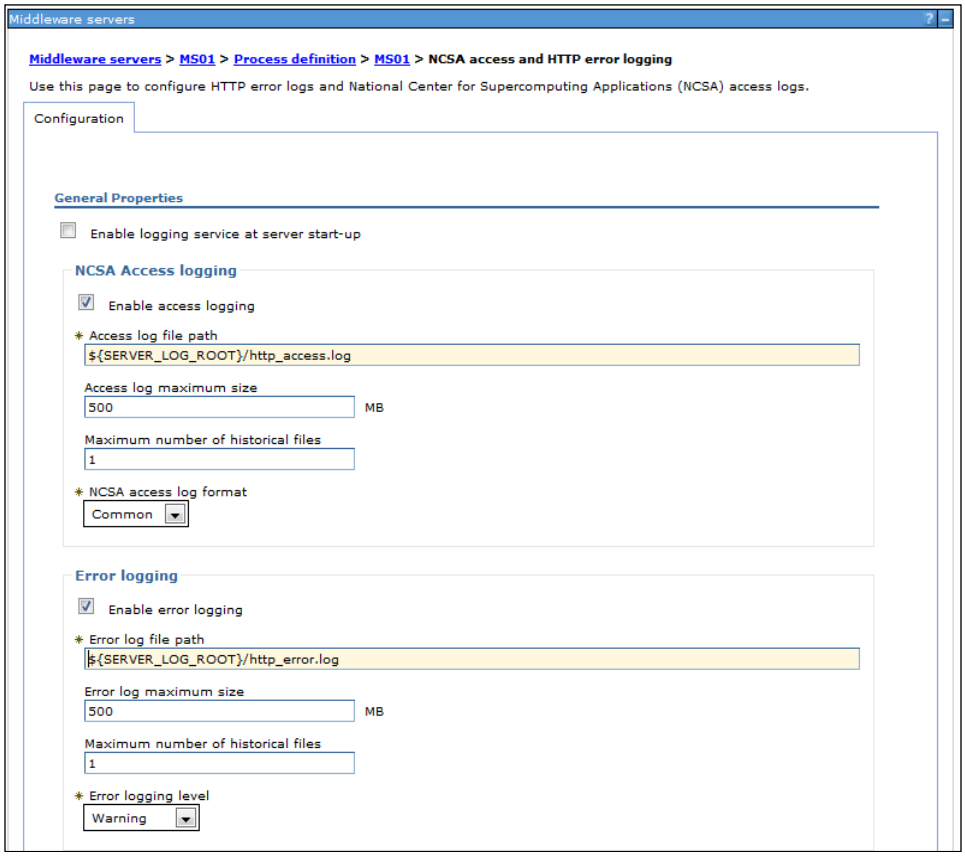

PARAMETER	VALUE
NCSA Access logging	True
Error logging Enable	True

5. Session Management

Web browsers and applications use HTTP to communicate. Since HTTP is a stateless protocol (meaning that each command is executed independently without any knowledge of the commands that came before it), there must be a way to manage sessions between the browser side and the server side. Session management is used to configure session manager properties to control the behaviour of Hypertext Transfer Protocol (HTTP) session support.

Navigation Path >> Enterprise Application->APPLICATION_NAME->Web Module Properties->Session Management

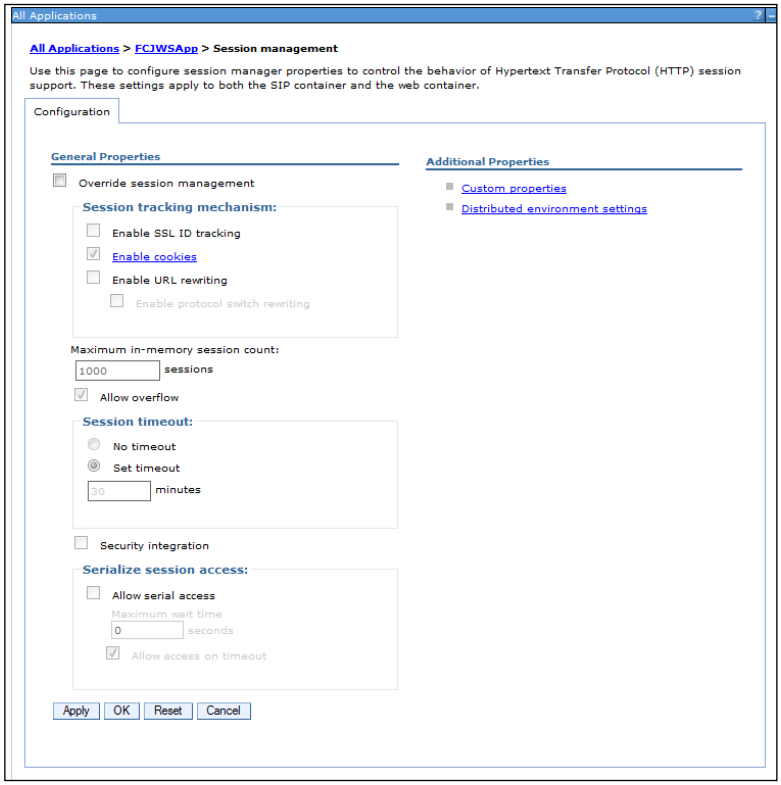

PARAMETER	VALUE
Enable SSL ID Tracking	False
Enable Cookies	True
Enable URL rewriting	False
Maximum in-memory session count	1000
Session timeout	30 Minutes
Security Integration	False

Serialization Session access	True

6. Appendix A: Frequently Encountered Errors

Error 1: Websphere is causing too many database locks.

Solution: In order to avoid problems with database locks issue, change the custom property "webSphereDefaultIsolationLevel=2". By default Webpshere uses Repeatable Read isolation level(4) for transactions. Before proceeding with the isolation level change make sure these locks are not caused by the application.

Possible values	JDBC isolation level	Isolation level
8	TRANSACTION_SERIALIZABLE	Repeatable Read (RR)
4 (default)	TRANSACTION_REPEATABLE_READ	Read Stability (RS)
2	TRANSACTION_READ_COMMITTED	Cursor Stability (CS)
1	TRANSACTION_READ_UNCOMMITTED	Uncommitted Read (UR)
0	TRANSACTION_NONE	No Commit (NC)

Navigation Path >> Resources-> JDBC->Data Sources ->DATA_SOURCE_NAME-> Custom Properties

If this property is found then update the value to 4 and Save, else create a new property

- Click New.
- Enter webSphereDefaultIsolationLevel for the name field.
- Enter 2 for the value field.

7. Appendix A: Frequently Encountered Errors

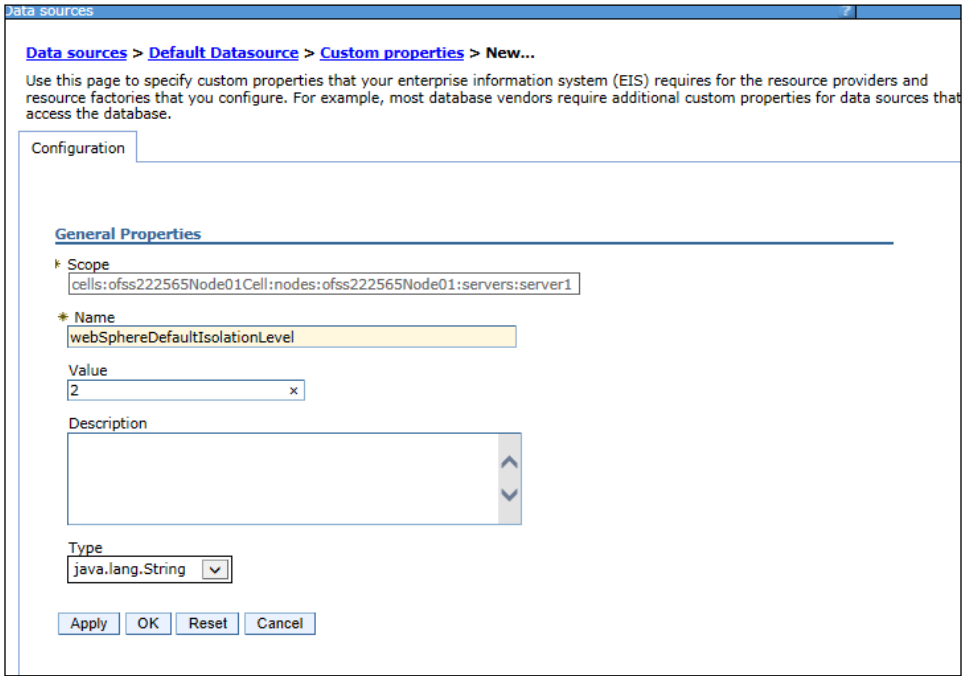

Web Sphere Server 9.0.5.1 Best Practices
[June] [2021]
Version 12.0.4.8.9

Oracle Financial Services Software Limited
Oracle Park
Off Western Express Highway
Goregaon (East)
Mumbai, Maharashtra 400 063
India

Worldwide Inquiries:
Phone: +91 22 6718 3000
Fax:+91 22 6718 3001
www.oracle.com/financialservices/