



## Yellowfin Clustering Guide

Version 7.1

March 2014

**Yellowfin®**  
Release 7.1  
Clustering Guide

Under international copyright laws, neither the documentation nor the software may be copied, photocopied, reproduced, translated or reduced to any electronic medium or machine-readable form, in whole or in part, without the prior written permission of Yellowfin International Pty Ltd, except in the manner described in the software agreement.

The information in this document is subject to change without notice. If you find any problems with this documentation, please report them to Yellowfin in writing at [support@yellowfin.bi](mailto:support@yellowfin.bi). Yellowfin does not warrant that this document is error free.

Copyright © Yellowfin International 2015 All rights reserved.  
Portions © Copyright Microsoft Corporation. All rights reserved.

Trademarks:  
Yellowfin and the Yellowfin Logo are registered trademarks of Yellowfin International.

All other product and company names mentioned herein are the trademarks of their respective owners.

Version: 1.5  
Published: March 2015

# Table of Contents

Yellowfin Clustering Overview .....	4
Application Level Messaging .....	4
Container Level Session Replication .....	4
Yellowfin Database Clustering .....	5
Licensing .....	5
Installing a Yellowfin Cluster .....	6
File System Replication .....	6
Reinstallation.....	6
Yellowfin for Application Messaging .....	7
Multicast Cluster Messaging (DYNAMIC mode).....	7
Webservice Cluster Messaging (LEGACY mode).....	7
Modifications to web.xml for DYNAMIC mode.....	8
Modifications to web.xml for LEGACY mode. ....	9
web.xml file parameters.....	10
Additional Parameters .....	11
Background Tasks .....	12
System Tasks .....	12
Task Scheduler .....	12
Container Level Session Replication .....	13
Caveats.....	13
Modifications to web.xml.....	14
Modifications to server.xml.....	14
Modifications to ROOT.xml .....	14

## Yellowfin Clustering Overview

Yellowfin can be clustered on multiple servers to allow for high-availability and load-balancing. The actual load-balancing (multiplexing requests from external requests) can be achieved with a hardware load-balancer, or load-balancing software. This guide will outline the modifications required to Yellowfin's configuration to enable clustering, but not the external environment that multiplexes the incoming requests.

It is required that the load-balancing infrastructure delivers network traffic to the Yellowfin application server transparently, as if the packets were connecting directly to the application server directly. If Session Replication is not enabled, the load-balancing infrastructure will need to provide “sticky-session” functionality, where traffic for a user’s session should be sent to the same node.

Yellowfin supports clustering at 2 levels:

### Application Level Messaging

Yellowfin will communicate with other nodes when application wide messages need to be sent. This communicates messages like refreshing caches, and logging out concurrent users. This is controlled by a background servlet in Yellowfin that will communicate with other nodes.

Yellowfin also has cluster safe webservice calls. Which allows a webservice Single-Sign-On request to come from an application to any cluster node, and return a token id. The token id will then be valid on any node that the token id is sent to.

### Container Level Session Replication

The Yellowfin application has also been written to allow user session information to be replicated onto multiple cluster nodes concurrently. User session information is the memory footprint that is stored on the server for each user session. Enabling Session Replication allows for the users session to continue, even if the cluster node they were connected to experiences a failure.

Without Session Replication, failure of a cluster node will destroy a users session, and they will need to login again to another node.

Session Replication is a function of the Java Container. An example of how this can be achieved with Tomcat will be discussed in this document.

## Yellowfin Database Clustering

A Yellowfin Cluster must share a single Yellowfin Database instance. The database could be a single database instance, that is shared across all Yellowfin nodes, or a database that is clustered itself. It is assumed that database clustering / replication is done by the client, as every database implements this in a different way. Database Clustering / Replication should be transparent to Yellowfin, as Yellowfin will have a single database URL for the entire cluster.

## Licensing

The Yellowfin Licence file is stored in the Yellowfin database, and because of this, the Licence must contain a reference to all hostnames in the cluster. A clustered licence can be requested from Yellowfin.

## Installing a Yellowfin Cluster

As a Yellowfin Cluster requires only one database instance, it means that the Yellowfin Installation process will differ slightly when installing a cluster.

The first Yellowfin Cluster node can be installed the same as standalone instance of Yellowfin. The command-line or GUI installer can be used.

Additional Yellowfin Cluster nodes do not require the Yellowfin database. There are two options to options for installing additional nodes:

### File System Replication

Assuming that the filesystem structure on each node is the same, a new node can be created by duplicating the Yellowfin installation directory onto another computer. The only change required to each node, are changes to the web.xml file which is discussed later in this document.

### Reinstallation

After the initial node is installed, and the Yellowfin database is created, the Yellowfin installer can be used for the installation of subsequent nodes, however, this process will create additional Yellowfin databases. The installer can be allowed to create these databases, but these can be removed once the installation process is finished. Each node's web.xml can then be modified to point to the database created with the first nodes installation.

## Yellowfin for Application Messaging

Each Yellowfin node must have the ClusterManagement servlet enabled for the node to become “Cluster Aware”. The ClusterManagement servlet is enabled by adding an additional configuration block to the web.xml file on each node.

The Application Messaging is performed using webservices, allowing different nodes to initiate cache clearing or changes to licensing details.

### Multicast Cluster Messaging (DYNAMIC mode)

Yellowfin inter-node messaging is now handled by a highly configurable multicast library called JGroups. Using this method will automatically find other nodes in the cluster sharing the same Yellowfin database.

Dynamic mode also allows for multiple cluster instances to run on the same server.

The default configuration of JGroups uses UDP multicast messages to determine group membership and find new nodes. There may be environments where these types of messages can not be sent. For example, Amazon does not allow multicast packets on its internal network between nodes. The Multicast Cluster Messaging adapter allows you to pass a XML configuration file to configure JGroups to use other methods for node discovery. This file can be referenced by passing the path to the BroadcastConfiguration servlet parameter within the ClusterManagement servlet.

### Webservice Cluster Messaging (LEGACY mode)

Yellowfin’s legacy cluster messaging is handled by AXIS webservices. This requires that all nodes be defined at startup, and that the service end-point, port, user and password be defined in each node’s web.xml file. Legacy mode does not allow cluster instances to reside on the same host.

## Modifications to web.xml for DYNAMIC mode.

The following servlet definition needs to be added to the web.xml on each node:

```
<!-- Cluster Management -->
<servlet>
  <servlet-name>ClusterManagement</servlet-name>
  <servlet-class>com.hof.mi.servlet.ClusterManagement</servlet-class>
  <init-param>
    <param-name>ClusterType</param-name>
    <param-value>DYNAMIC</param-value>
  </init-param>
  <init-param>
    <param-name>SerialiseWebserviceSessions</param-name>
    <param-value>>true</param-value>
  </init-param>
  <init-param>
    <param-name>ChecksumRows</param-name>
    <param-value>>true</param-value>
  </init-param>
  <init-param>
    <param-name>EncryptSessionId</param-name>
    <param-value>>true</param-value>
  </init-param>
  <init-param>
    <param-name>EncryptSessionData</param-name>
    <param-value>>true</param-value>
  </init-param>
  <init-param>
    <param-name>AutoTaskDelegation</param-name>
    <param-value>>true</param-value>
  </init-param>
  <load-on-startup>11</load-on-startup>
</servlet>
```

## Modifications to web.xml for LEGACY mode.

The following servlet definition needs to be added to the web.xml on each node:

```
<!-- Cluster Management -->
<servlet>
  <servlet-name>ClusterManagement</servlet-name>
  <servlet-class>com.hof.mi.servlet.ClusterManagement</servlet-class>
  <init-param>
    <param-name>ServiceUser</param-name>
    <param-value>admin@yellowfin.com.au</param-value>
  </init-param>
  <init-param>
    <param-name>ServicePassword</param-name>
    <param-value>test</param-value>
  </init-param>
  <init-param>
    <param-name>ServiceAddress</param-name>
    <param-value>/services/AdministrationService</param-value>
  </init-param>
  <init-param>
    <param-name>ServicePort</param-name>
    <param-value>80</param-value>
  </init-param>
  <init-param>
    <param-name>ClusterHosts</param-name>
    <param-value>
      192.168.4.184
    </param-value>
  </init-param>
  <init-param>
    <param-name>SerialiseWebserviceSessions</param-name>
    <param-value>true</param-value>
  </init-param>
  <init-param>
    <param-name>ChecksumRows</param-name>
    <param-value>true</param-value>
  </init-param>
  <init-param>
    <param-name>EncryptSessionId</param-name>
    <param-value>true</param-value>
  </init-param>
  <init-param>
    <param-name>EncryptSessionData</param-name>
    <param-value>true</param-value>
  </init-param>
  <load-on-startup>11</load-on-startup>
</servlet>
```

## web.xml file parameters

The following properties outline the options in the ClusterManagement servlet definition.

Property	Value
ClusterType	Either DYNAMIC or LEGACY. DYNAMIC will use multicast messaging and automatically find other nodes in the cluster. LEGACY is the default, and will use webservices to communicate with a defined list of cluster nodes.
BroadcastConfiguration	A JGroups configuration file. This allows for a custom Jgroups configuration to be used in environments where multicast networking is not available. This is for DYNAMIC mode only and is optional. By default, JGroups will use the configuration defined in udp.xml.
ServiceUser	User that will validate the webservice connection to other nodes. For LEGACY mode only.
ServicePassword	Password for the ServiceUser. For LEGACY mode only.
ServicePasswordEncrypted	True/False. Is the contents of the ServicePassword encrypted.
ServiceAddress	Address of the Yellowfin webservice. For LEGACY mode only.
ServicePort	Port on which Yellowfin is running. For LEGACY mode only.
ClusterHosts	This is a comma-separated list of all nodes in the Cluster. These can include IP addresses or hostnames. For LEGACY mode only.
SerialiseWebserviceSessions	True/False. This is required if using Single Sign-On on the Cluster. It can serialise tokens to the database so that the token can be accessed from any node.
ChecksumRows	True/False. Security option to check sum the serialised webservice session records in the database. This helps prevent modification to the table which could lead to the creation of unauthorised sessions in Yellowfin.
EncryptSessionId	True/False. Security option to encrypt the serialised webservice session id in the database. This helps prevent modification to the table which could lead to the creation of unauthorised sessions in Yellowfin.
EncryptSessionData	True/False. Security option to encrypt the serialised webservice session records in the database. This helps prevent modification to the table which could lead to the creation of unauthorised sessions in Yellowfin.
AutoTaskDelegation	True/False. This is for DYNAMIC mode only and is optional. When enabled, the Cluster will automatically assign a node to perform background tasks. When this is turned on you do not need to manually configure a node to run background tasks. Yellowfin will turn off System

	Tasks on a node at startup if a another node is already delegated to perform this task.
SessionReplication	True/False. Set this to true if container level Session Replication is enabled. This will modify the logic used for destroying sessions in a cluster.

## Additional Parameters

Under some circumstances DYNAMIC mode may not work in some environments. This will usually be due to the networking configuration on the server. By default, JGroups will use IPv6 if it is available. IPv4 can be forced, which may allow for it to work correctly. Add the following to catalina.sh/catalina.bat file:

### Catalina.sh

```
JAVA_OPTS="$JAVA_OPTS -Djava.net.preferIPv4Stack=true"
```

### Catalina.bat

```
set JAVA_OPTS=%JAVA_OPTS% -Djava.net.preferIPv4Stack=true
```

When using the default Jgroups configuration (udp.xml), the multicast address and port can also be configured. This may help in environments where cluster nodes are not discovering each other.

### Catalina.sh

```
JAVA_OPTS="$JAVA_OPTS -Djgroups.udp.mcast_addr=228.0.0.5 -  
Djgroups.udp.mcast_port=47885"
```

### Catalina.bat

```
set JAVA_OPTS=%JAVA_OPTS% -Djgroups.udp.mcast_addr=228.0.0.5 -  
Djgroups.udp.mcast_port=47885
```

## Background Tasks

Each Yellowfin node will be configured by default to run background tasks, which includes the broadcasting of reports. This potentially could send reports multiple times to end users. Because of this it is recommended that background tasks be enabled on one node only.

## System Tasks

Disable the background system tasks by removing (or commenting-out) the following XML block from the web.xml. This will disable system tasks, like Group Flattening, LDAP synchronization, Event Archiving, Document Cleanup and Average Run time calculations.

```
<servlet>
  <servlet-name>SystemTaskManager</servlet-name>
  <servlet-class>com.hof.servlet.SystemTaskManager</servlet-class>
  <load-on-startup>8</load-on-startup>
</servlet>
```

## Task Scheduler

Disable the task scheduler by adding the following to the web.xml file, inside the MIStartup Servlet block. This will disable all tasks visible in the Task Schedule menu within the Yellowfin interface. This includes broadcasting, populating cached filters, scheduled reports and populating cached views.

```
<init-param>
  <param-name>DisableTaskScheduler</param-name>
  <param-value>TRUE</param-value>
</init-param>
```

# Container Level Session Replication

Container Level Session Replication allows Yellowfin to be deployed in a fail-safe environment where in the event of a node in the cluster failing, a currently connected Yellowfin user can continue their session on another node. This is achieved with Session Replication at the Application Server level, and is not a function of Yellowfin itself, and hence is not a supported component of Yellowfin. This configuration also requires a load-balancer to evenly distribute incoming requests to different cluster nodes, which will not be addressed in this guide.

Different Application Servers may have different ways of implementing this functionality. This example is shown using Tomcat, which is shipped with Yellowfin.

Setting up session replication consists of the following steps:

1. Modifications to web.xml
2. Modifications to server.xml
3. Modifications to ROOT.xml

More information regarding Tomcat Session Replication can be found at <http://tomcat.apache.org>

## Caveats

There are several caveats to using Session Replication with Yellowfin (or any Java web application). The two main downsides of this are Session Replication overhead, and memory usage.

There is additional overhead in replicating each session to all nodes in the cluster. This happens after each user's request, and by default, only the changed session elements are copied to other nodes. This means all nodes will be doing additional work each time a user does something on any connect node.

The function of Session Replication is to replicate a user's memory footprint on the server, to all nodes in the cluster. This means the sessions for all users in the cluster will be stored on each node, taking up additional memory. Expanding the cluster with additional nodes is then restricted by the amount of memory that each node has.

It should be considered if this overhead is worth the added fail-safe functionality that it brings. If Session Replication is not used, and a node fails, a user's session will be destroyed, but they'll be able to login to another node to continue their work.

## Modifications to web.xml

The text “<distributable/>” needs to be added to the web.xml file on each node underneath the “<webapp>” line.

```
<web-app>
  <distributable/>
  <!-- System Event and Debug classes -->
  <listener>
    <listener-class>com.hof.servlet.SysSessionListener</listener-class>
  </listener>
```

## Modifications to server.xml

Changes need to be made to the server.xml file in the Yellowfin/appserver/conf directory.

Add the additional XML entry within the “<host>” XML block.

Add:

```
<Cluster className="org.apache.catalina.ha.tcp.SimpleTcpCluster"/>
```

The IP address “228.0.0.4” is a multicast address that session replication messages will be received on. This does not need to be modified, but a linux/unix server may need to register a valid multicast route to receive broadcasts on this address. This can be done by running the following command on the console:

```
sudo route add -net 224.0.0.0 netmask 224.0.0.0 dev eth0
```

This syntax may be different for different flavours of linux / unix. This example uses “eth0” as the network device that will be used for this route.

## Modifications to ROOT.xml

Remove the lines from the ROOT.xml in the Yellowfin/appserver/conf/Catalina/localhost directory:

```
<Manager className="org.apache.catalina.session.PersistentManager" debug="0"
distributable="false" saveOnRestart="false">
  <Store className="org.apache.catalina.session.FileStore" />
</Manager>
```

These lines were previously added to suppress session serialisation to disk in Yellowfin.