



IBM Worklight Foundation

IBM Worklight Foundation V6.2.0

Scalability and Hardware Sizing

17 February 2015

Copyright Notice

© Copyright IBM Corp. 2012, 2014

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Trademarks

IBM, the IBM logo, ibm.com, WebSphere, and Worklight are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at www.ibm.com/legal/copytrade.shtml.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company products or service names may be trademarks or service marks of others.

This document may not be reproduced in whole or in part without the prior written permission of IBM.

About IBM®

See <http://www.ibm.com/ibm/us/en/>.

Contents

1	Architecture	1
2	Scalability Tests	2
2.1	Tests.....	2
2.2	Test Flow.....	2
2.3	Hardware.....	3
2.4	Results	3
3	Using the hardware sizing calculator	4
3.1	About sizing.....	4
4	Push Notification	5
4.1	What is a push notification?.....	5
4.2	Possible architectures for push notifications	5
4.3	Broadcast, tag-based, and unicast (1:1) notifications.....	5
4.4	Testing assumptions.....	6
4.5	Test results.....	6
4.6	Version notes	6
5	Reporting	7
5.1	Database-based reports.....	7
5.2	IBM Worklight Analytics Platform	7
6	Direct Update	9
6.1	What is Direct Update?.....	9
6.2	Direct Update results.....	9
6.3	Direct Update scalable solutions	9
7	Geolocation	10
7.1	What is geolocation?	10
7.2	How geolocation affects the server performance	11
8	Database usage and size	12
8.1	IBM Worklight Foundation database usage profile.....	12
8.2	Database size.....	12
9	Application Center	13
	Appendix A - Notices	14
	Appendix B - Support and comments	16

Figures and Tables

Figure 1-1: Typical IBM Worklight Foundation topology.....	1
Figure 7-1: Location services architecture	10
Table 2-1: Scalability tests performed	2

About this document

This document contains information about scalability-related parameters of the IBM® Worklight® Foundation mobile platform, and the scalability tests performed. It serves as a guide to the included hardware sizing calculator that is used to determine the required hardware for an IBM Worklight Foundation deployment.

This document is not a replacement for the IBM Worklight Foundation user documentation, which describes a Worklight Server architecture, server setup, and daily maintenance requirements.

1 Architecture

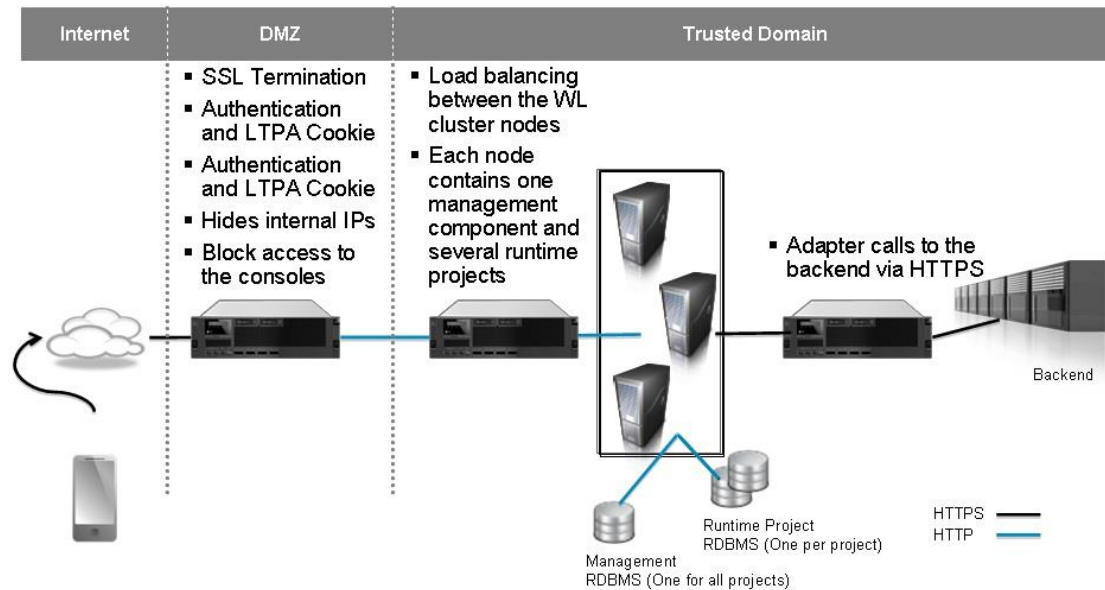


Figure 1-1: Typical IBM Worklight Foundation topology

The IBM Worklight Foundation topology is based on the following principles:

- Worklight Server is installed in the organization LAN, connecting to various enterprise back-end systems.
- Worklight Server can be clustered for high availability and scalability.
- It is possible to install different Worklight project WAR files in the same application server. Each Worklight project configuration must use a different Worklight database or schema. The database is used for storing push notification information, devices information, and statistics for reporting and analytics. A single instance of the database is shared by the Worklight runtime environment in all Worklight Server cluster nodes. It is possible to cluster the database by using the tools and methods that are provided by the database vendor.
- Worklight Server is installed behind a web authentication infrastructure, which acts as a reverse proxy and provides SSL.
- The database server and the Worklight server must be installed on separate operating systems.

Topologies can vary depending on the corporate network architecture, disaster recovery, different back-end systems, and so on.

2 Scalability Tests

IBM conducted a series of scalability tests before releasing IBM Worklight Foundation. The following table contains a short summary of the tests performed.

2.1 Tests

#	Test	Goal
1	Single user performance test	This test creates a performance baseline for future tests from a typical user scenario.
2	Basic stability	This test checks the basic stability and endurance of the server under some load. No crashes, memory leaks, or errors for 12 hours.
3	Regression from former version	This test verifies that the system performs at least as well as the former version.
4	Multiple user performance test	This test checks the performance of a typical user scenario under typical load.
5	Scalability tests	This test proves that IBM Worklight Foundation can linearly scale to multiple servers under a single database.
6	Endurance testing	This test proves that the server can stay up without memory leaks for a duration of 72 hours under a typical load scenario.
7	Stress testing	This test increases the load until the server reaches 100% CPU and starts failing. Reduce the load and see that the server is recovering.
8	Recoverability	This test stops various components of the IBM Worklight Foundation solution (internet, database, and back end). It starts components again and ensures that the server recovered.

Table 2-1: Scalability tests performed

2.2 Test Flow

The test, which is run repeatedly, involves the following actions:

1. A typical banking application is simulated. It logs in, makes multiple calls for account list and transaction lists, and then logs out.
2. The IBM Worklight Foundation application server uses an HTTP adapter to call a simulated back end. On average, a client-server call generates 1.1 requests that are sent to the back end.
3. The back end is called via HTTP and returns with a predefined response.

2.3 Hardware

The following hardware was used for IBM Worklight Foundation application server:

- x3250 M4, Xeon 4C E3-1220v2 69W 3.1GHz/1600MHz/8MB

Note: With the specified number of cores, IBM Worklight Foundation performs thousands of transactions per second. If your throughput is low, you can use fewer cores than the specified number.

- 32 GB 1333 MHz

Note: Not all RAM was used. Use the accompanying calculator to compute required RAM.

- NetExtreme II BCM5790 Gigabit Ethernet
- 2 X 500 3.5in SATA, SR C100

The following hardware was used during testing for the IBM Worklight Foundation database:

- 2 Intel Xeon 8C 2.3 GHz/20 MB
- 20 GB DDR3 RAM
- 6-7 500 GB 7.2 Krpm SATA disks
- 1 RAID 5 controller

IBM Worklight Foundation Client:

- Simulated by twenty virtual machines that run multiple client threads. Running Ubuntu 10.04 64-Bit 4 GB RAM.
- The client load is achieved with Apache jMeter, an open source server performance testing tool by Apache.

2.4 Results

Tests were successful in achieving the goals that are defined in Table 2.1. Specifically, the final test on a single IBM Worklight Foundation server ran for three days, without memory loss or performance loss. Throughput was measured at **2,000** requests per second per server with an average CPU usage of 50%.

3 Using the hardware sizing calculator

3.1 About sizing

The hardware sizing calculator is based on a scenario that was tested in IBM Labs. This scenario is based on a real customer scenario.

However, Worklight Server is a platform that is augmented with server-side developer code. Therefore, the sizing is largely influenced not only by client calls, but also by the efficiency of the server-side code. Hence, you must consider the calculator as a baseline from which you can proceed to customer-specific scalability tests.

This document is accompanied by a [Microsoft Office Excel spreadsheet](#). The purpose of this spreadsheet is to help organizations calculate the required hardware for an IBM Worklight Foundation deployment.

The calculator computes:

- The number of required servers. The server hardware baseline spec is defined in the preceding test.
- The amount of RAM required for each server.
- The Worklight Server average request latency.
- The required database size.
- The push notification throughput.
- The analytics server disk space and the number of servers.

4 Push Notification

4.1 What is a push notification?

A push notification is an alert that indicates a change or an update that the Worklight Server sends to users' mobile devices. Push notifications are popular with all mobile operating systems. IBM Worklight Foundation has APIs that maintain the lists of mobile devices and their tag subscriptions, and send push notifications. IBM Worklight Foundation supports push notifications for the following vendors:

- Android – [GCM](#)
- Apple – [APNS](#)
- Microsoft, Windows Phone – [MPNS](#)

IBM Worklight Foundation does not currently support push notification API for BlackBerry or Windows 8.

4.2 Possible architectures for push notifications

Two architectures are possible for push notifications. They differ by the method that the enterprise back end uses to provide push notification data to the Worklight Server:

- The Worklight Server pulls data by using a Worklight JMS adapter. The enterprise back end puts the push notifications in the JMS queue.
- Push notifications are sent to the Worklight Server by a call to an IBM Worklight Foundation procedure.

For more information about these alternate architectures, see [the IBM Worklight Foundation user documentation](#).

The primary difference between the two methods is that the first method is limited to a single server, which is pulling data from the JMS queue. With the second method, the procedures in the Worklight Server are called through a load balancer. Therefore, the architecture can scale to more than one server.

4.3 Broadcast, tag-based, and unicast (1:1) notifications

In IBM Worklight V6.1.0 and earlier, a device or user had to subscribe to an "event source" to receive push notifications. This is called "event source" based notification. In addition to event-source based notifications, IBM Worklight Foundation V6.2.0 supports broadcast, tag-based, and unicast (1:1) notifications.

- Broadcast notifications are sent to all the registered mobile devices of a specific application.
- Tag-based notifications are sent to all the devices that subscribed to a specific tag.
- Unicast notifications, like event-source based notifications, are sent to the devices of specific users. However, no subscription to an "event source" is required. Instead, the "userID" (or) "deviceID" is specified in the newly introduced server-side API `WL.Server.sendMessage`, which enables to send push notification messages to a specified target.

For more information about these features, see [the IBM Worklight Foundation user documentation](#).

4.4 Testing assumptions

- Push providers were simulated with a delay that is typical to these services:
 - APNS: 0.05-0.1 milliseconds
 - GCM: 20-25 milliseconds
 - MPNS: 20-25 milliseconds
- IBM Worklight Foundation can associate multiple devices per user. On average, each user is assumed to have 1.1 devices.
- With IBM Worklight Foundation V6.2.0, devices can subscribe to tags or topics of interest. Performance of tag-based notifications was tested with tags that had 600,000, 300,000, and 150,000 devices subscribed to them.

4.5 Test results

For broadcast or tag-based push notifications, a Worklight server can push a notification message to over 22,500 iOS devices, 50,000 Android devices, and 350 Windows Phone devices per second. This is with all of the device types together.

For unicast or event-source based push notifications, an instance of Worklight Server can push 1,750 notifications per second to APNS, and 350 notifications per second to GCM and MPNS, when it is under load. The load includes 2,400 regular transactions per second, as well as proportional device calls for subscribing to, and unsubscribing from, the push service.

The CPU and the memory are mostly unaffected for unicast and event-source based push notifications. However, for broadcast and tag-based push, if the number of mobile devices that would register for push notifications is large (for example, more than 100,000), then there could be bursts of 100% CPU usage and higher memory usage for a short duration during which the notification message would be sent to all the registered mobile devices. If your requirements are such that you do not want to delay your mobile users because of push notifications, you might want to have an instance of Worklight Server dedicated to handle all the push notification needs, and configure your back-end server to send all the push notification requests directly to this dedicated server.

The database is unaffected, except for the storage of devices and tag subscriptions, as mentioned in *section 8. Database usage and size*.

4.6 Version notes

The testing data for event-source based push notification that is provided in this document relates to IBM Worklight V6.0.0 and onwards.

The testing data for broadcast, tag-based, and unicast push notification that is provided in this document relates to IBM Worklight Foundation V6.2.0 and onwards.

Customers who use earlier versions of IBM Worklight must be aware of known scalability issues. Some of these issues have workarounds. Therefore, if you intend to use push notification with earlier versions of IBM Worklight, contact IBM customer support or see the [technote about improving performance of IBM Worklight push notifications](#).

5 Reporting

IBM Worklight Foundation contains the following reporting mechanisms:

- Reports that are based on databases.
- Reports that are based on IBM Worklight Analytics Platform.

You can find a comparison between these reporting options in section [Comparison of operational analytics and reports features](#) of the user documentation.

Note: These reports are not enabled by default. Customers must decide which of these mechanisms they want to use.

5.1 Database-based reports

Database-based reports are the older Worklight reports. These reports use a single table to store large quantities of raw data in the database. This table is then aggregated periodically to produce reports.

Sizing for these reports is provided in section 8, *Database usage and size*.

5.2 IBM Worklight Analytics Platform

IBM Worklight Analytics Platform is a WAR file that is bundled with Worklight Server. The IBM Worklight Analytics Platform is based on file system storage rather than database storage.

The hardware that is required to run the IBM Worklight Analytics Platform relies on several factors including:

- The frequency of connectivity to the server
- The verbosity set for logging
- The number of days that data is kept before being deleted

5.2.1 Data types

There are three distinct data types:

- Analytics collected from server network connections
- Client Logs
- Server Logs

The analytics data is always collected and forwarded to the IBM Worklight Analytics Platform.

Worklight Server produces debug logs that are optionally forwarded to the IBM Worklight Analytics Platform. Client logs are sent when the client app is optionally configured to collect and send logs.

An event of network connection to the analytics server is recorded in the following cases:

- A device successfully authenticates with the Worklight Server instance.
- A device makes an adapter invocation.
- A push notification is sent to a device.
- A client-side API is called to create a custom activity.

Server network connections are always collected and forwarded to the IBM Worklight Analytics Platform.

Server network connections occur in the following situations:

- A device successfully authenticates with the Worklight Server instance.
- A device makes an adapter invocation.
- A push notification is sent to a device.
- A client-side API is called to create a custom activity.

Server network connections typically take up most of the disk space in the IBM Analytics Platform. Device authentications and adapter invocations are the most common types of network connections that are recorded in the Analytics Platform.

When you estimate hardware requirements, it is important to take into account the number of devices that will be connecting to your server and the number of adapter invocations each device will be making.

Client logs are sent when the client app is optionally configured to collect and send logs.

By default, client logs are set to “FATAL”, which means that only crash logs are forwarded to the Worklight Server instance. With this default setting, client logs use less space than analytics data that are created by server network connections.

Changing this level and adding custom client logs drastically affects the disk space used by client logs. With more custom logs and higher verbosity log levels, client logs can consume much more disk space (by several magnitudes) than server network connections. When you estimate hardware requirements, you must take into consideration the size client logs and the frequency of sending these logs.

Server logs, which are produced by Worklight Server and are saved to disk, are also forwarded to the Worklight server by default.

Server logs will have the least impact on the disk space sizing.

5.2.2 Purging data

On IBM Worklight Analytics Platform, you can configure automatic purging of data. For more information, see the IBM Worklight Foundation [user documentation](#). Different data types can be purged at different intervals. This means that if you wish to keep a year's worth of server network event data but only three weeks of log data, you can set these values separately.

Log data takes up more space than server network event data. Depending on the configured verbosity, Worklight Server could be sending a very large amount of log data. If disk space is limited, you must consider purging data more frequently or adding nodes to your IBM Worklight Analytics Platform cluster.

5.2.3 Clustering

You can add a new node to the analytics cluster with zero downtime. Therefore, you might want to start with a single node and add nodes as necessary. The total number of nodes that you can add to the cluster is equal to the number of shards that were created when the first node was installed. For more information about shards and how to set up a cluster, see the IBM Worklight Foundation [user documentation](#).

6 Direct Update

6.1 What is Direct Update?

When the administrator redeploys an app to the Worklight Server instance without changing its version, Worklight Server updates the web resources folder. Then, when an app connects to Worklight Server, Worklight Server directly pushes the web resources (HTML, JavaScript™, and CSS) of the newly deployed app to the device. The push happens only when an app contains different web resources and the same native code.

6.2 Direct Update results

An instance of Worklight Server can push 250 MB/s. For example, if the application size is 5 MB, Worklight Server can serve 50 direct updates per second. This means that if the cluster contains four Worklight Server nodes, the cluster can serve 200 direct updates per second.

Note: You must validate that the network bandwidth is not the bottleneck.

6.3 Direct Update scalable solutions

The best scalable solution for Direct Update is to use a content delivery network (CDN) in front of the Worklight Server instance. When caching the direct update URL, the request is served from the CDN and does not get to Worklight Server. The expiration policy must be specified in the CDN. For more information about Direct Update, see the [Serving direct update requests from CDN](#) blog post on developerWorks.

7 Geolocation

7.1 What is geolocation?

Location services in IBM Worklight Foundation provide you with the opportunity to create differentiated services that are based on a user location, by collecting data about geolocation and WiFi, and by feeding the location data and triggers to business processes, decision management systems, and analytics systems.

Application code on the mobile device controls the collection of data from device sensors. The collected data is referred to as the *device context*. When a change occurs in the device context, such as a change in the geolocation of the device, or the fact that it has just entered a WiFi zone, triggers can be activated. The triggers specify that an action should occur: either a callback function is called, or an event is sent to the server, based on the device context.

Events are created by triggers and application code and include a snapshot of the device context at the time of their creation. Events are buffered on the client and are transmitted to the server periodically.

To handle the events, the server uses IBM Worklight Foundation adapter application code. This code sets up event handlers on the server, which filter event data and pass matching events to a callback function. The code also accesses the device context of the client, that is to say its location, and WiFi network information. It then sets an application context. Server activities and received events are logged, together with the device and application contexts, for future reporting and analytics.

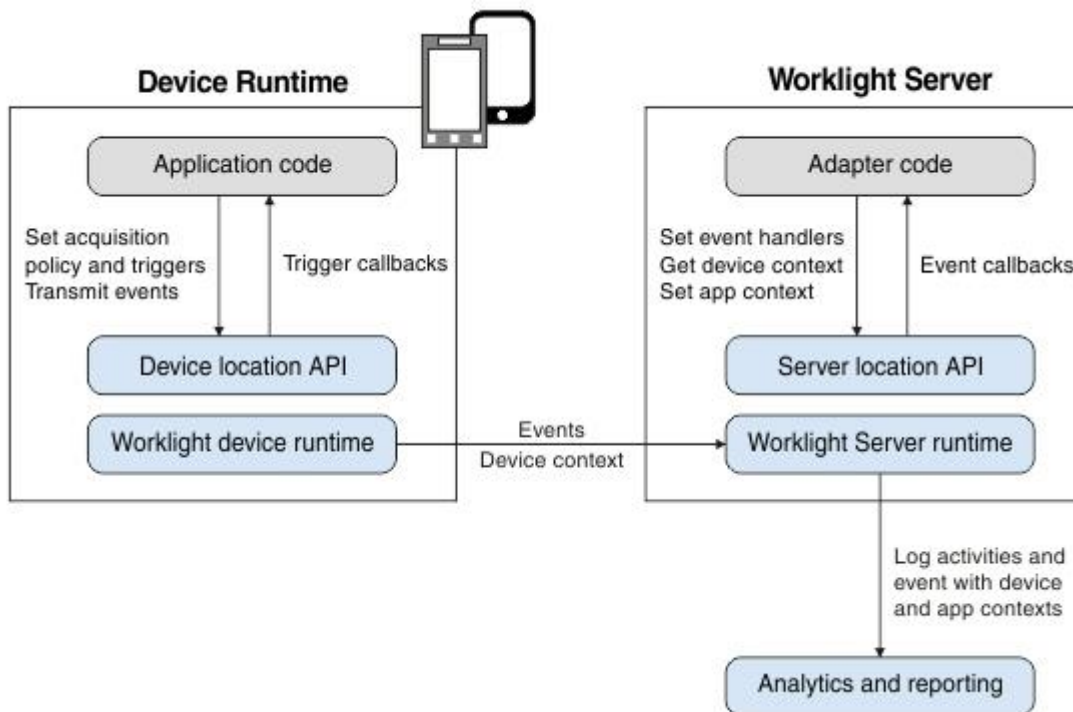


Figure 7-1: Location services architecture

7.2 How geolocation affects the server performance

7.2.1 Analytics and reports

Additional entries go to analytics and reports. Entries are added each time an event handler is run.

7.2.2 Additional data passed when `invokeProcedure` is called

Additional data is passed when the `invokeProcedure` method is called from the device, to transfer the device context of the client to the server. The amount of data that is transferred is usually in the tens of bytes. This amount of data might be larger, depending on how the WiFi acquisition of the device is configured and what WiFi access points are visible to it. Events are also transmitted occasionally, according to the event transmission of the client.

7.2.3 Adapters implementing event handlers

One or more event handlers might be run each time a geolocation event is received.

8 Database usage and size

8.1 IBM Worklight Foundation database usage profile

An IBM Worklight Foundation project is a set of apps, adapters, and project-specific configuration. A project is associated with a WAR file.

IBM Worklight Foundation requires one database schema per Worklight runtime environment, which is called `WRKLGHT`, and one database schema for all Worklight runtime environments, which is called `WLADMIN`. The `WLADMIN` database stores all the management data, for example the applications and adapters, and their state. Optionally, you can add a database schema for database-based reports, which is called `WLREPORT`. The reports database is used only with the classic IBM Worklight Foundation reports. For more information, see section 5, *Reporting*.

In addition, when the Application Center is used, one more database schema, called `APPCNTR`, is used. This database schema is global, which means that it is not project-dependent.

For Oracle and DB2, the default configuration uses three schemas instead of three actual database instances.

The IBM Worklight Foundation database is rather small, with a number of exceptions:

- Push notification information: the list of devices that registered for push notifications, and the subscriptions from each device.
- SSO state information: SSO (Single Sign-On) is the ability for business customers to have multiple apps from the same organization. When they log in with one app, they automatically log in to the other device. Turning on the SSO feature puts a larger load on the database because SSO state information is read and updated to the database at the start of each client access to the server.
- The Reports database contains a single table that is called `APP_ACTIVITY_REPORT`, which is append-oriented. Each transaction that is called from a mobile device generates an `INSERT` statement into this single table. This data feed accumulates a large amount of data on what the user does. The data is aggregated at specific intervals. The default interval is 20 minutes.
- Device licensing is used to count licenses and control user access to applications. This feature is used with B2E customers (internal enterprise apps) but is enabled by default for all customers and must be explicitly disabled for B2C. The feature uses a few more database queries and stores additional information in the database.

8.2 Database size

Database size is determined almost solely by the size of the data feed that is described earlier. The hardware sizing calculator attempts to help you determine the expected size of this feed, as well as push notifications and SSO.

Notes on the report tables:

- The application developer might also choose to add custom log messages, in which case the table size becomes bigger.
- The size of the aggregated tables is negligible when compared to the raw data.

Important note: It is the responsibility of the customer to purge reporting data regularly, typically every 1-2 days.

9 Application Center

The Application Center is the IBM Worklight Foundation application store. The concept of the Application Center is similar to the concept of the Apple public App Store or the Google Play Store, except that it targets only private usage within a company. By using the Application Center, users from the same company or organization download applications to mobile phones or tablets from a single place that serves as a repository of mobile applications.

The Application Center and Worklight Server can be installed on different systems. This installation choice ensures that the setup of the Application Center app store does not affect the performance of IBM Worklight Foundation applications on the server, when it is installed on a different hardware.

The performance of the Application Center is mainly tied to the quality and bandwidth of the network, and to the number of users. The distribution of large mobile applications must be done preferably through 3/4G or WiFi networks. On low networks, the Application Center client tries to solve problems that are related to timeout. If the Application Center cannot reach the server, it tries again several times. The download of an application might also be interrupted when the connection to the network is lost. In such case, the download can restart from where it stopped (for Android, iOS).

The Application Center requires one database schema to store the application catalog and the user information. The size of the database depends mainly on the average size of the applications that are stored in the catalog and the number of users that are connected to the system. This is reflected in the Hardware Calculator document.

Appendix A - Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Dept F6, Bldg 1
294 Route 100
Somers NY 10589-3216
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

Privacy Policy Considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect session information (generated by the application server). These cookies contain no personally identifiable information and are required for session management. Additionally, persistent cookies may be randomly generated to recognize and manage anonymous users. These cookies also contain no personally identifiable information and are required.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent. For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Appendix B - Support and comments

For the entire IBM Worklight Foundation documentation set, training material and online forums where you can post questions, see the IBM website at:

<http://www.ibm.com/mobile-docs>

Support

Software Subscription and Support (also referred to as Software Maintenance) is included with licenses purchased through Passport Advantage and Passport Advantage Express. For additional information about the International Passport Advantage Agreement and the IBM International Passport Advantage Express Agreement, visit the Passport Advantage website at:

<http://www.ibm.com/software/passportadvantage>

If you have a Software Subscription and Support in effect, IBM provides you assistance for your routine, short duration installation and usage (how-to) questions, and code-related questions. For additional details, consult your IBM Software Support Handbook at:

<http://www.ibm.com/support/handbook>

Comments

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this document. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state.

Thank you for your support.

- Submit your comments in the IBM Worklight Foundation forums at:
<https://www.ibm.com/developerworks/mobile/worklight/connect.html>

If you would like a response from IBM, please provide the following information:

- Name
- Address
- Company or Organization
- Phone No.
- Email address

