

Enterprise Caching in a Mobile Environment

IBM Redbooks Solution Guide

In the current global enterprise business environment, with the millions of applications running across Apple iOS, Android, Windows Phone, and other platforms, mobile marketing brings more business opportunities for value added services but also brings more challenges to existing systems. Data is being accessed at a rapidly increasing rate to service mobile applications and the workload is becoming a burden to back-end servers. This guide shows how you can use IBM® WebSphere® eXtreme Scale or the WebSphere DataPower® XC10 Appliance to cache frequently used data, thereby improving the performance of mobile applications and reducing the traffic to back-end applications.

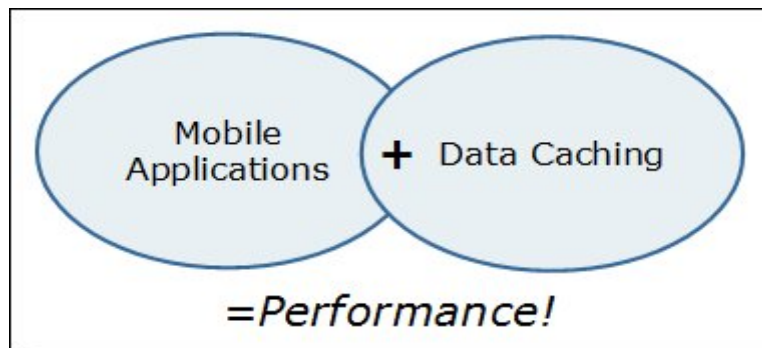


Figure 1. Mobile applications benefit from caching data

Did you know?

By the end of 2012, there were over 1 billion smart phones in use worldwide. In the next three years, that number is expected to reach 2 billion.

Business value

Similar to traditional applications, mobile applications need to cache large amounts of data in memory, which can affect performance. Yet the requirements for better performance, high availability, and linear scalability are critical to mobile services. Improving the user experience is the key issue faced by mobile businesses.

IBM Worklight is part of the IBM Mobile Foundation family and provides an open, comprehensive, and advanced mobile application platform. WebSphere eXtreme Scale is one of the caching infrastructures for the WebSphere family of products that has the ability to use distributed object caching to manage application data and business logic for multiple servers. By integrating these elastic cache capabilities of WebSphere eXtreme Scale with IBM Worklight, you can accelerate mobile businesses and provide a better experience for the user.

Solution overview

Mobile applications are often unique among application types. With mobile applications, extremely high and unpredictable workloads can occur because the applications can be accessed at any time and from anywhere. There can be multiple, frequent peak times. Consider also that the data that is exchanged is small in size, but frequently requested. In this type of scenario, back-end systems can become a bottleneck when the workload is heavy. Caching can relieve the pressure on the back-end system. In this solution, a mobile application caches data that is retrieved from back-end applications in order to respond to subsequent requests for the same data. This increases response time to the user and decreases the number of requests to the back-end server.

WebSphere eXtreme Scale servers and Worklight servers utilize WebSphere Application Server V8.5 Liberty profile servers. The Liberty profile is a composable and lightweight runtime application server; WebSphere eXtreme Scale features can be enabled quickly and easily using the Liberty profile. Without complex configuration and with only minimal code changes, mobile applications can benefit from the WebSphere eXtreme Scale cache seamlessly.

This solution uses the WebSphere eXtreme Scale client to access remote WebSphere eXtreme Scale grids that are then invoked by the Worklight HTTP adapter. As such, this solution focuses on using the WebSphere eXtreme Scale as a side cache.

Solution architecture

An overview of the solution architecture is shown in Figure 2.

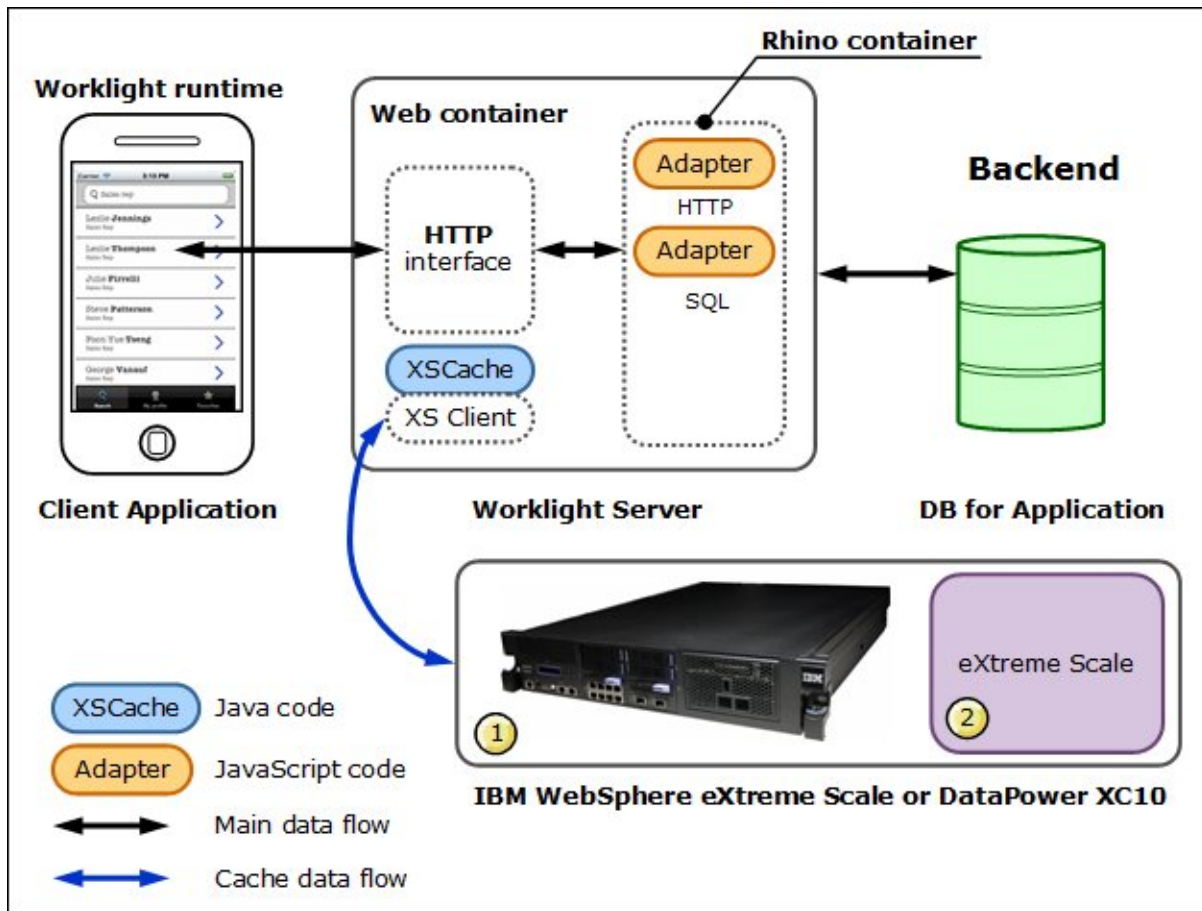


Figure 2: Solution architecture

This solution includes the following products:

- WebSphere Application Server V8.5 and the Liberty profile
- WebSphere eXtreme Scale or DataPower XC10
- IBM Worklight
- IBM Worklight Studio

Building the cache

The cache is implemented by defining a WebSphere eXtreme Scale grid that spans multiple Liberty profile servers. One set of servers runs the WebSphere eXtreme Scale catalog and the other set runs the grid containers. The container servers hold the application data stored in the cache. A catalog server provides administrative services to the container servers in its domain. The catalog server manages the placement of the data in the cache and keeps track of where the data is. When a request is made for data from the cache, the catalog server directs the requester to the proper location. Multiple catalog servers are used for failover purposes.

Each Liberty profile server is configured using simple XML and property files that define the features that your applications will use, including the WebSphere eXtreme Scale server feature and the location of the catalog servers. The WebSphere eXtreme Scale cache on these servers is also configured using simple XML files. When the Liberty profile servers are started, the WebSphere eXtreme Scale configuration files are detected and the caching grid is made available for use by applications.

Figure 3 shows a sample topology, with three catalog servers and four grid containers. Use one Worklight server as shown, or you can scale out to use a Worklight cluster for your production system.

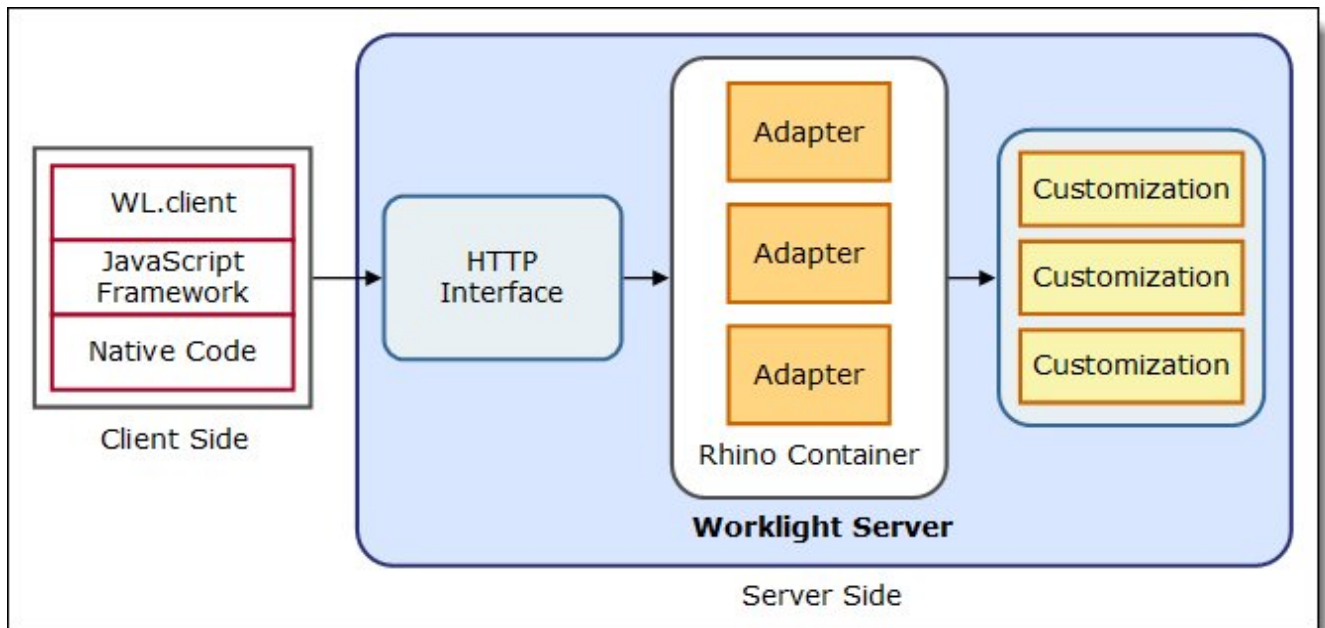


Figure 3: Components of a Worklight application

Building the IBM Worklight server on a Liberty profile server

Beginning with V5.0, Worklight server ships with a Liberty profile, making implementation simple. The Liberty profile server is configured with an XML file that includes information about the Worklight server application and the Worklight application center. The XML file also enables the WebSphere eXtreme Scale client feature, giving applications on the server access to the cache.

When the server is started, you can access the Worklight console shown in Figure 4.

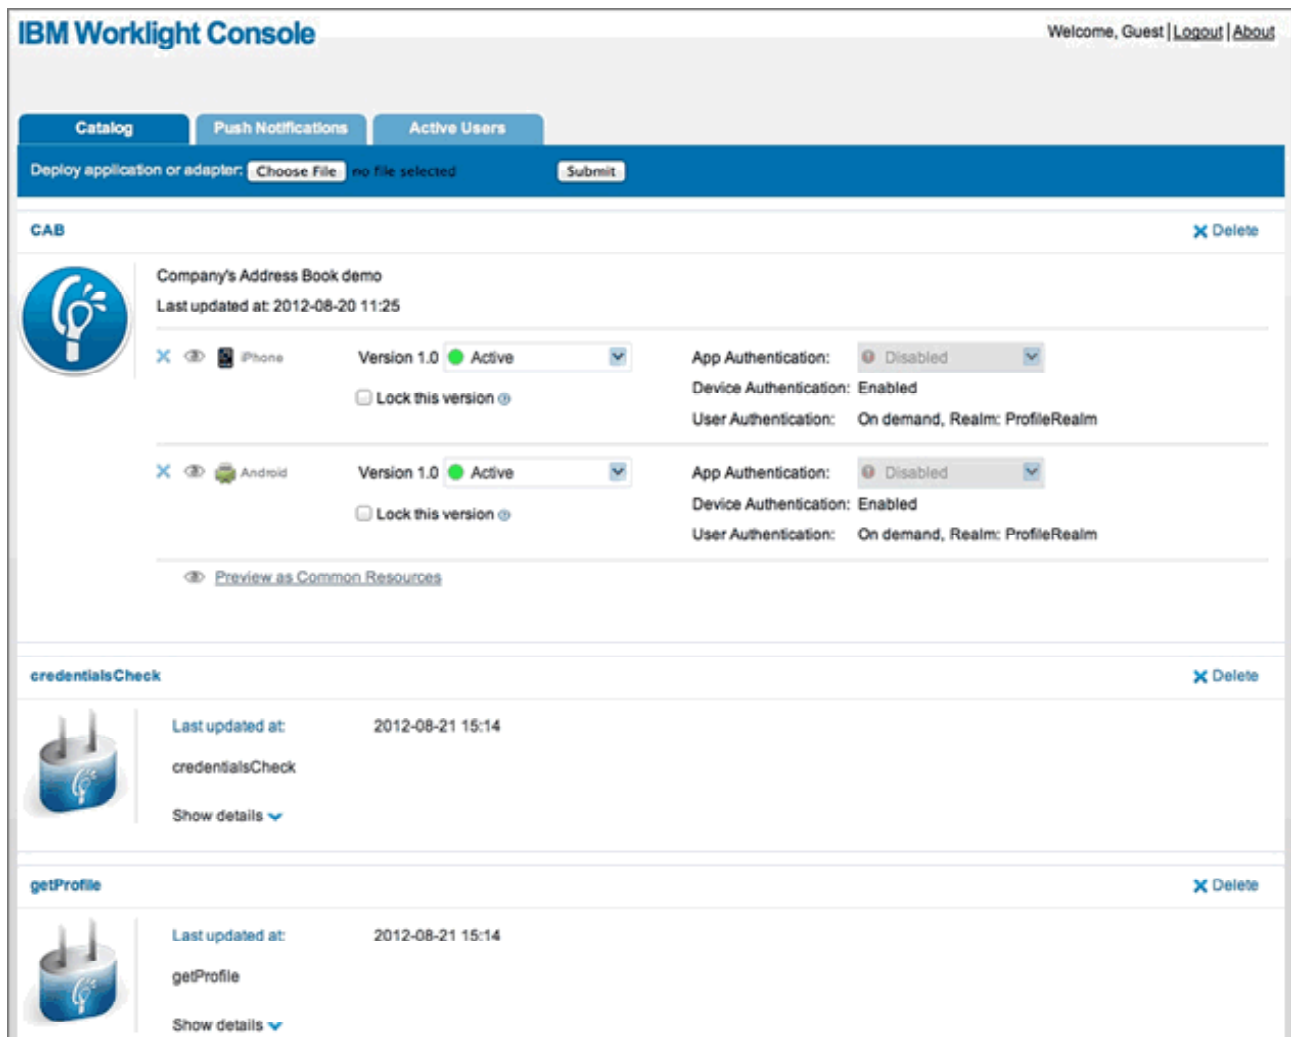


Figure 4: Worklight console

You can perform the following tasks on the Worklight console:

- Deploy a Worklight application
- Deploy a Worklight adapter
- Preview and test your Worklight server-side application with a browser

Using the cache from mobile applications

IBM Worklight provides Worklight Studio, an Eclipse-based IDE that supports the development of native, web, or hybrid applications for iOS, Android, Windows Mobile, and BlackBerry. Applications can be created that use JavaScript, Ajax, HTML5, and CSS.

Worklight Studio contains an embedded Worklight server for development. To develop a mobile application with Worklight Studio:

1. Create a Worklight project and develop your mobile application. Typically, the logic is implemented in the JavaScript code and is common for all mobile platforms. Calls to the cache are coded into the application.

2. Generate the artifacts for a specific device.
3. Build and deploy your application for testing, selecting a target device (for example, an iphone).
4. Run and test your mobile application with a simulator.

A typical Worklight application topology looks similar to that shown in Figure 5 and includes a 3-tier architecture with a client, a Worklight Server, and a back-end system, which might be a database, a service bus, and so on.

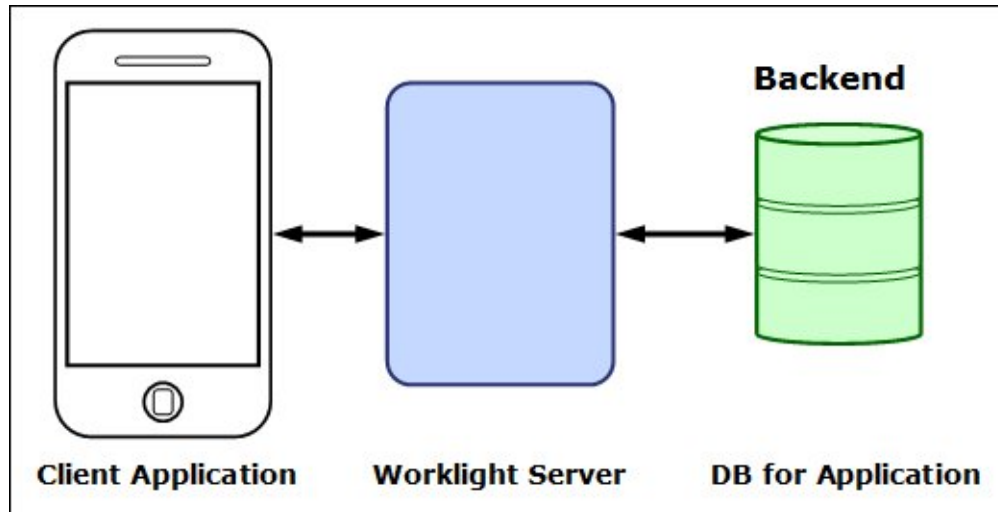


Figure 5: Worklight application topology

In this solution, the client communicates with the Worklight server through the HTTP protocol. The Worklight server uses JDBC to communicate with a back-end database.

For the client side, web technologies are used to develop the Worklight client application. It is deployed on the consumer device (such as a smartphone or tablet). Worklight provides a JavaScript framework to facilitate the development. Most client logic is implemented in JavaScript, which also allows the generation of unique native assets for various mobile platforms, such as iOS, BlackBerry, Windows Mobile, and Android. Developers can also take advantage of native device capabilities, such as camera, GPS, accelerometer, and so on.

On the server side, Worklight provides an interface to handle the incoming HTTP requests. If the application needs to talk with a back-end system, you can use a Worklight adapter. The adapters can also call server-side Java code; this solution uses this capability.

Usage scenarios

Consider a simple mobile application called Company Address Book (CAB). From the main screen, the user can search a company address book. Each search result is a list of records that satisfies the search criteria. The user can then drill down into a person's data, including a photo of the person. The application obtains all of this data from the company's back-end database. Search results are cached. For example, many people can search for sales reps in the organization or for the CEO. Caching the results of these searches eliminates the need to query the database again for corresponding requests.

The CAB application includes generic code that is common to all platforms and specific code for iOS and Android. The application also contains a server-side adapter for connectivity and a search adapter. All

requests from a device go to the search adapter, which parses requests and is invoked by other adapters on behalf of the user.

To use caching, a customization class is created so that every request for data passes from the search adapter through the class. The class checks the cache for the requested data, and, if found, the data is returned from the cache. Otherwise, the class invokes another adapter to retrieve the data from the database, stores it in the cache, and returns the data to the search adapter. Then the data is sent back to the client device. This customization class works either for WebSphere eXtreme Scale or DataPower XC10. All configuration data is stored in Worklight configuration files and the cache server addresses can be changed dynamically without recompiling the code.

The data flow is shown in Figure 6.

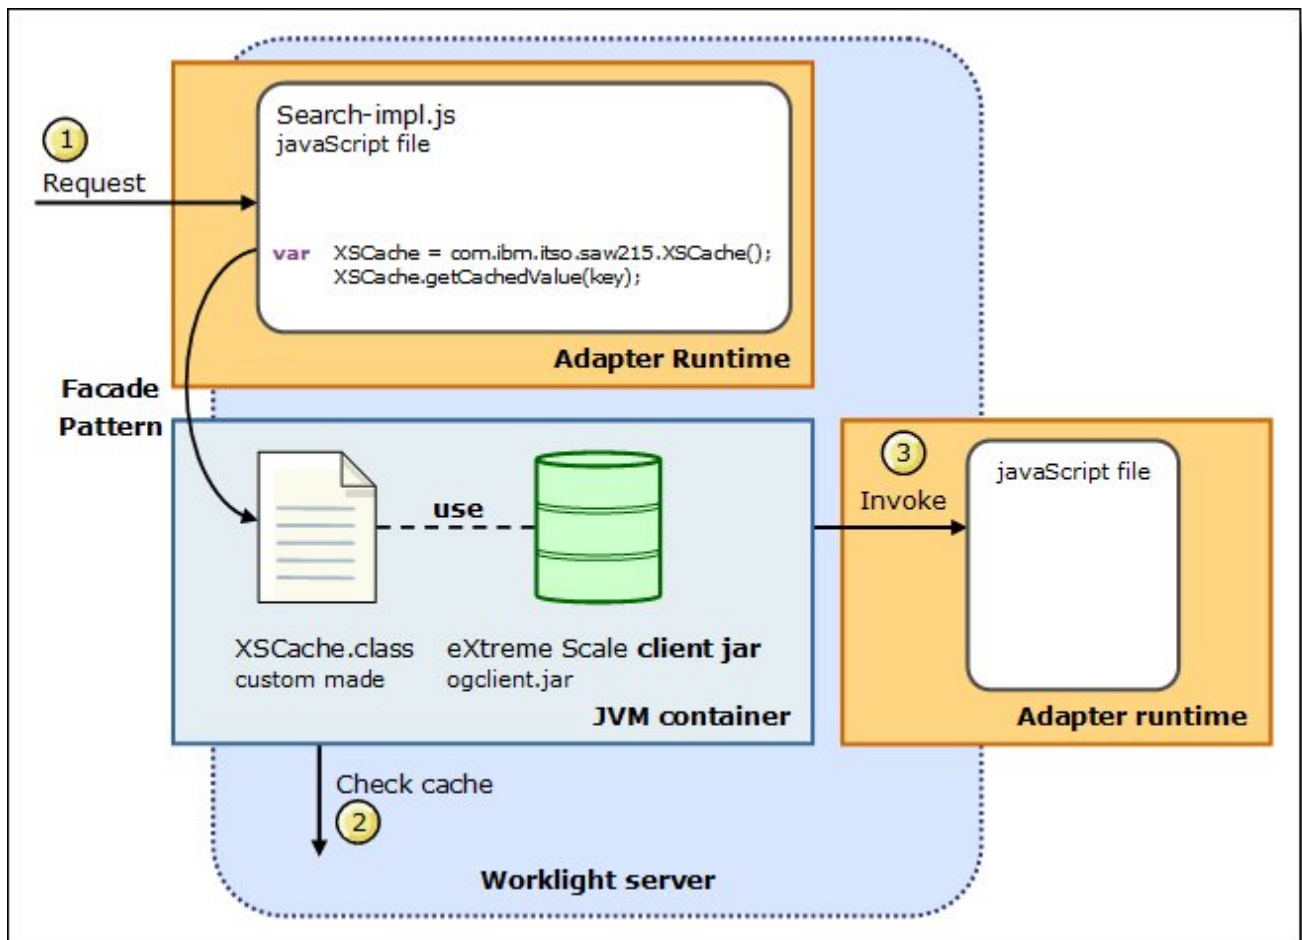


Figure 6: Application workflow with WebSphere eXtreme Scale caching

Integration

An alternative to using WebSphere eXtreme Scale server running on a Liberty profile, is the use of the DataPower XC10 as a cache. Among other features, the XC10 can be used as a simple data grid. In this solution, you would simply define the cache on the XC10 and when you configure your Worklight server for the WebSphere eXtreme Scale features, you would point to the XC10 catalog and cache.

Supported platforms

IBM WebSphere eXtreme Scale is supported on AIX®, HP, Linux, Linux on System z®, Solaris, Windows, z/OS® and z/OS.e. Details about platform support can be found at <http://www-01.ibm.com/software/webservers/appserv/extremescale/sysreqs/index.html>.

IBM Worklight Server is supported on AIX, HP, Linux, Solaris, and Windows. IBM Worklight Studio is supported on Linux, Mac OS, and Windows. Details about platform support for both IBM Worklight and Worklight Studio can be found at <http://www-01.ibm.com/support/docview.wss?uid=swg27024838>.

Ordering information

Ordering information is shown in the following table.

Table 1. Part numbers and ordering information

Program name	PID number	Ordering information
WebSphere eXtreme Scale V8.6 (available only via IBM Passport Advantage®)	5724-X67	http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=897/ENUS212-370&infotype=AN&subtype=CA&appName=wwwsearch#h2-ordinfx
WebSphere eXtreme Scale V8.6 for z/OS (available only via Passport Advantage)	5655-VX8	http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appName=gpatteam&supplier=897&letternum=ENUS212-367#h2-ordinfx
WebSphere Application Server V8.5 (available only via Passport Advantage)	5724-J08	http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS212-109&appName=USN#h2-ordinfx
WebSphere Application Server V8.5 Network Deployment (available only via Passport Advantage)	5724-H88	http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS212-109&appName=USN#h2-ordinfx
WebSphere Application Server for z/OS V8.5	5655-W65	http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/897/ENUS212-106/index.html&lang=en&request_locale=en#h2-ordinfx
IBM WebSphere DataPower XC10 Appliance V2.1	Machine type: 7199 Model: 92X	http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=897/ENUS112-075&infotype=AN&subtype=CA&appName=wwwsearch#h2-prodnox
IBM Mobile Foundation Enterprise Edition V5.0.5 (available only via Passport Advantage)	5725-I41	http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/8/897/ENUS212-368/index.html&lang=en&request_locale=en#h2-ordinfx
IBM Worklight V5.0.5	5725-I43	http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/8/897/ENUS212-368/index.html&lang=en&request_locale=en#h2-ordinfx

Related information

For more information, see the following documents:

- IBM Redbooks: Enterprise Caching Solutions using WebSphere DataPower SOA Appliances and WebSphere eXtreme Scale, SG24-8043
<http://www.redbooks.ibm.com/abstracts/sg248043.html>
- WebSphere eXtreme Scale V8.6 Sales Manual
http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/7/897/ENUS5724-X67/index.html&lang=en&request_locale=en
- WebSphere eXtreme Scale home page
<http://www-01.ibm.com/software/webservers/appserv/extremescale/>
- WebSphere DataPower XC10 Appliance Sales Manual
http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/3/897/ENUS9235-h03/index.html&lang=en&request_locale=en
- WebSphere DataPower XC10 Appliance home page
<http://www-01.ibm.com/software/webservers/appserv/xc10/>
- IBM Worklight home page
<http://www-01.ibm.com/software/mobile-solutions/worklight/>
- IBM Mobile Foundation home page
<http://www-01.ibm.com/software/mobile-solutions/>

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 give 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.

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. 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. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

© Copyright International Business Machines Corporation 2012. All rights reserved.

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

This document was created or updated on December 27, 2012.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips0953.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
DataPower®
IBM®
Passport Advantage®
Redbooks (logo)®
System z®
WebSphere®
z/OS®

The following terms are trademarks of other companies:

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

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

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

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