Tuning Mobile on IBM z Systems with Linux in z/VM

Session ID: 17775

Wilhelm Mild
IBM Executive IT Architect
IBM Laboratory Germany
wilhelm.mild@de.ibm.com

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

Copyright (c) 2015 by SHARE Inc. Except where otherwise noted, this work is licensed under http://creativecommons.org/licenses/by-nc-sa/3.0/

What about the mainframe?

The mainframe…

- Home to business critical applications and data

- How do we bridge the gap?

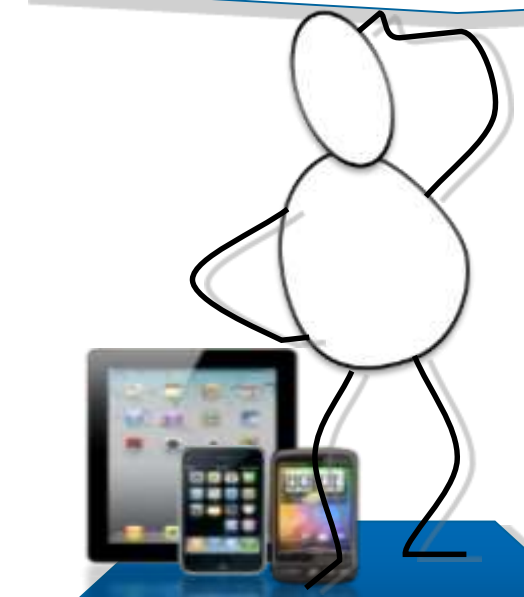

Mobile is changing the way information is used

Mobile First

Information developed using multiple platforms and transformed into web services

Information restricted and developed in the data center

IBM MobileFirst Enterprise Blueprint

A Guideline to Defining Your Optimized MobileFirst Strategy

Mobile Transformation

Transform the business by creating new opportunities

Extend existing business capabilities to mobile devices

Relevant Mobile Scenarios

Business Scenarios - B2C, B2E, B2B, M2M – Industries – Target User Groups

Mobile Device Selection

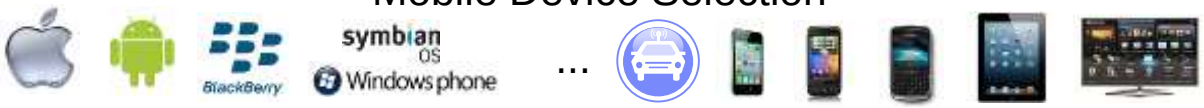

Mobile Application & Data Platform

Build, test and maintain mobile applications

Connect to backend services and data in support of mobile

Management

Manage mobile devices and applications

Security

Secure my mobile applications and business

Analytics

Optimize my mobile applications and business

Mobile Infrastructure

Deploy and operate on-premise, Cloud, Managed Services

Scenario based Discovery and Architecture Definition, Leading to an Optimized Mobile Strategy

Mobile First - Mobile Applications are Not Miniature PC applications…

	 Mobile Applications	 PC Applications
Usage Context	• User may be in the middle of some other activity (e.g. shopping in a supermarket) • Interactions are short and may be interrupted • Users are very impatient	• Using the application is the primary activity • Interactions are longer and more focused • Users are impatient
Mode of Interaction	• Non-keyboard: touch prevalent, also speech • Typing should be minimized • Screen size/real-estate is small	• Keyboard and mouse • Typing is okay • Larger screen size for presenting information
Other considerations	• Integration with device capabilities (e.g. camera, GPS, accelerometer) • Offline behavior	

Multi-tier Mobile Apps - Specific Challenges

Client Tier Devices

Middle Tier Server

Back-end Data & Services

Mobile-specific challenges:

- ✓ Lots of device targets
- ✓ Provisioning rules and artifacts
- ✓ Curated App Stores
- ✓ Dependent upon backend service versions

The Mobile-specific challenges are mainly:

1.Dealing with the **specific issues in the Mobile Client tier**

2.And subsequently **coordinating separate pipelines** for each tier:

- **Mobile Client**
- Middleware
- Back-end data and services

z Systems - Bridge Systems of Record and Systems of Engagement

Systems of Engagement

Systems of Engagement are cloud-based, decentralized, support rapid app development

Systems of Record

Systems of Record are well integrated, trusted repositories

Existing Web Apps

Mobile Apps

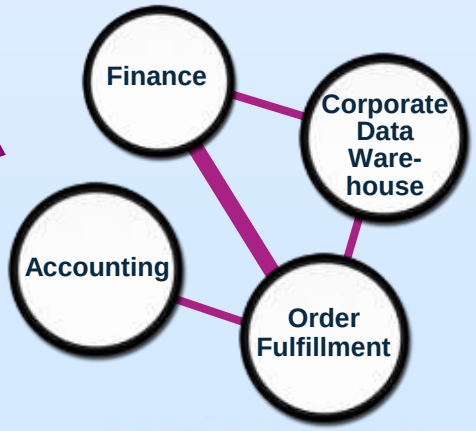

SHARE in Orlando 2015

Mobile Reference Architecture Overview Diagram for z Systems

Mobile Device

Mobile Applications

MobileFirst Studio and Rational Developer for z Systems

Security Layer 1

Caching Service 1 (Optional)

Linux on z Systems

WebSphere Application Server on Linux on z

MobileFirst Server

MobileFirst App Code

MobileFirst Server and Console

MobileFirst Adapters

DB

Caching Service 2 (Optional)

Security Layer 2

z/OS

CICS, IMS

DB2

TPF

z/VSE

z/VM

IBM Information Bus (IIB)

Mobile application overview diagram

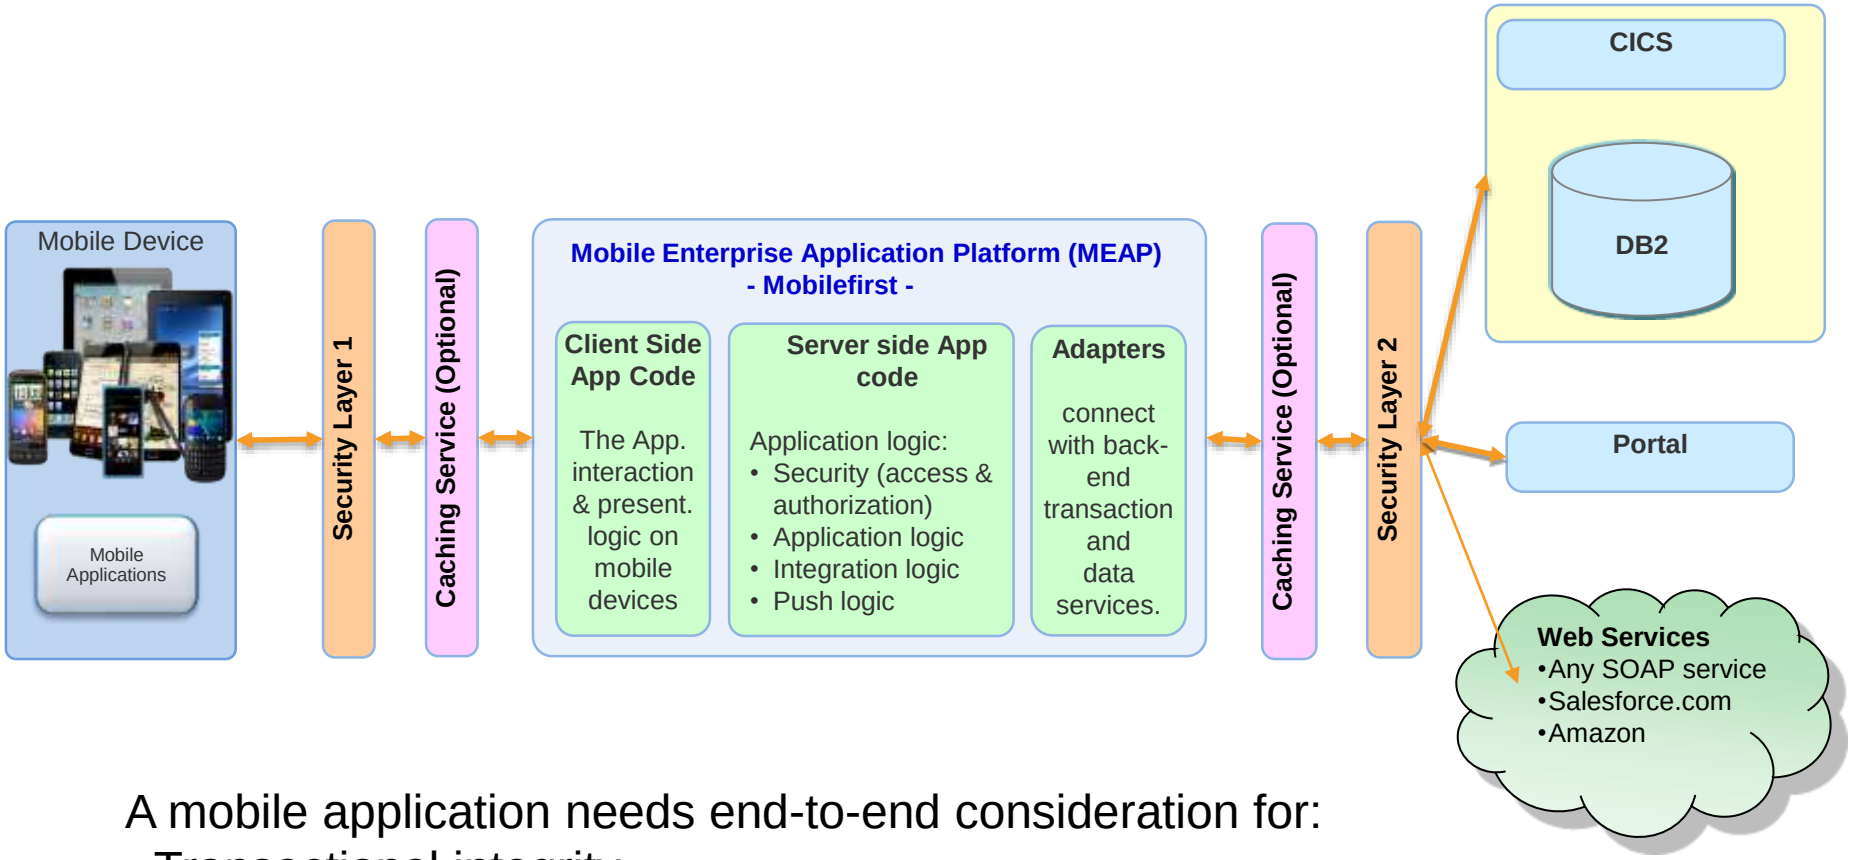

A mobile application needs end-to-end consideration for:
- Transactional integrity
- Data integrity
- Security

Mobile application integration is realized with MobileFirst Adapters

Tuning a MobileFirst environment on Linux on z

- Implementations for increased End-to-End Performance
- z/VM and Linux OS considerations
- WebSphere for MobileFirst
- MobileFirst database access
- Adapter connectivity
- End-to-end security considerations

Tuning a MobileFirst environment on Linux on z

- Implementations for increased End-to-End Performance
  - Consider a flexible implementation
  - Consider HA and DR
- z/VM and Linux OS considerations
- WebSphere for MobileFirst
- MobileFirst database access
- Adapter connectivity
- End-to-end security considerations

Linux on z and IBM MobileFirst Server – A Cloud environment on z Systems for mobile devices

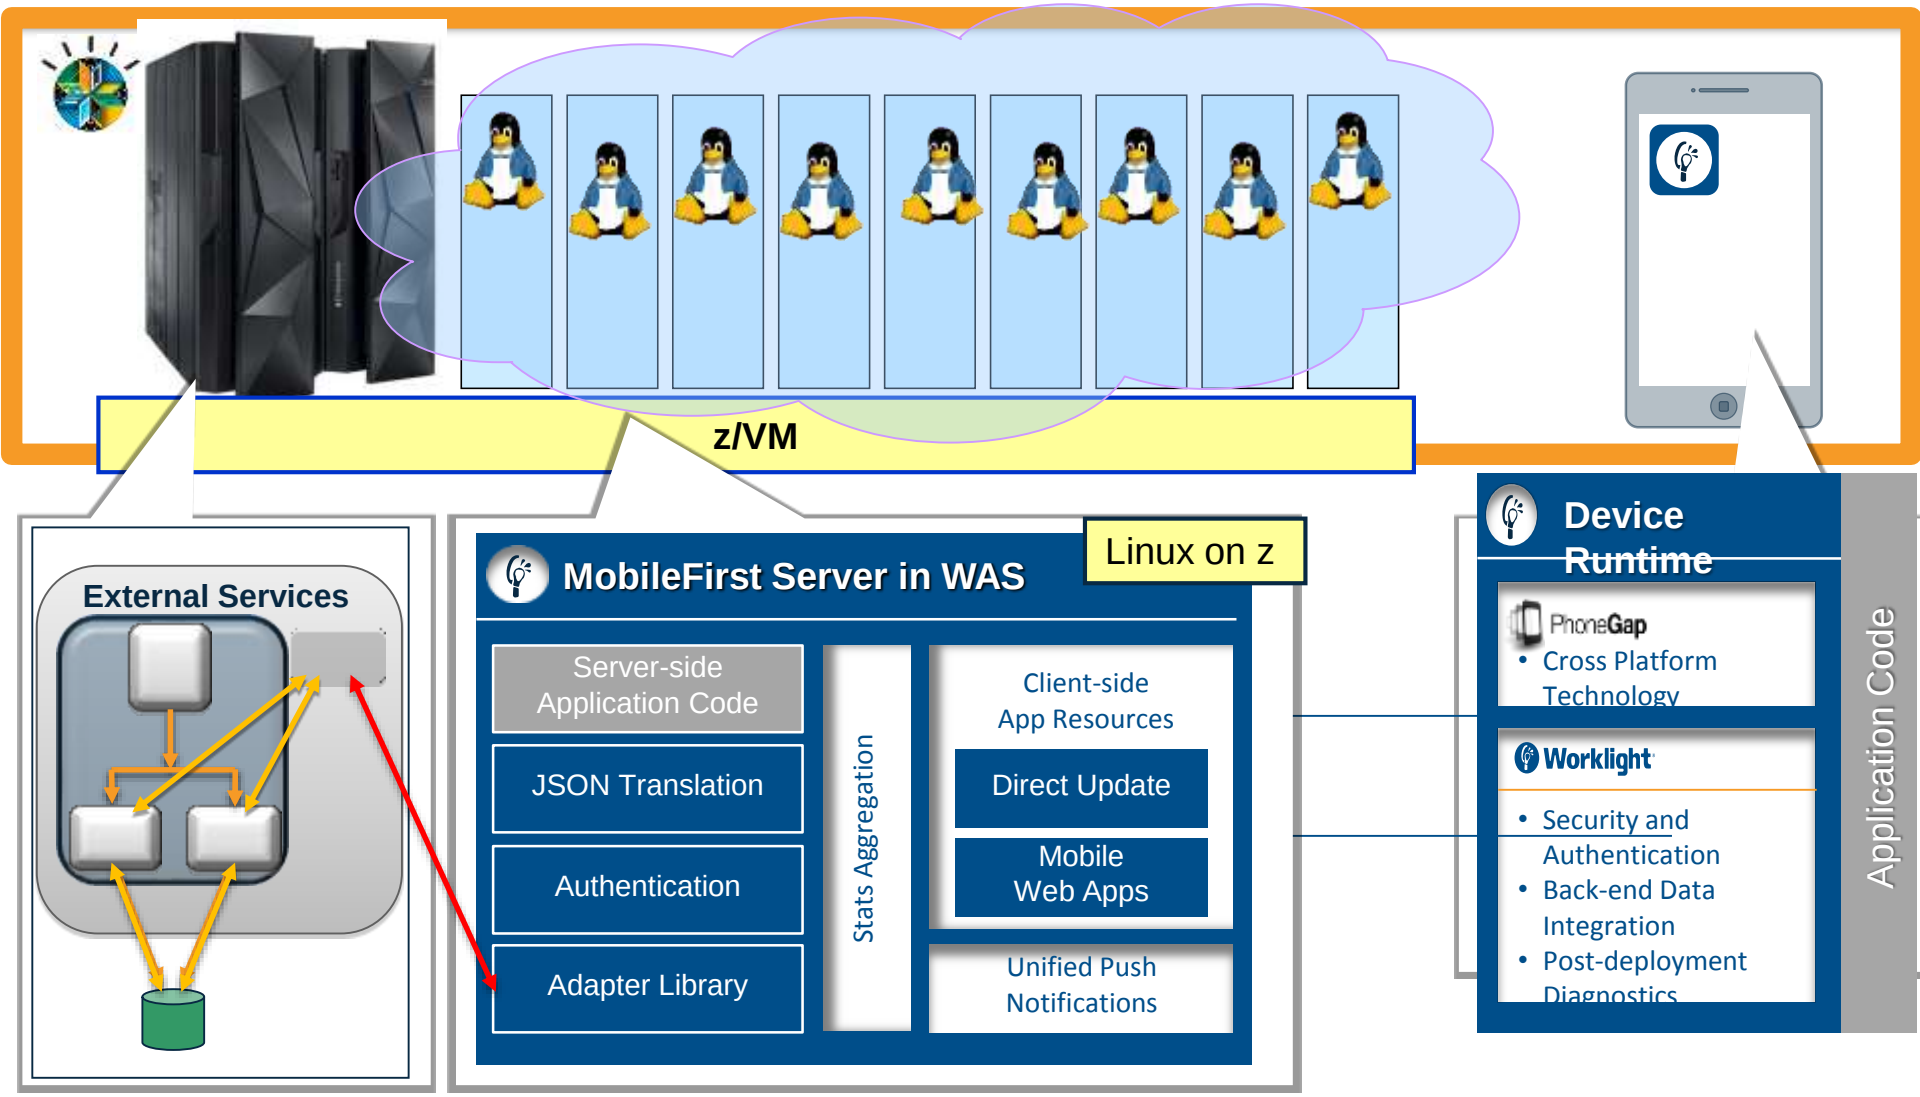

MobileFirst Server on WebSphere on Linux on z Systems Production High Availability

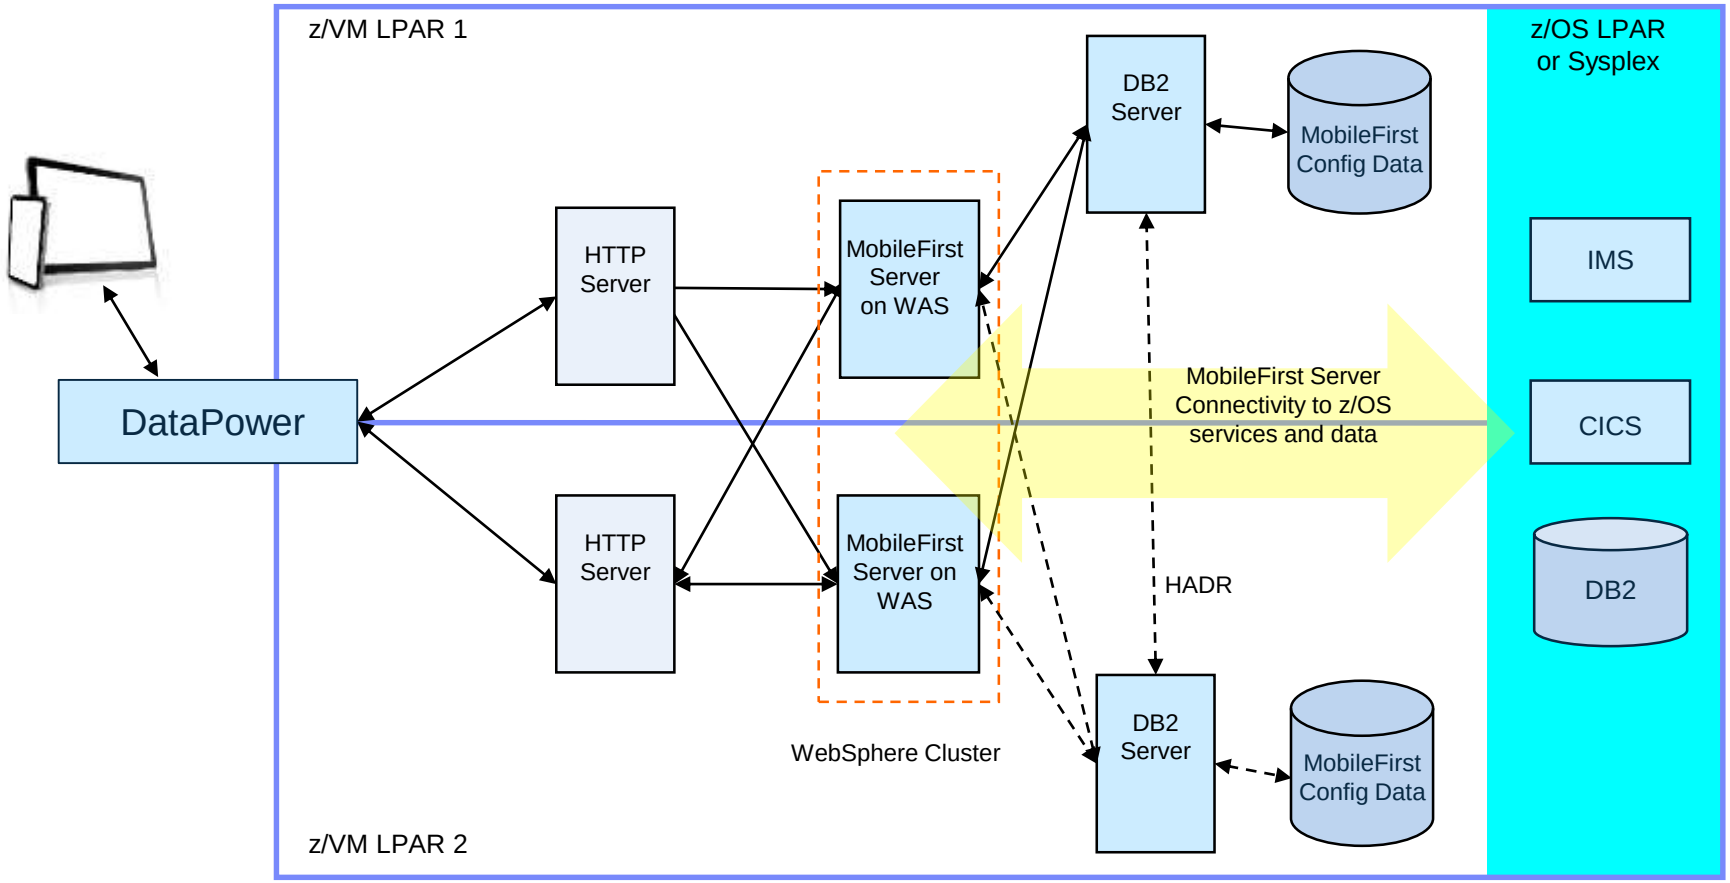

Solid Lines denote primary data path, dashed lines denote backup data path.

SHARE in Orlando 2015

z/VM SSI and GDPS support with xDR

z/VM or LPAR Maintenance - Continuous Availability of z/VM Guests

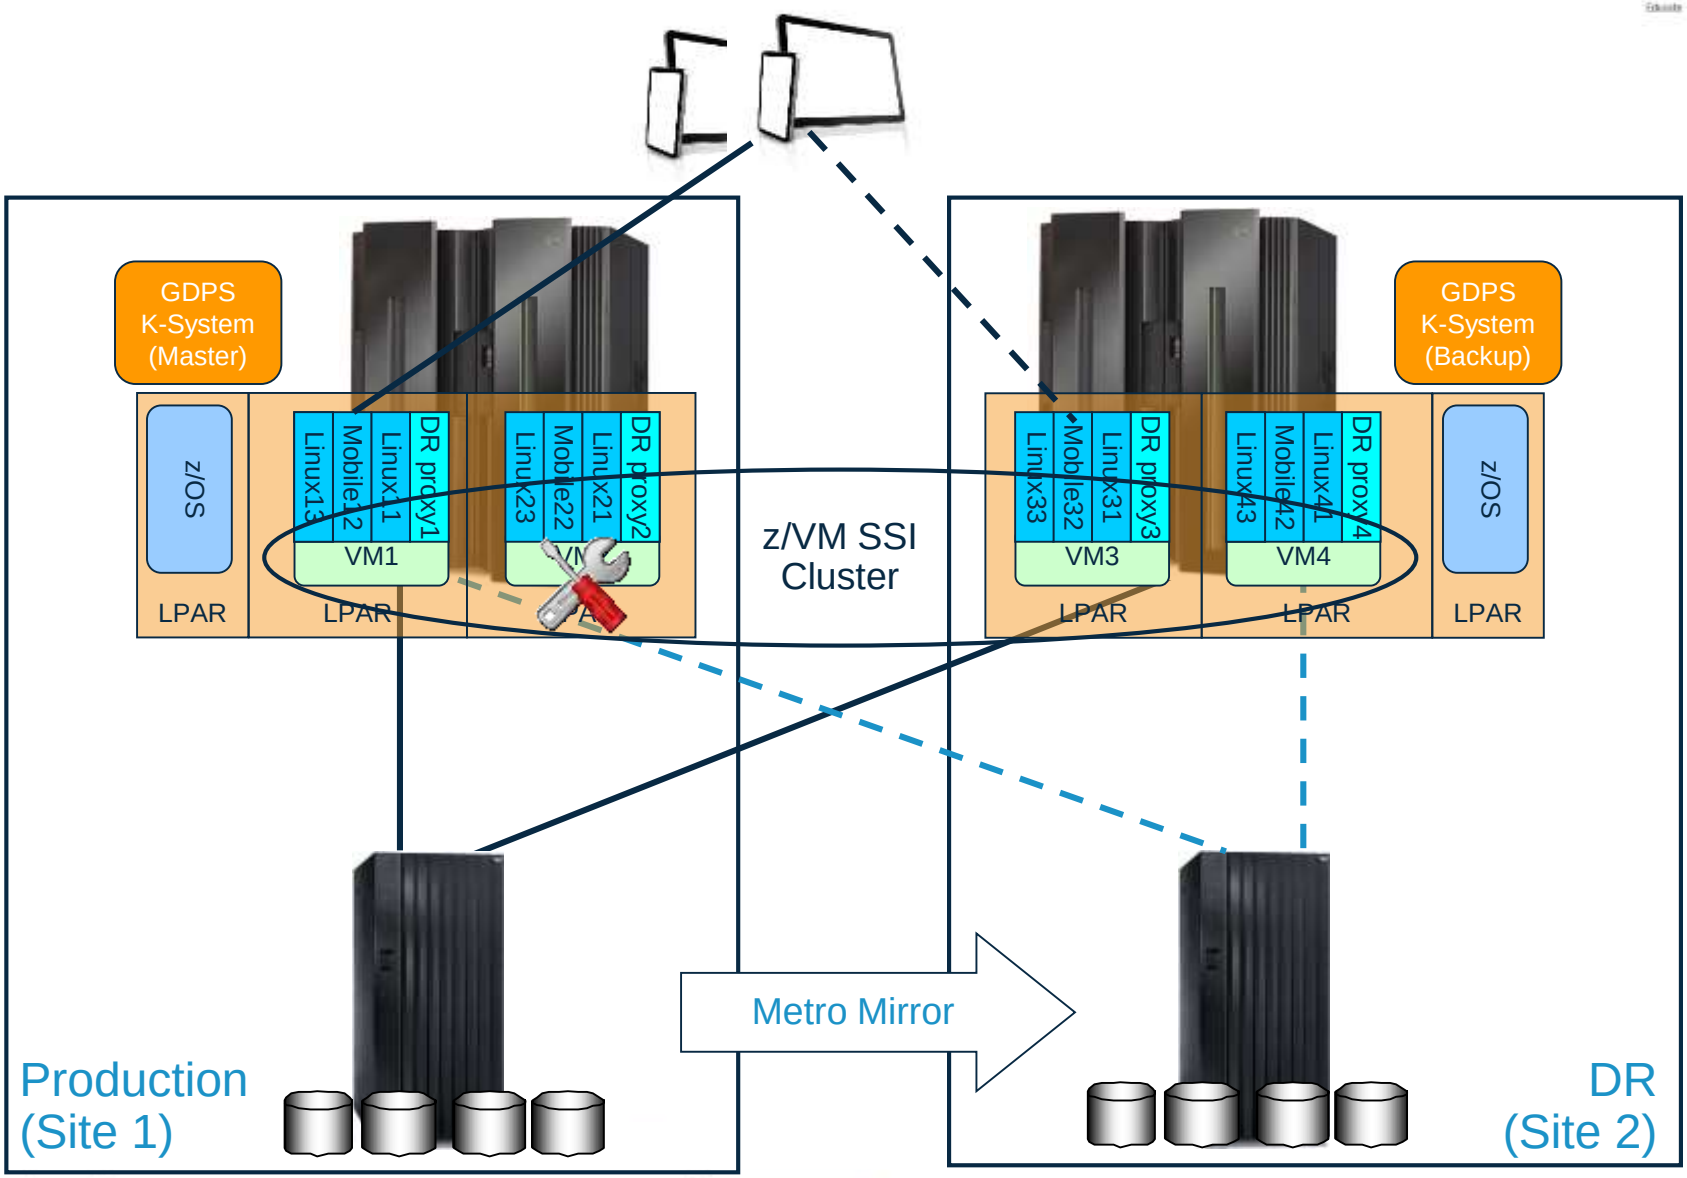

z Systems Ha and DR scenario

Mobile Architecture Overview for z Systems

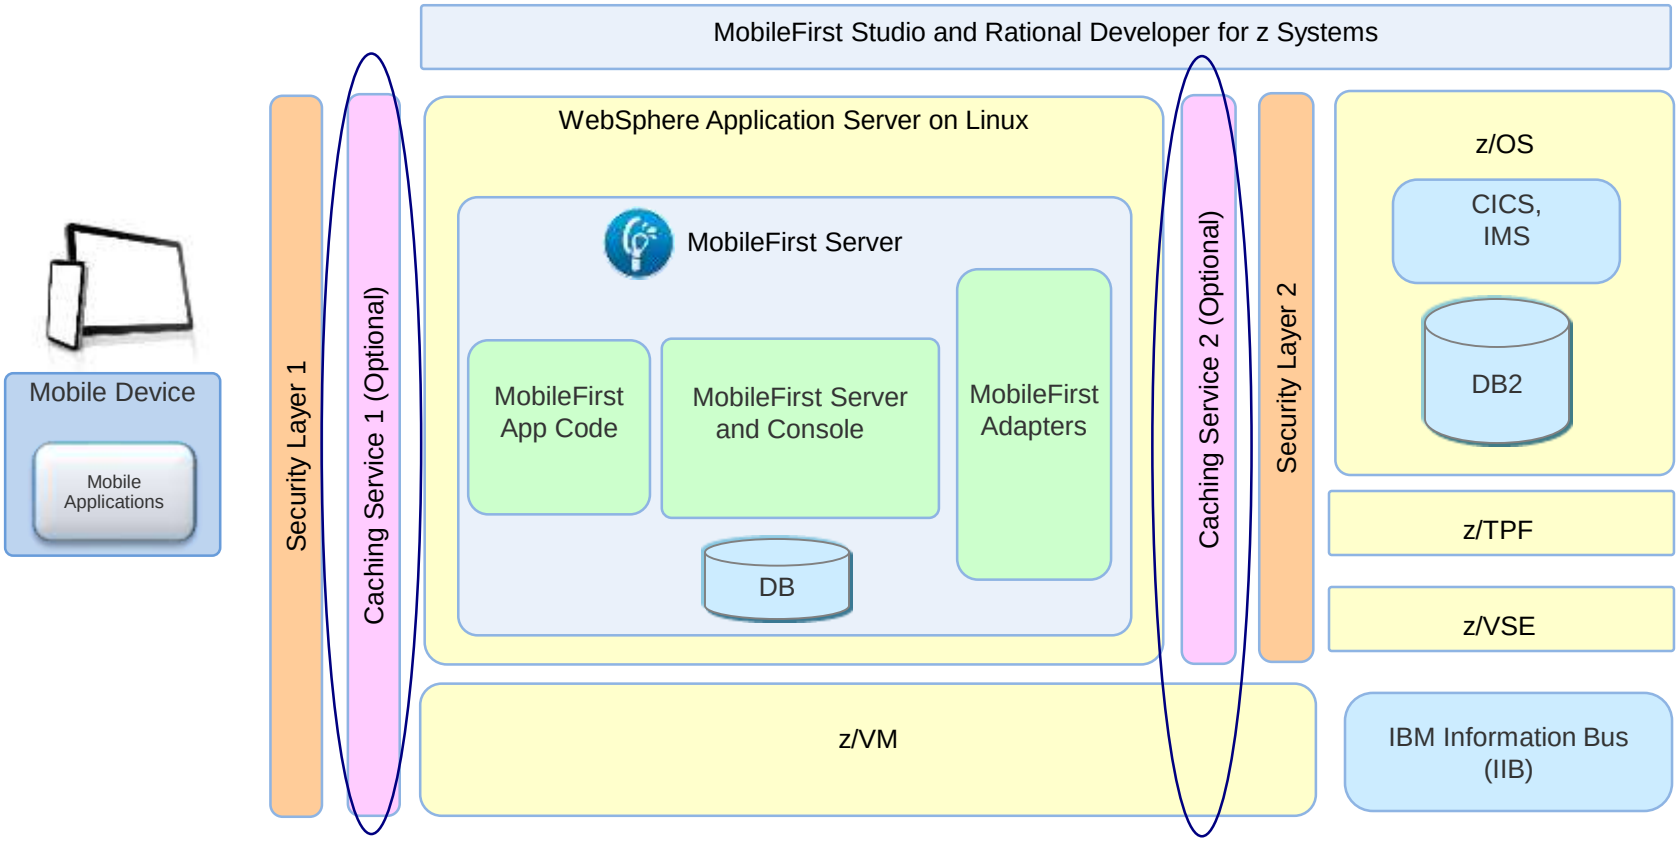

MobileFirst caching

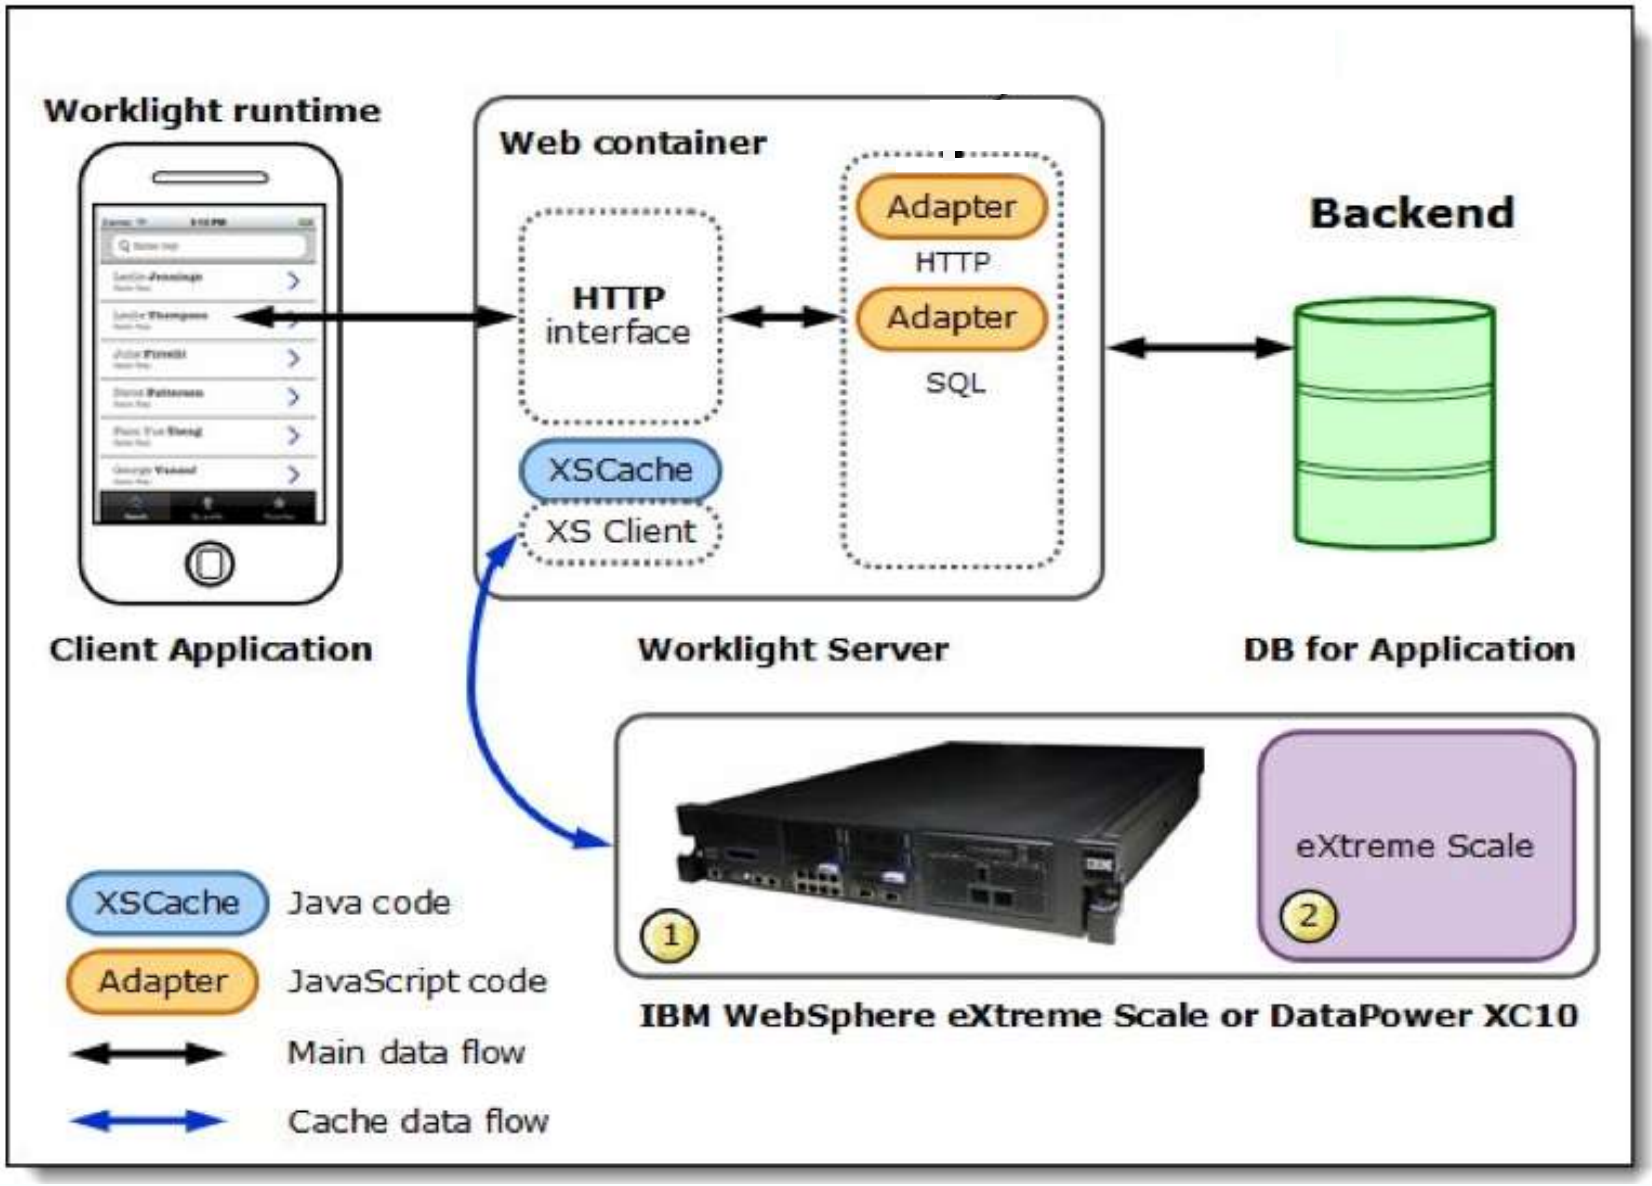

MobileFirst Adapter integration with WebSphere Extreme Scale

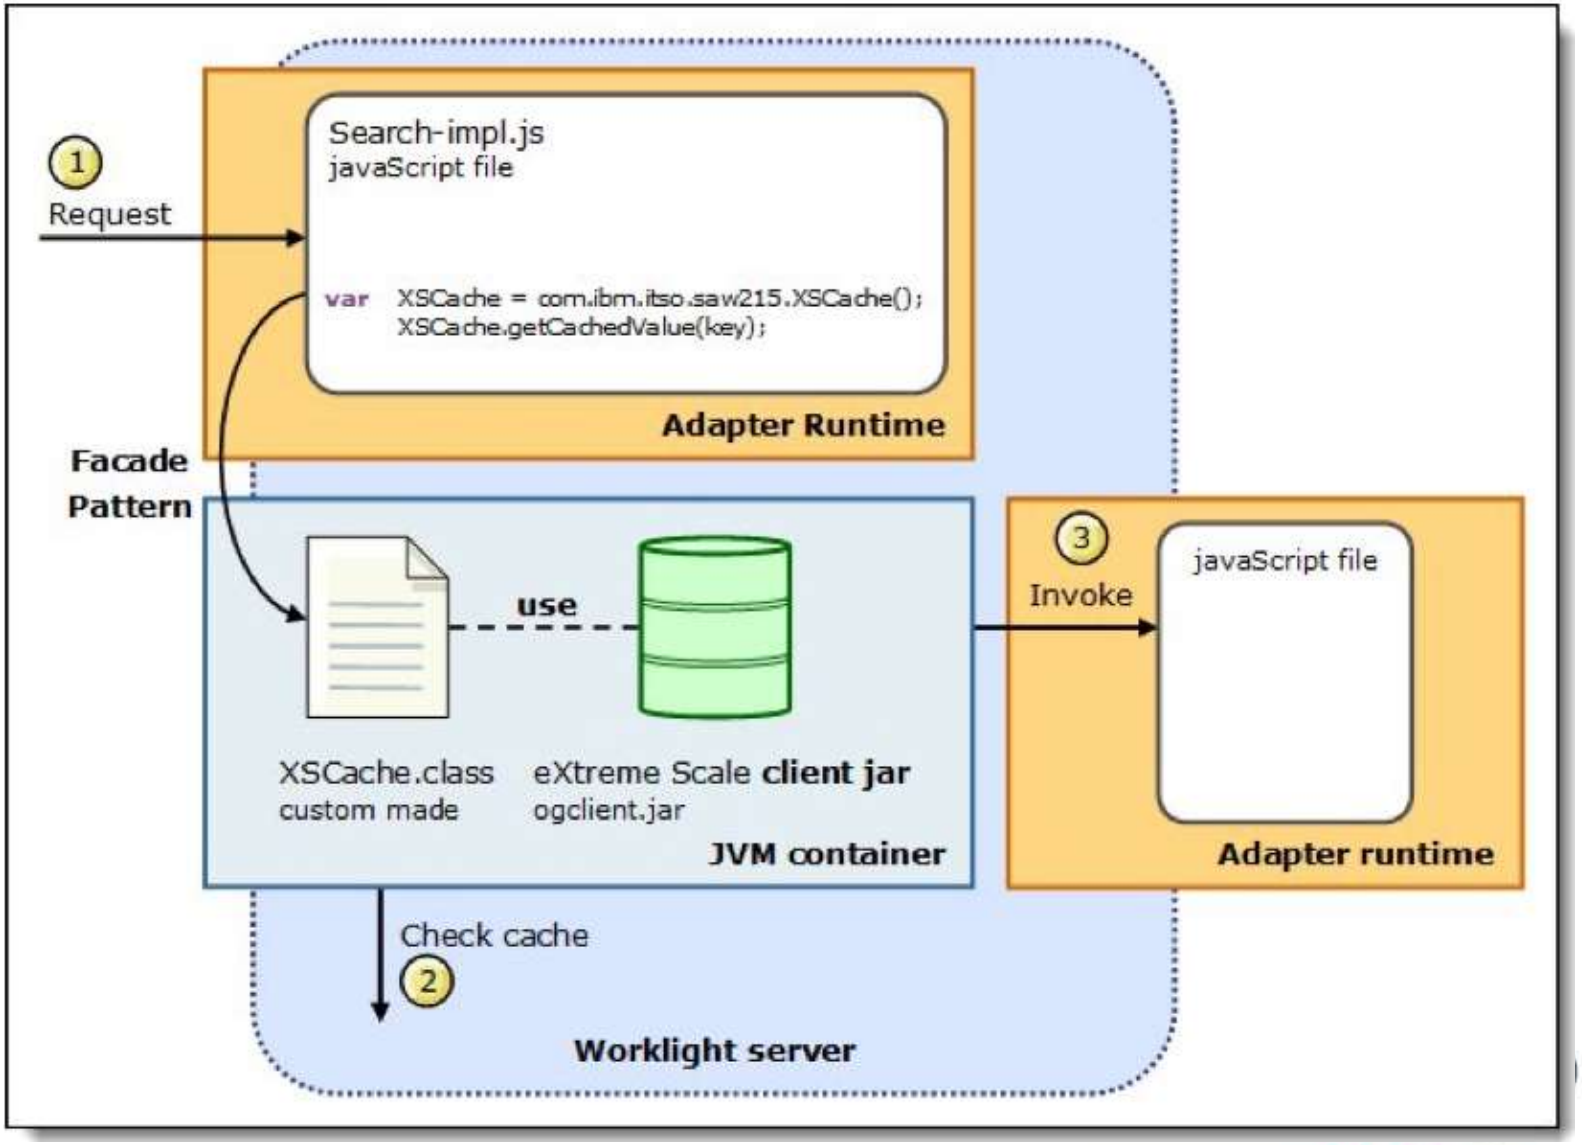

Caching solution decision

Options for caching	Rationale and decision points
When use WebSphere Extreme Scale (WXS)?	WXS is a general-purpose scalable cache. It can be added to any java application running in the mid-tier without requiring changes to any transactions running in the back-end. JavaScript code has to be implemented in the mobile application source to take full benefit of WXS.
When use DataPower XC10 appliance?	Out-of-the box caching appliance that can deliver benefits without adaption of (mobile) application needed. Just configure the network topology to point to the XC10 Typically placed in DMZ to cache static data.
Why use front end caching?	In cases where static data like images, user profiles, product description and HTML are to be cached. Front end caching makes it possible to cache a large set of data, for all requests for (back end) services are processed here. Performance improvement tends to be more of an entry point.
Why use back end caching?	Typically to off-load back end queries in cases where inquiries are made but no business relevant transactions are performed.

Optimization

Tuning a MobileFirst environment on Linux on z

- Implementations for increased End-to-End Performance
- z/VM and Linux OS considerations apply
  - z/VM environment
  - Guest Memory
  - Disk attachements
  - Network options for scalability and HA
- WebSphere for MobileFirst
- MobileFirst database access
- Adapter connectvity
- End-to-end security considerations

Single System Image Feature and cpuplugd
Clustered Hypervisor with Live Guest Relocation and capacity variation

- LGR - dynamically move Linux guests from one member to another with Live Guest Relocation
  - Reduce planned outages
  - Enhance workload management
  - Non-disruptively move work to another z/VM Node

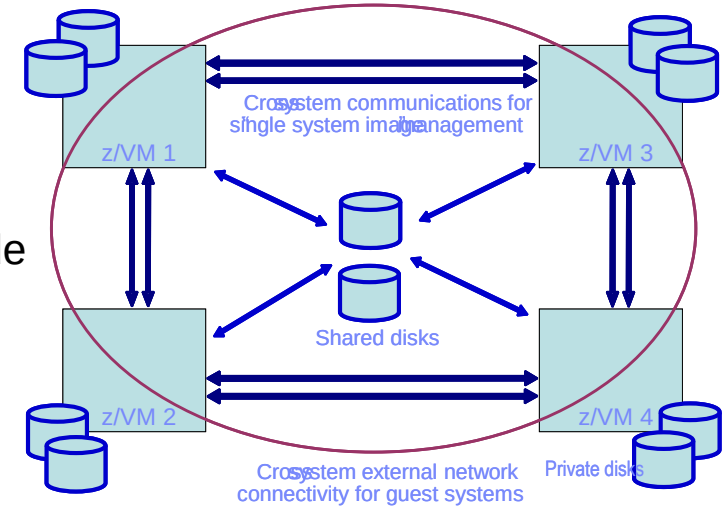

- cpuplugd - changes the number of used processors on the fly, depending on the current overall utilization and load

- When combined with Capacity Upgrade on Demand, Capacity Backup on Demand, and Dynamic Memory Upgrade, you will get the best of both worlds

Bring additional resources to the workload!

Move the workload to the resources!

http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliaag%2Fl0cpup00_2012.htm

Improvements in case where the default (high) number of CPUs is not needed

- Improvements in case where the default (high) number of CPUs is not needed
- up to 40% more throughput,
- up to 40% CPU cost savings

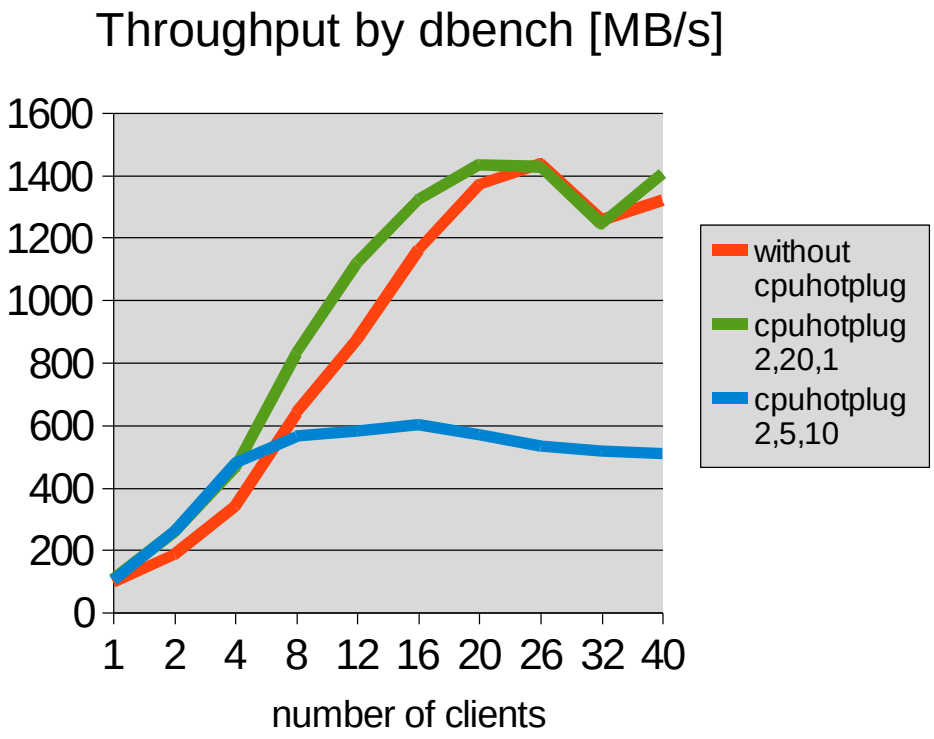

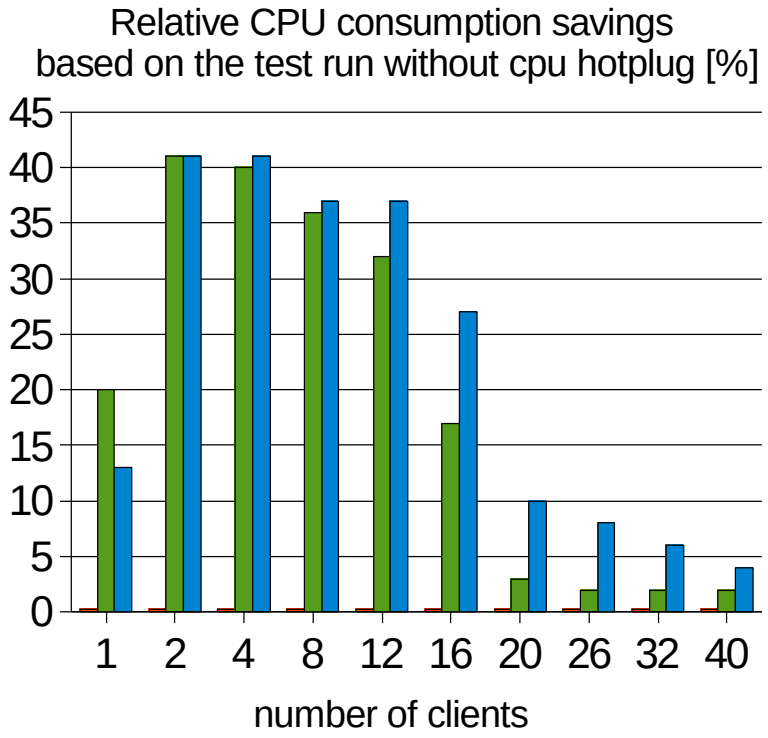

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102139

SRM – memory and paging considerations for Linux guests

Default SRM settings (**q srm**)

LDUBUF : Q1=100% Q2=75% Q3=60%
STORBUF: Q1=125% Q2=105% Q3=95%

LDUBUF – assign the system's paging resources

- 60% for Q3 means: All Q3 guests together can use at maximum 60% of the paging resources (if already used → eligible list)

Recommendation:
SET SRM LDUBUF 100 100 100
to allow all Q3 guests to allocate the whole paging space

STORBUF – to partition main memory

- 95% for Q3 means: All Q3 guests together can use only 95% of the system memory

→ This prevents memory overcommitment when running Linux guests

Recommendation:
SET SRM storbuf 300 250 200
to allow all Q3 guests to allocate twice the amount of real memory
Depending on the level of overcommitment and amount of active/inactive guests, it might be necessary to go even higher, e.g. SET SRM storbuf 300 300 300

Ensure to have sufficient paging resources allocated!

Memory overcommitment

- Memory overcommitment is often mentioned as a major benefit of virtualized environments

  Memory overcommitment is the ability to use more virtual memory as physically available memory. It is based on the assumption that not all guests use their memory at the same time and/or some guests are over-sized in their memory setup.

- Recommendations / limits are missing
  Very different definitions exist for the level of memory overcommitment
  - they are independent of the used middle ware
  - “Common” System z ratio for Non-production guests is 3:1 virtual to physical memory
    E.g. run guests defined with a total of 3GB on 1 GB real memory
  - Performance can be heavily degraded when the memory overcommitment level is to high

- Goal: Proved memory – overcommitment recommendations
- Identify “rules”
  - Determine the minimum amount of physical memory required to run with an acceptable performance
  - Identify a dependency on the used middle ware / applications

z Systems internal network alternatives

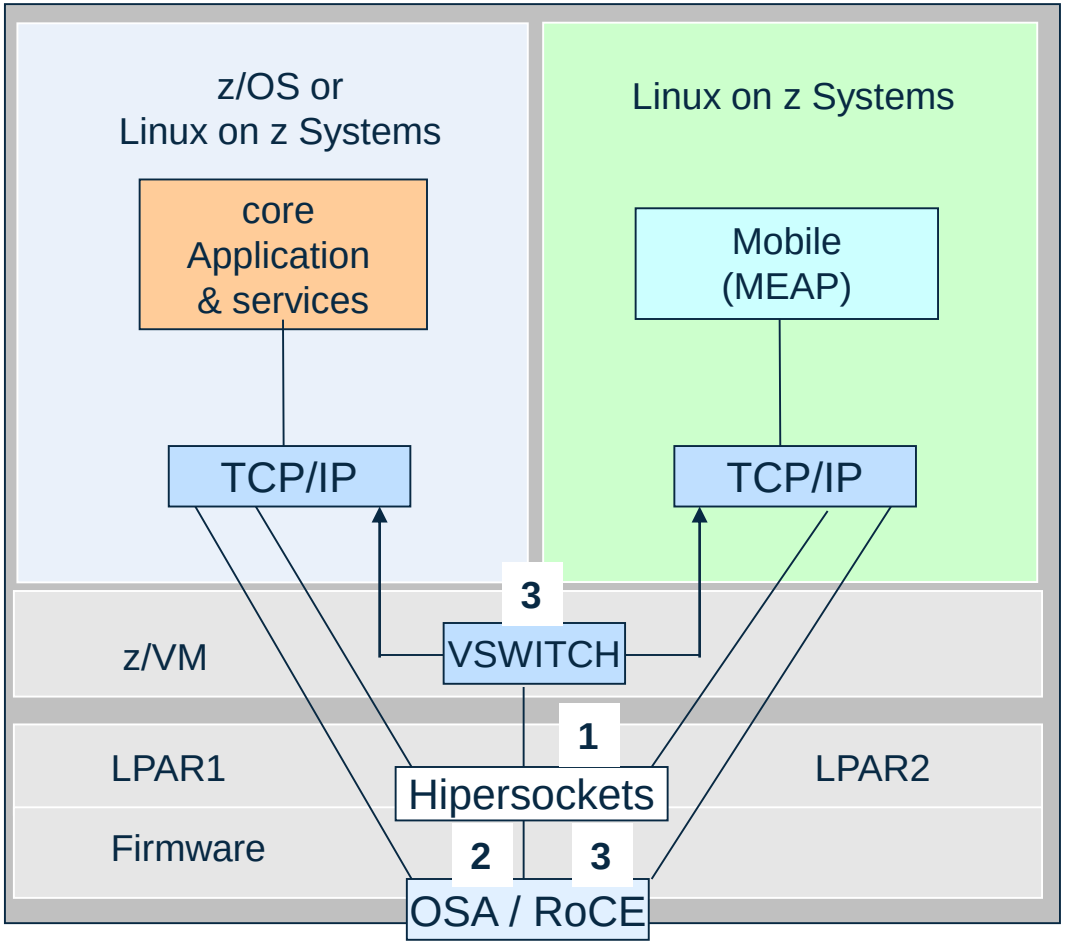

1 - Hipersockets
2 - shared OSA
3 – RoCE
4 - z/VM VSWITCH

Network Bandwidth, Redundancy and Automated Failover

Resource Virtualization:

OSA Channel Bonding in Linux

build-in HA and network scalability

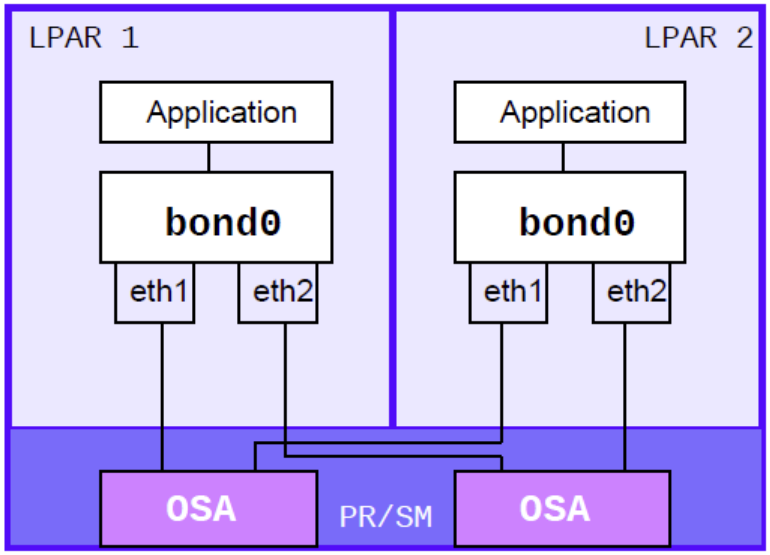

- Linux *bonding* driver enslaves multiple OSA connections to create a single logical network interface card (NIC)
- Detects loss of NIC connectivity and automatically fails over to surviving NIC
- **Active/backup & aggregation modes**
- **Separately configured for each Linux**

Network Virtualization:

z/VM VSWITCH - Link aggregation

build-in HA and network scalability

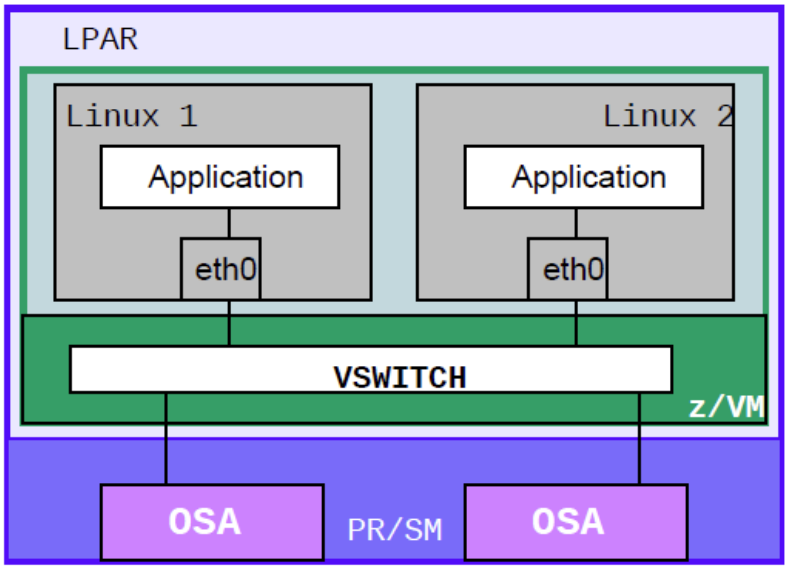

- z/VM *VSWITCH* enslaves multiple OSA connections. Creates virtual NICs for each Linux guest
- Detects loss of physical NIC connectivity and automatically fails over to surviving NIC
- **Active/backup & aggregation modes**
- **Centralized configuration benefits all guests**

Optimization

Tuning a MobileFirst environment on Linux on z

- Implementations for increased End-to-End Performance
- z/VM and Linux OS considerations apply
- WebSphere for MobileFirst
  - Clustering for WAS scalability and HA
  - Java tuning and heap size
  - Parallel threads, memory consumption
- MobileFirst database access
- Adapter connectvity
- End-to-end security considerations

MobileFirst tuning areas on z Systems

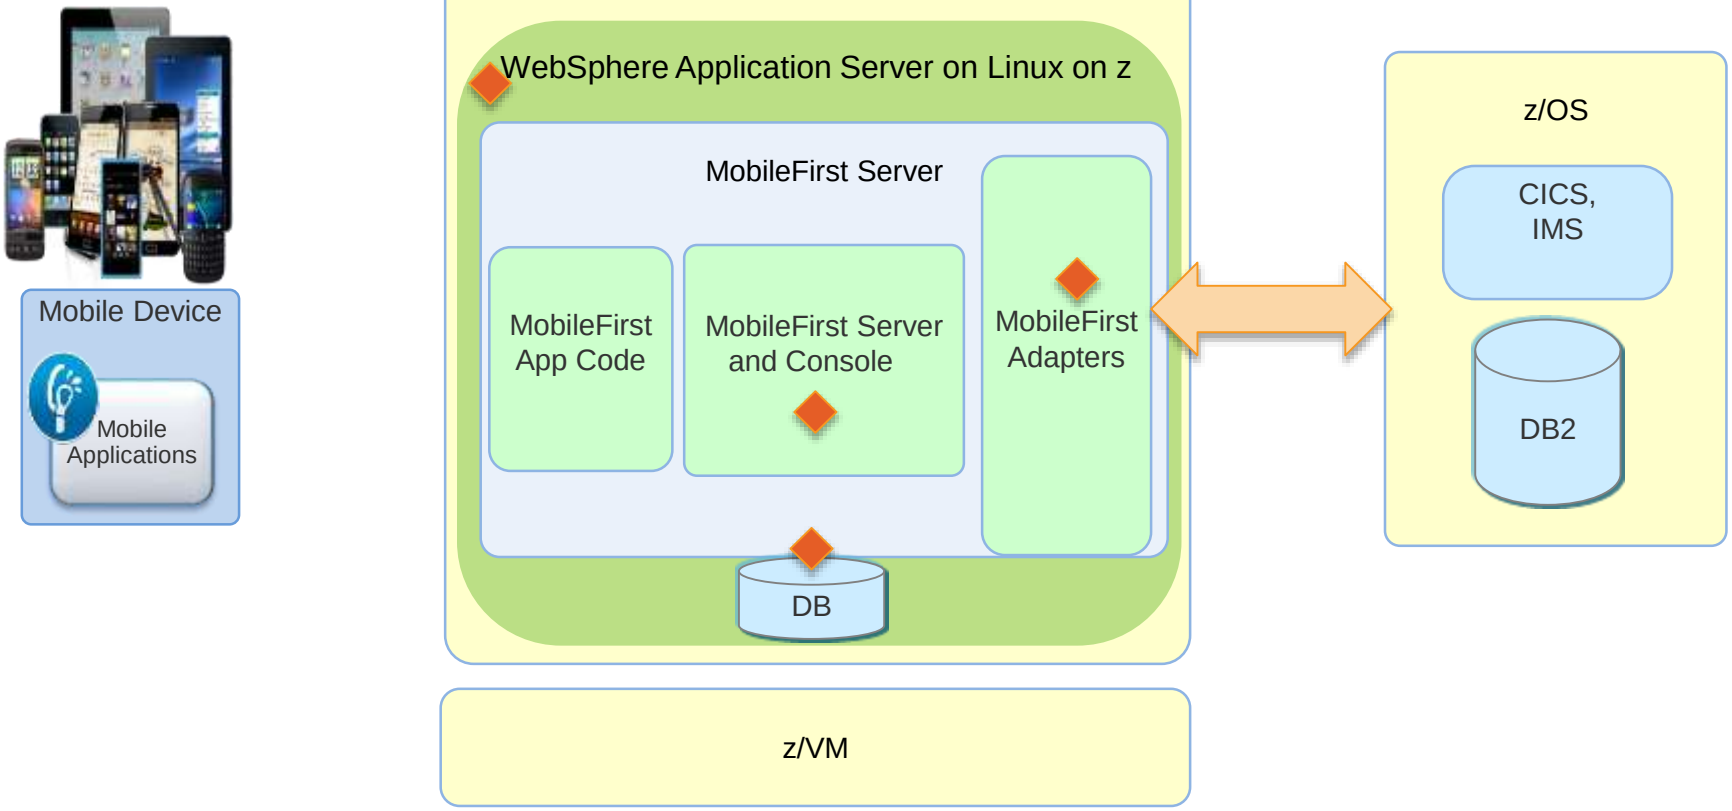

Application server memory sizing for MobileFirst server

The MobileFirst Server can utilize different Application Servers like Apache Tomcat, WebSphere Application Server (WAS) in different flavors WAS Liberty profile, WAS Full Profile, or WAS ND.

IBM MobileFirst server must be installed on a 64-bit operating system with all software at 64 bit.

JVM memory allocation

- Set the JVM to have at least 2GB memory
- For a production environment, it is recommended, setting the minimum and maximum heap size to the same value to avoid heap expansion and contraction.

Where to set application server configuration:

- ***Apache Tomcat***:

Find the Catalina script and set JAVA_OPTS to inject memory.

- ***WebSphere Application Server:***

Log in to the admin console. Go to Servers > Server types > WebSphere application servers: choose each server and set Java memory settings under Java Process definition > JVM arguments

- ***WebSphere Liberty***

Adoptions have to be made in profile jvm.options

Consult:
http://pic.dhe.ibm.com/infocenter/wasinfo/v8r5/index.jsp?topic=%2Fcom.ibm.websphere.wlp.core.doc%2Fae%2Ftwlp_admin_customvars.html

For general guidelines for JVM Memory allocations, you can use the hardware calculator at
http://www.ibm.com/developerworks/mobile/worklight/getting-started.html

Application server thread thresholds

Execution thread behavior

- Each incoming request requires a thread for the duration of that request.
  - Depending on workload or connection type this varies
- Simultaneous requests are handled by the currently available request processing threads
  - Additional threads will be created up to the configured maximum.

Application server configuration:

- **Apache Tomcat:**
  - By default the maximum number of threads is 200.
  - For details consult: http://tomcat.apache.org/tomcat-7.0-doc/config/http.html

- **WebSphere Application Server:**
  - By default the maximum number of threads is 50.
  - Verify via admin console. (Go to Servers > Server types > WebSphere application servers > server_name > Web container)

- **Liberty see executer section in:**
  - By default the maximum number of threads is unbounded.
  - For Details see: http://pic.dhe.ibm.com/infocenter/wasinfo/v8r5/index.jsp?topic=%2Fcom.ibm.websphere.wlp.nd.multiplatform.doc%2Fautodita%2Frwlp_metatype_4ic.html
  - even though the maximum number of threads is theoretically unbounded, the executor service makes informed choices about whether adding another thread will actually be useful.

There are several considerations when setting http threads configuration:

- Analyze request behaviors.
  - if the longest call takes 500 ms and you have maximum of 50 threads, you can have about 100 requests per second.
- Back-end connection behavior influences thread execution
  - For slow back-end services you will need to increase the number of default threads.
  - In addition increase the number of back-end connection threads (set maxConcurrentConnectionsPerNode as shown below).
- For high number of concurrent users, increase the number of default threads

MobileFirst session timeouts

Mobile clients have a ‘heart beat‘which allows the mobile client to ping the server while the App is in the foreground so that the server session will not time out.

Also note that when a Mobile App is moved into the background, it no longer interacts with the server or sends a “heartbeat” leading the server session to stop after the specified server session timeout.

Parameter for session control:

- serverSessionTimeout – Client inactivity timeout, after which the session is invalidated.
  - Default session timeout is 10 minutes. The default can and should be configured.
    - It is recommended to set it from 3 to 10 minutes.
  - This parameter affects the server memory consumption.
  - A session is an object stored in the server memory for each connecting device (with its authentication information)
  - Active sessions are determined by the number of sessions opened vs. the sessions timing out due to lack of activity

- Example with 10 min session timeout: Suppose every minute 1,000 users start a session against the server. Even if they exit the application after 3 minutes, their session will remain active on the server for 10 minutes, leaving us with 10 x 1,000 = 10,000 sessions.

MobileFirst server background tasks

MobileFirst background tasks perform several actions on the database and/or file system. They can be controlled via parameters in MobileFirst.properties file.

Important parameters for background tasks:

- cluster.data.synchronization.taskFrequencyInSeconds – The parameter controls the sync interval of the file system with the database content.
  - default is 2 seconds
  - application and adapter files are stored in the database for synchronization of the deployment data between all cluster nodes
  - every 2 seconds every MobileFirst server node checks the Database to see if a new adapter or application was deployed in another MobileFirst server node and will deploy the adapter/application to local node & file system
  - Increasing this fequency number will cause fewer queries on the database, however it will also increase the unsychncronized MobileFirst server nodes
- deployables.cleanup.taskFrequencyInSeconds – Delete unused deployables from the file system.
  - default is 24 hours.
- sso.cleanup.taskFrequencyInSeconds –The SSO (Single Sign on) mechanism stores session data in a database table.
  - this parameter defines the interval for the SSO cleanup task
  - default is 5 seconds (every 5 seconds accounts are checked for inactivity - idle for more than serverSessionTimeout)
- push.cleanup.taskFrequencyInSeconds – Delete inactive push notification subscriptions, currently implemented only for Apple APNS.
  - Default is 60 minutes.

Optimization

Tuning a MobileFirst environment on Linux on z

- z/VM and Linux OS considerations apply
- WebSphere for MobileFirst
- MobileFirst database access
  - Database concurrent access considerations
  - Database clustering for availability and HA
- Adapter connectivity
  - Adapter scalability considerations
- Implementations for increased End-to-End Performance
- End-to-end security considerations

Tuning MobileFirst database connections

- Configure in the data source, the number of connection threads from the server to its database.
- Two MobileFirst features rely heavily on the Database connection threads
  - Single Sign On (SSO)
  - MobileFirst Reporting / Analytics feature

Limitations

- Each node in MobileFirst server cluster has max of
  - MAX_DB_INCOMING_CONNECTIONS & NUM_OF_CLUSTER_NODES connection threads
    - MAX_DB_INCOMING_CONNECTIONS is the maximum incoming connections defined in the database server
    - NUM_OF_CLUSTER_NODES represents the number of MobileFirst server nodes in the cluster

As rule of thumb, set the number of database connections equally with the number of http threads in the application server

***Data source* configuration in application servers:**

- **For WebSphere Application Server see:**
  - http://pic.dhe.ibm.com/infocenter/wasinfo/v7r0/index.jsp?topic=%2Fcom.ibm.websphere.nd.multiplatform.doc%2Finfo%2Fae%2Fae%2Fudat_conpoolset.html
- **For WebSphere Liberty see datasource section in:**
  - http://publib.boulder.ibm.com/infocenter/radhelp/v8r5/index.jsp?topic=%2Fcom.ibm.websphere.wlp.nd.multiplatform.doc%2Fautodita%2Frwlp_metatype_4ic.html
- **Apache Tomcat:**
  - http://tomcat.apache.org/tomcat-7.0-doc/jndi-datasource-examples-howto.html

Adapter back-end connection tuning

Define the maximum number of concurrent requests from the MobileFirst server to the back-end services application node with:

maxConcurrentConnectionsPerNode – in the *adapter.xml* at connectivity level

There are two considerations when setting this parameter:

- ***If no limitation in the back-end about the incoming connections***
  - set the number of connection threads per adapter to be the number of http threads in the application server
  - For more precise setting, set the number respectively to each back-end service (HTTP, SOA, Database service)

- ***The back-end with limitation on the incoming connection threads depend on:***
  - BACKEND_MAX_CONNECTIONS
  - NUM_OF_CLUSTER_NODES
    - BACKEND_MAX_CONNECTIONS is the maximum incoming connections define in the back-end server
    - NUM_OF_CLUSTER_NODES is the number of MobileFirst server nodes in the cluster

Optimization

Tuning a MobileFirst environment on Linux on z

- Implementations for increased End-to-End Performance
- z/VM and Linux OS considerations apply
- WebSphere for MobileFirst
- MobileFirst database access
- Adapter connectivity
- End-to-end security considerations
  - End-2-end Security
  - DMZ authentication for high number of requests

Secure the Users & Devices and every transaction from Mobile to the Enterprise transactions and data

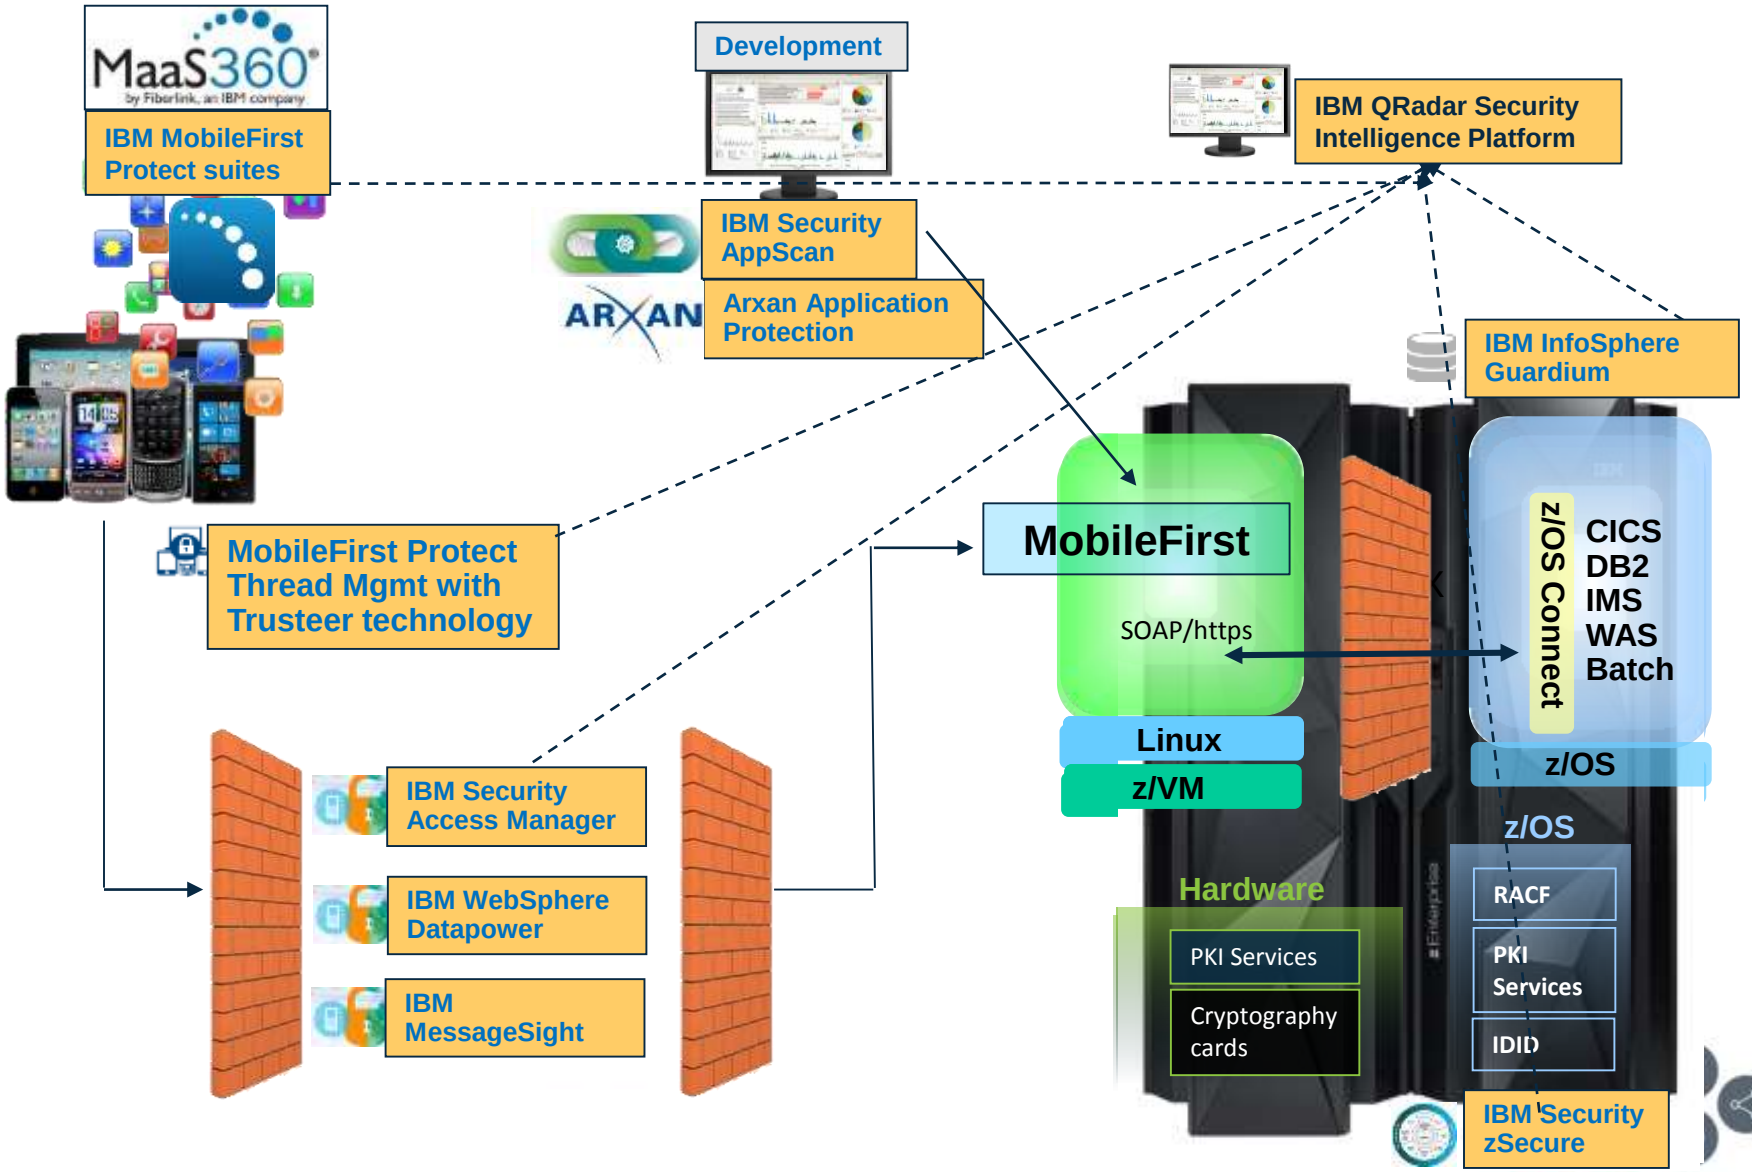

DataPower Mobile Security Features

Available as a physical or virtual appliance

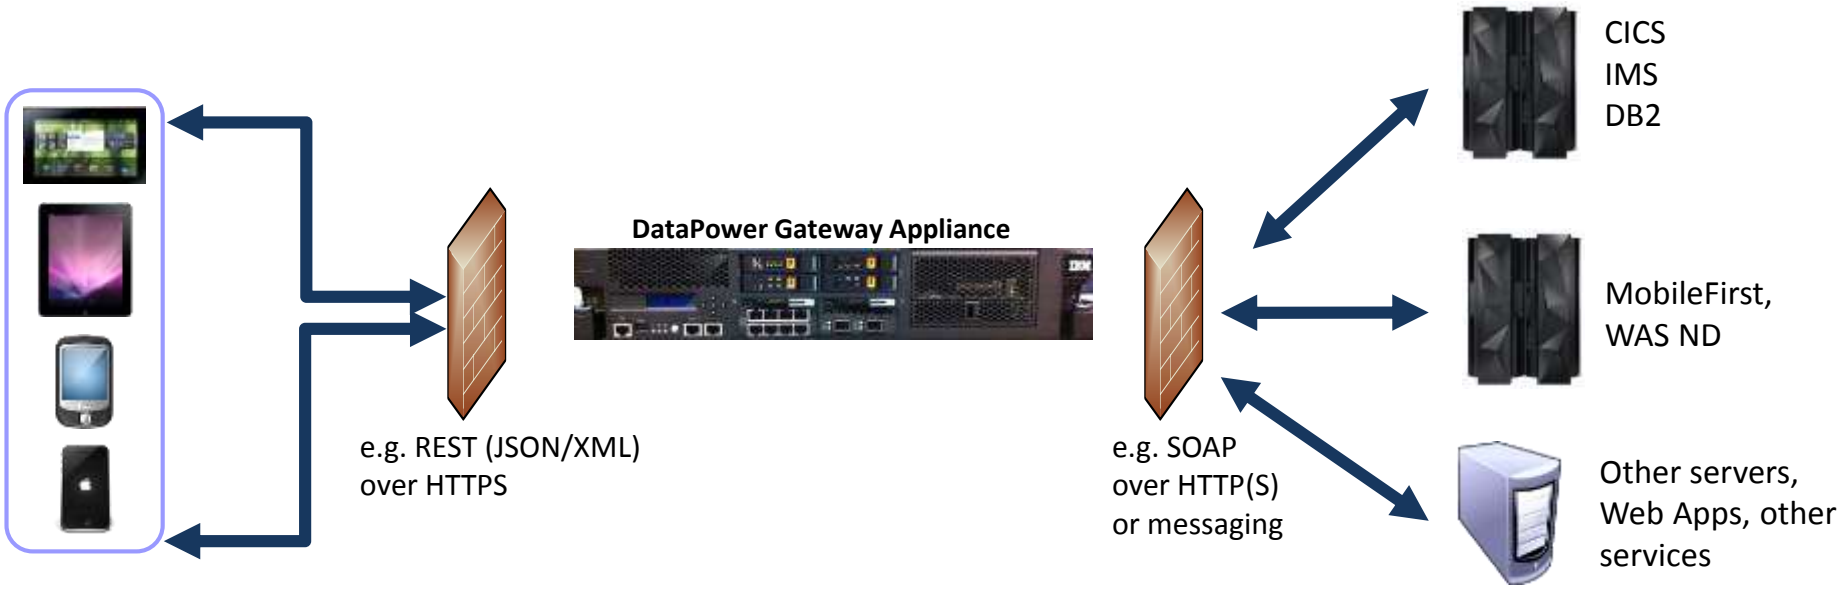

- Security, Control, Integration & Optimization of mobile workload
- Enforcement point for centralized security policies
- Authentication, Authorization, **SAML**, OAuth 2.0, Audit
- Threat protection for XML and JSON
- Message validation and filtering
- Centralized management and monitoring point
- Traffic control / Rate limiting
- Integration with MobileFirst

© 2015 IBM Corporation

DataPower: Authentication, Authorization, Auditing (AAA)

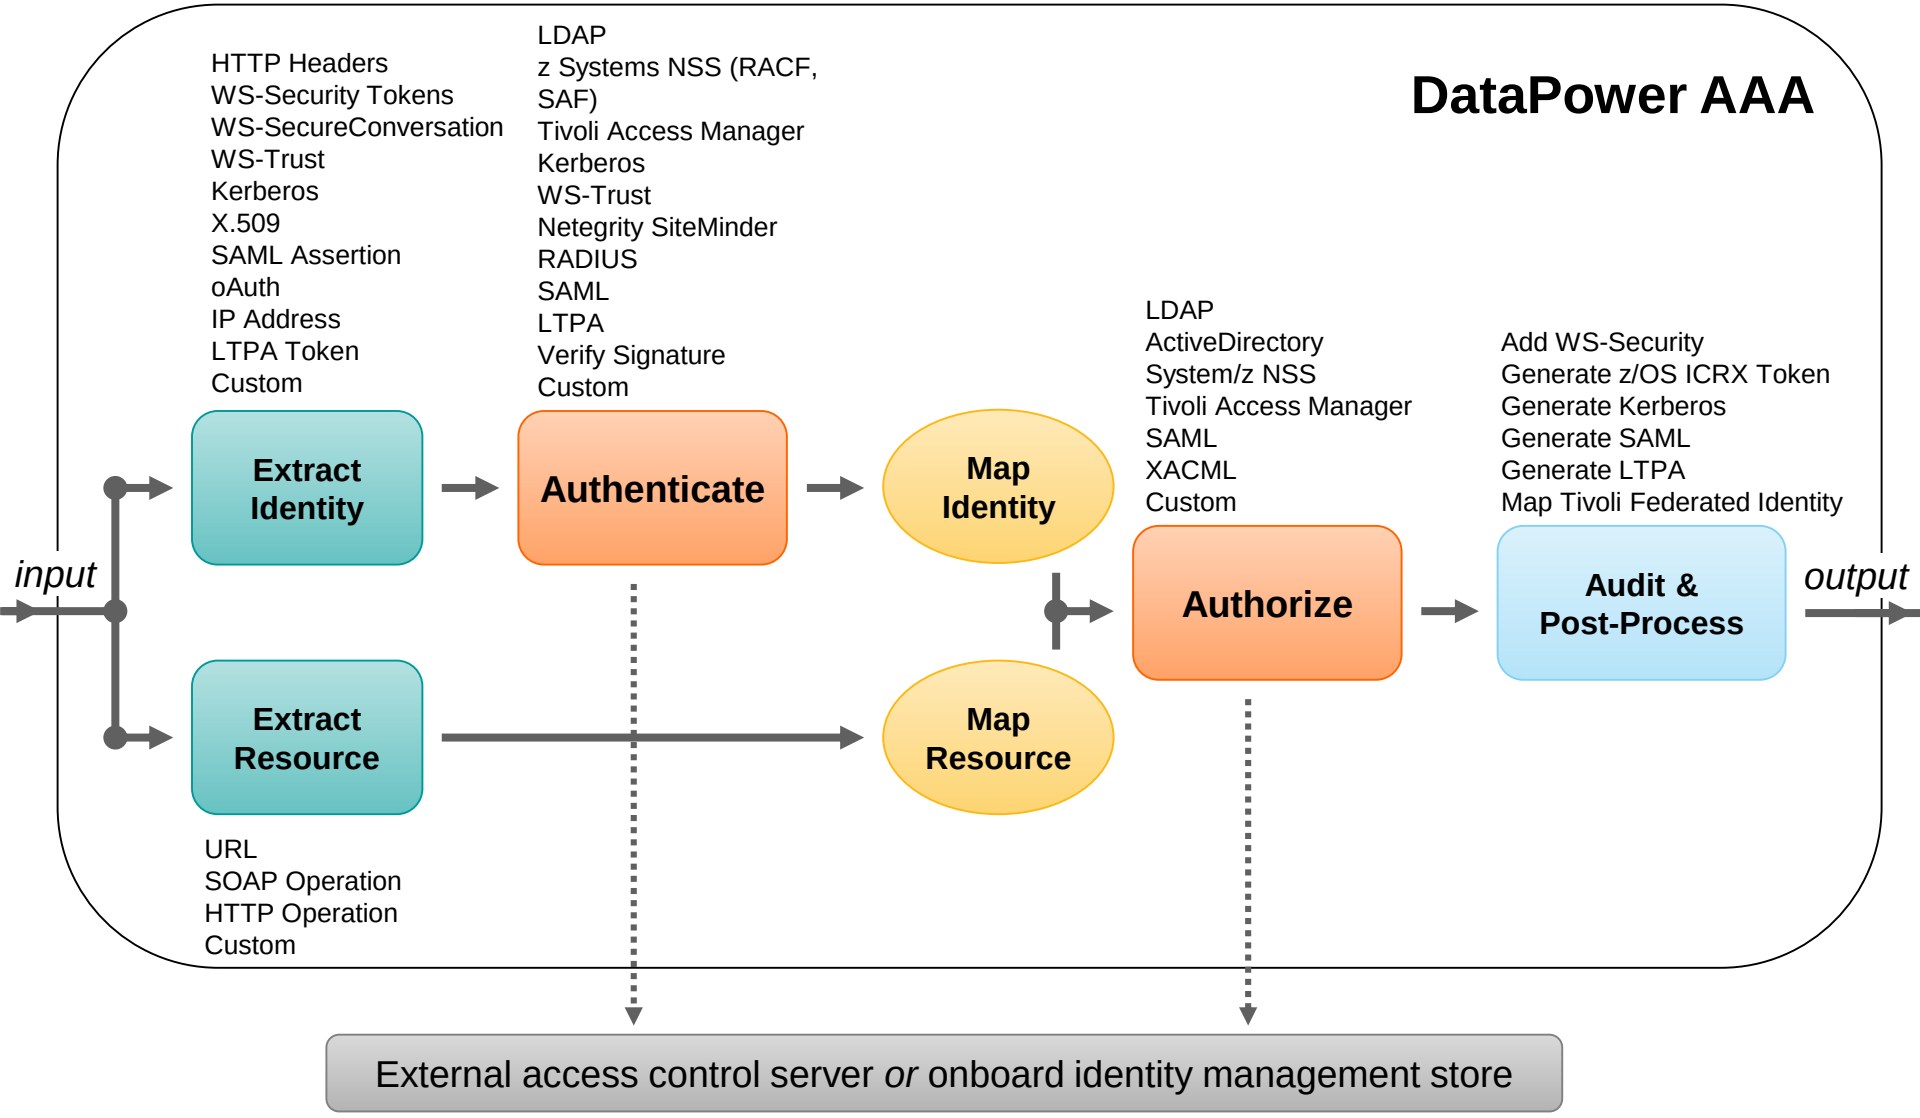

© 2015 IBM Corporation

Access to z/OS via z/OS Connect with increased security

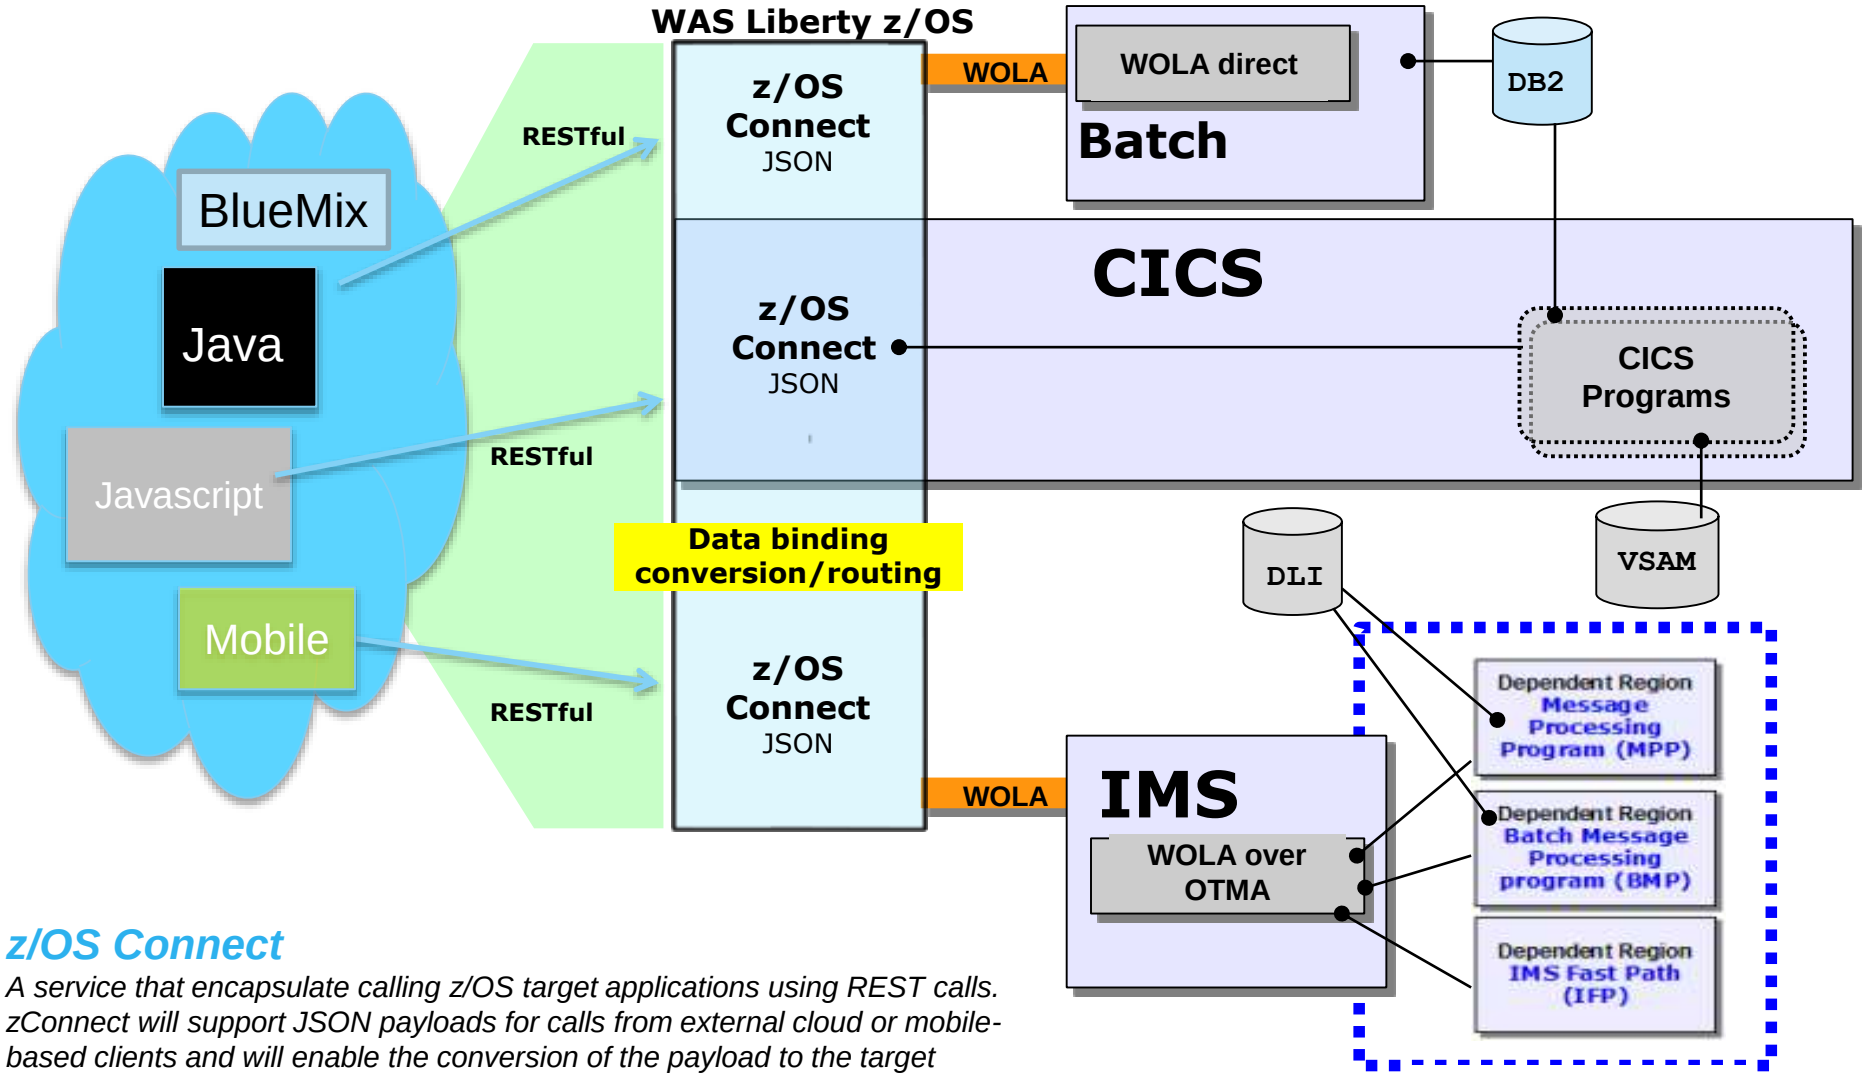

z/OS Connect

A service that encapsulate calling z/OS target applications using REST calls. zConnect will support JSON payloads for calls from external cloud or mobile-based clients and will enable the conversion of the payload to the target program's expected format (WOLA – WebSphere Optimized Local Adapters). It will also provide the response payload conversion from a byte array into JSON format before returning the response to the caller.

© 2015 IBM Corporation

z/OS Connect Security

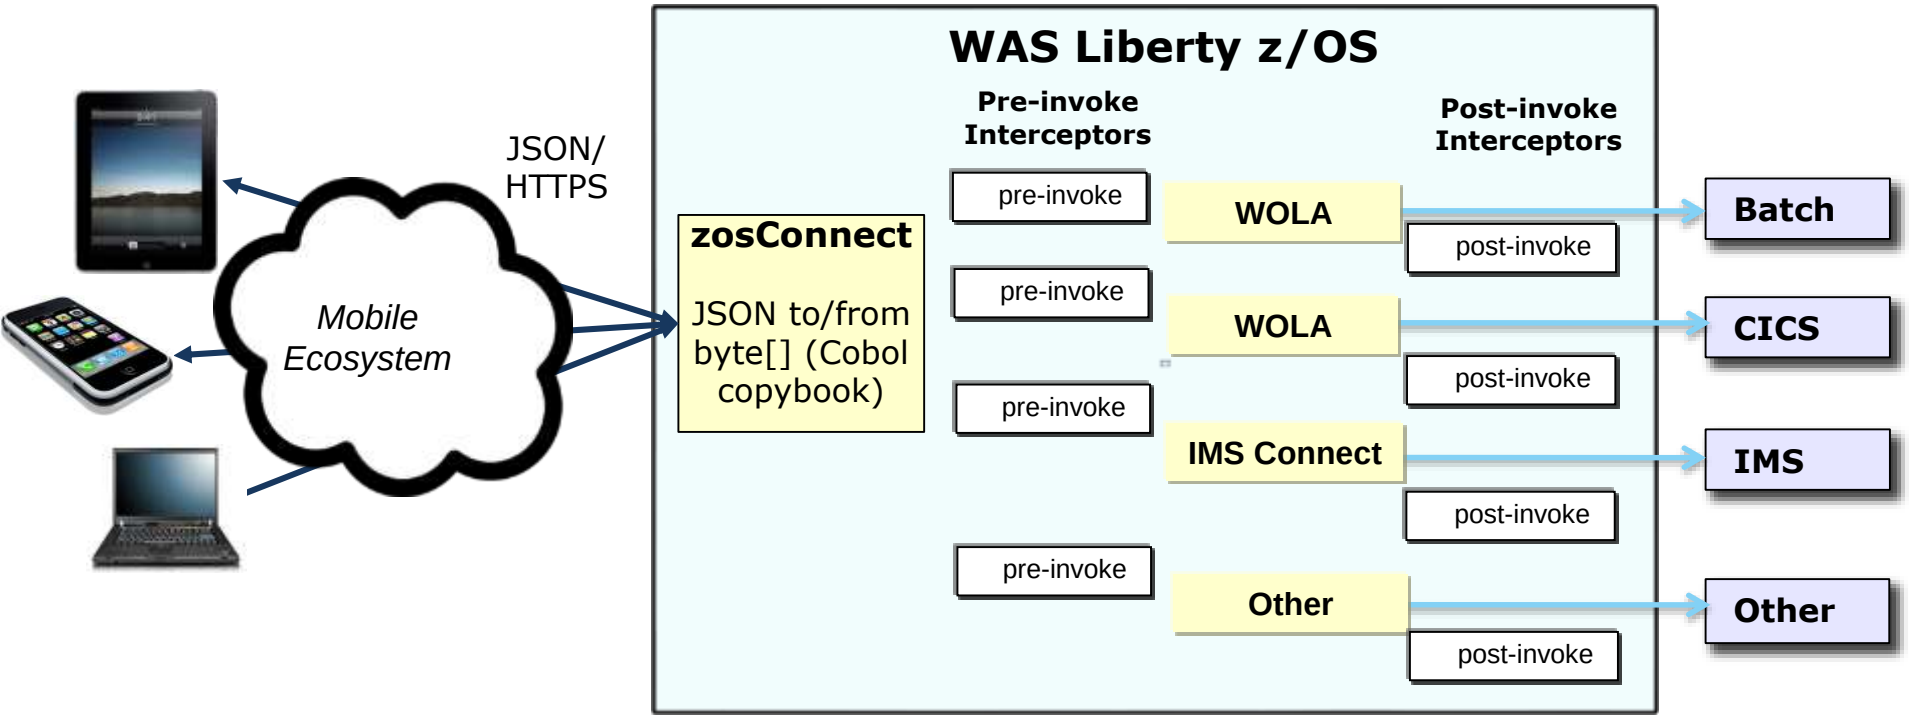

- Framework that allows interceptors to be executed around the invocation of the service
- Authentication with RACF or LDAP
- Authorization interceptor e.g is user in 'Invoke' group for requested service)
  - **com.ibm.wsspi.zos.connect.Authorization()**
- Audit interceptor for SMF-based auditing
  - **com.ibm.wsspi.zos.connect.Audit()**

© 2015 IBM Corporation

Implementation with security flow- options

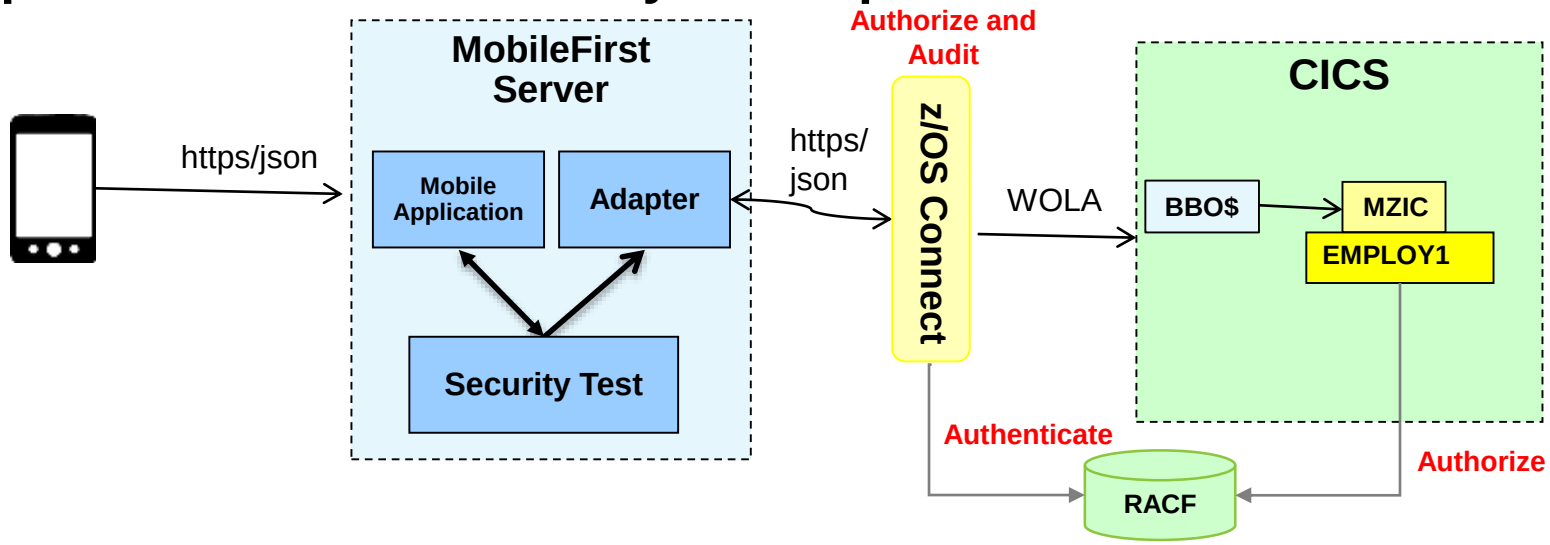

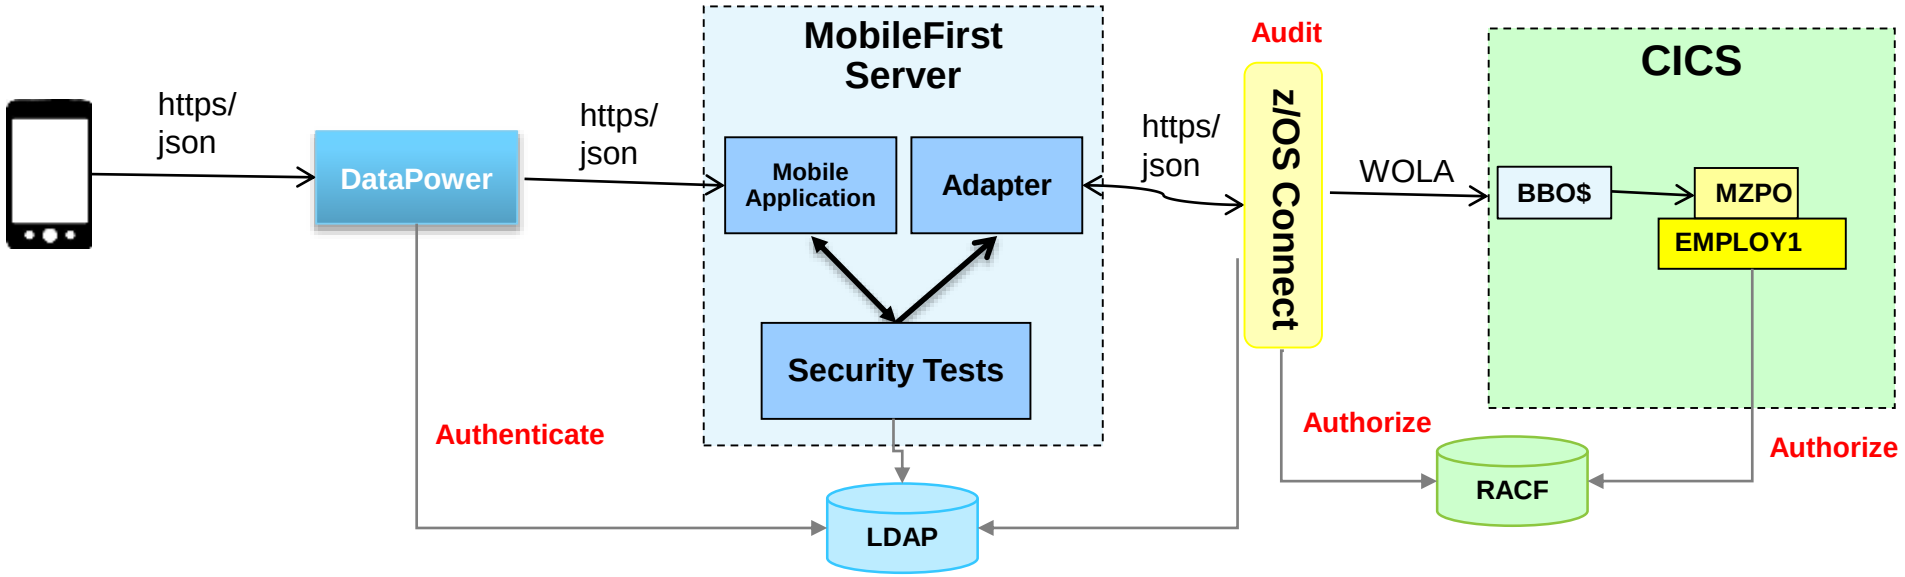

© 2015 IBM Corporation

‘Tuning‘ of Mobile Workload Pricing – up to 60% savings for Mobile workload

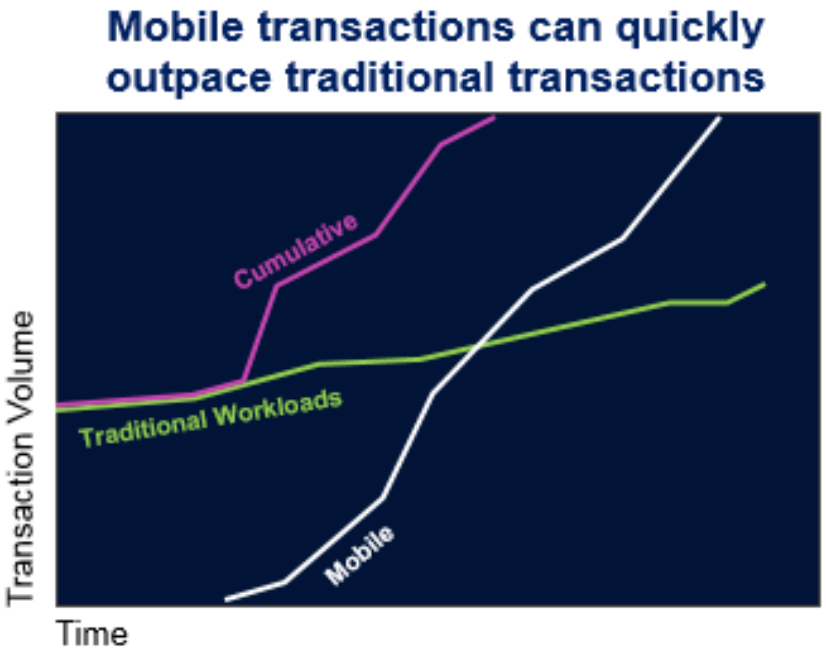

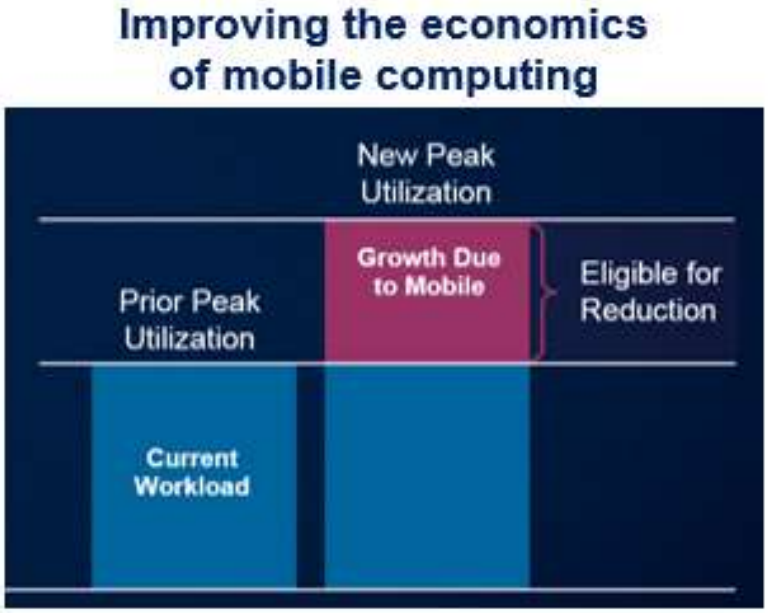

MWP improves the value of growth for mobile transactions processed in System z environements, such as CICS; IMS; MQ; DB2; it enhances the sub-capacity pricing; no infrastrucutre changes required, no LPAR needed.

Announcement letter:

http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/0/877/ENUSZP14-0280/index.html&lang=en&request_locale=en

© 2015 IBM Corporation

Mobile Environment on zEnterprise connecting to Core Systems

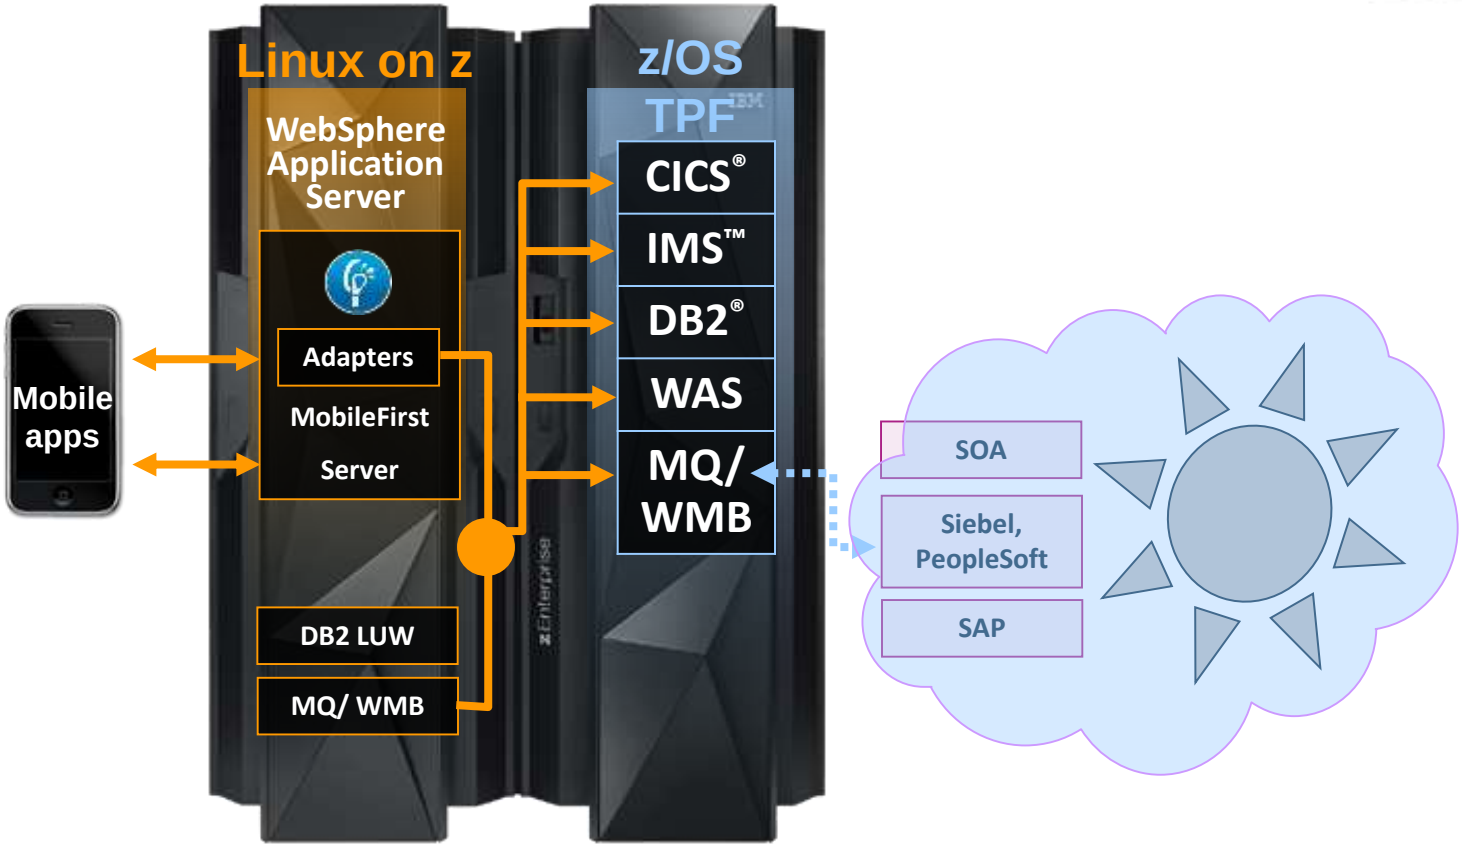

IBM zEnterprise®

- **Server side software components and adapters for channeling z Systems to mobile devices** with IBM MobileFirst Server
- **Mobile application support** with WebSphere Application Server on z Systems
- **Mobile protocol connectivity with core z Systems applications** including CICS, IMS, TPF, MQ, WMB and DB2

Additional information in Mobile Redbooks

- ***Transform Your Organization into a Mobile Enterprise with IBM Worklight***, *Solution Guide,* published 9 October 2013
- ***Extending Your Business to Mobile Devices with IBM Worklight***, SG24-8117-00 ***Red****books,* published 12 August 2013
- ***IBM MobileFirst Strategy Software Approach***, SG24-8191-00 *Draft* ***Red****books, 5* December 2013
- ***IBM z Systems in a Mobile World***, *Solution Guide,* published 21 February 2014
- ***z Systems in a Mobile World***, REDP-5088-00, *Point-of-View,* 24 January 2014
- ***Implementing IBM CICS JSON Web Services for Mobile Applications***, TIPS1066 *Solution Guide,* 9 September 2013
- ***Securing Your Mobile Business with IBM Worklight***, SG24-8179-00, 7 October 2013
- ***Enabling Mobile Apps with IBM Worklight Application Center***, REDP-5005-00 ***Red****papers,* 1 June 2013
- ***Responsive Mobile User Experience Using MQTT and IBM MessageSight***, SG24-8183-00 *Draft* ***Red****books,* last update 18 December 2013
- ***Mobilizing Employees with IBM Notes Traveler***, *Solution Guide,* published 19 February 2013

Questions?

Wilhelm Mild

IBM Executive IT Architect

IBM Deutschland Research & Development GmbH
Schönaicher Strasse 220
71032 Böblingen, Germany

Office: +49 (0)7031-16-3796
wilhelm.mild@de.ibm.com

Wellcome to the Mobile era !

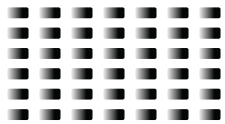

Notices

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

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

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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates 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. 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 IBM's application programming interfaces.

Trademarks

- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

  http://www.ibm.com/legal/copytrade.shtml