
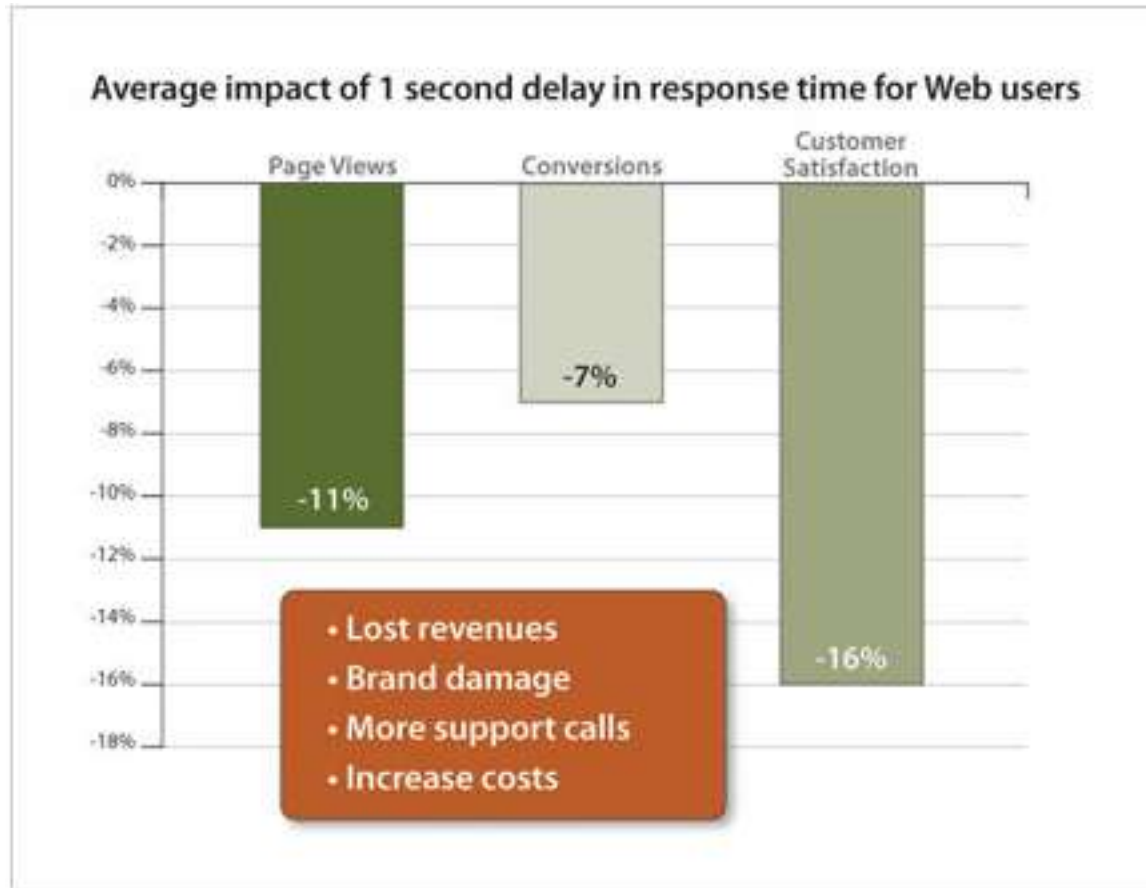


# Data Grids in Your World: WebSphere eXtreme Scale DataPower XC10



# Context



***Aberdeen Group found that an average of one second delay in Web page response time negatively affected page views, conversions and customer satisfaction.<sup>1</sup>***


1. "The Performance of Web Applications: Customers Are Won or Lost in One Second," Bojan Simic, Aberdeen Group, November 2008.

## First, what's a cache?

- A database cache? A page fragment cache? A service Cache?

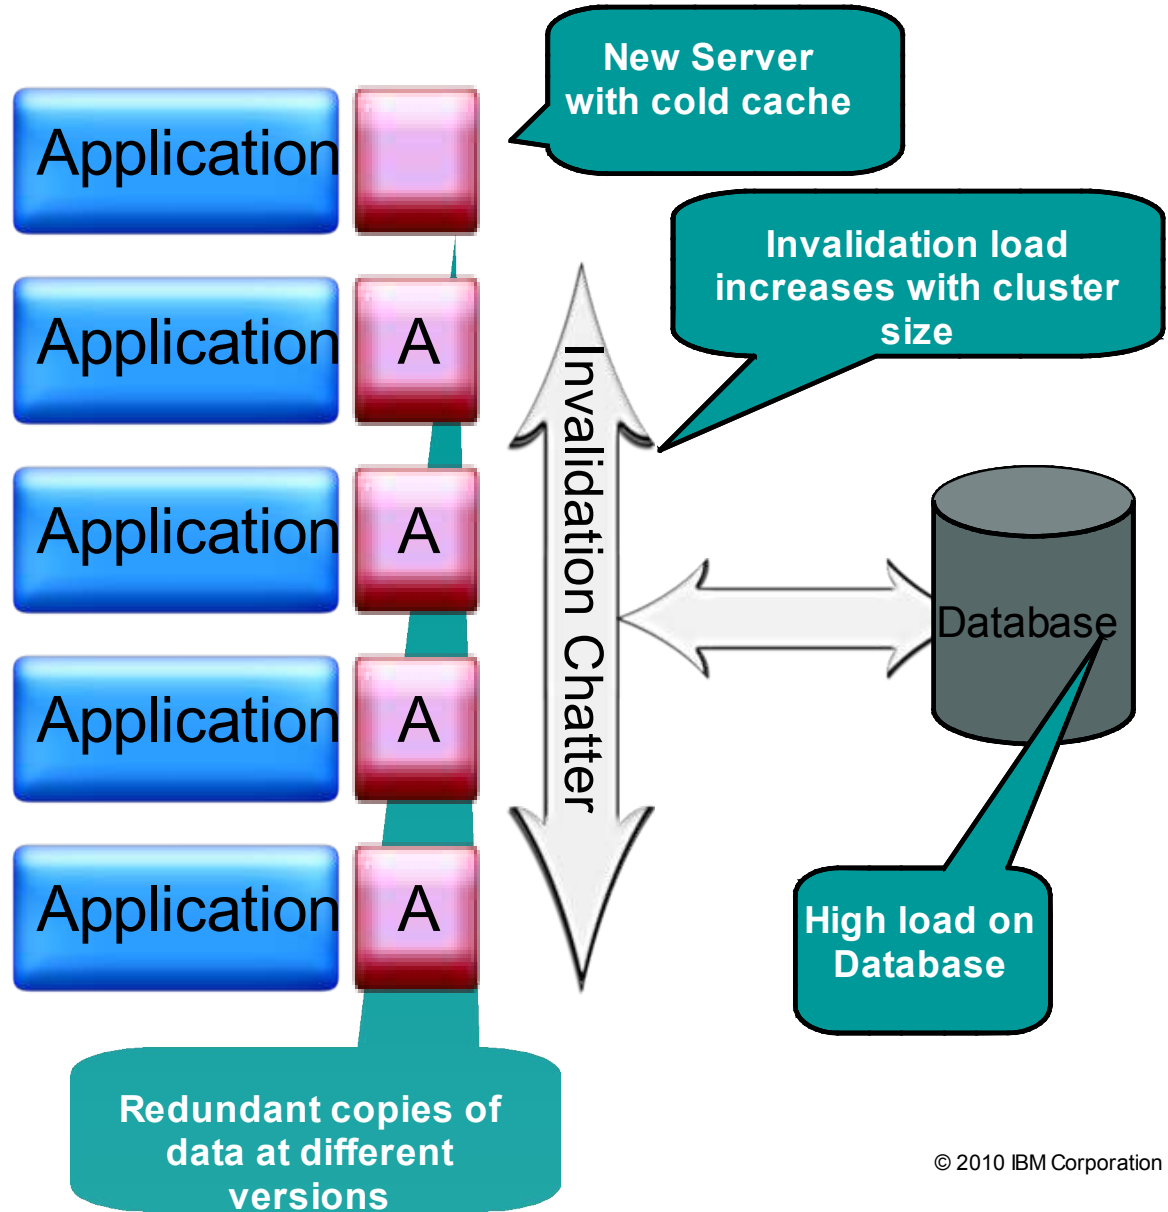
### TOO SPECIFIC!

- A cache is a tool for **reducing application path length**
- OR the **distance data has to travel** before it gets to the customer/ data sink



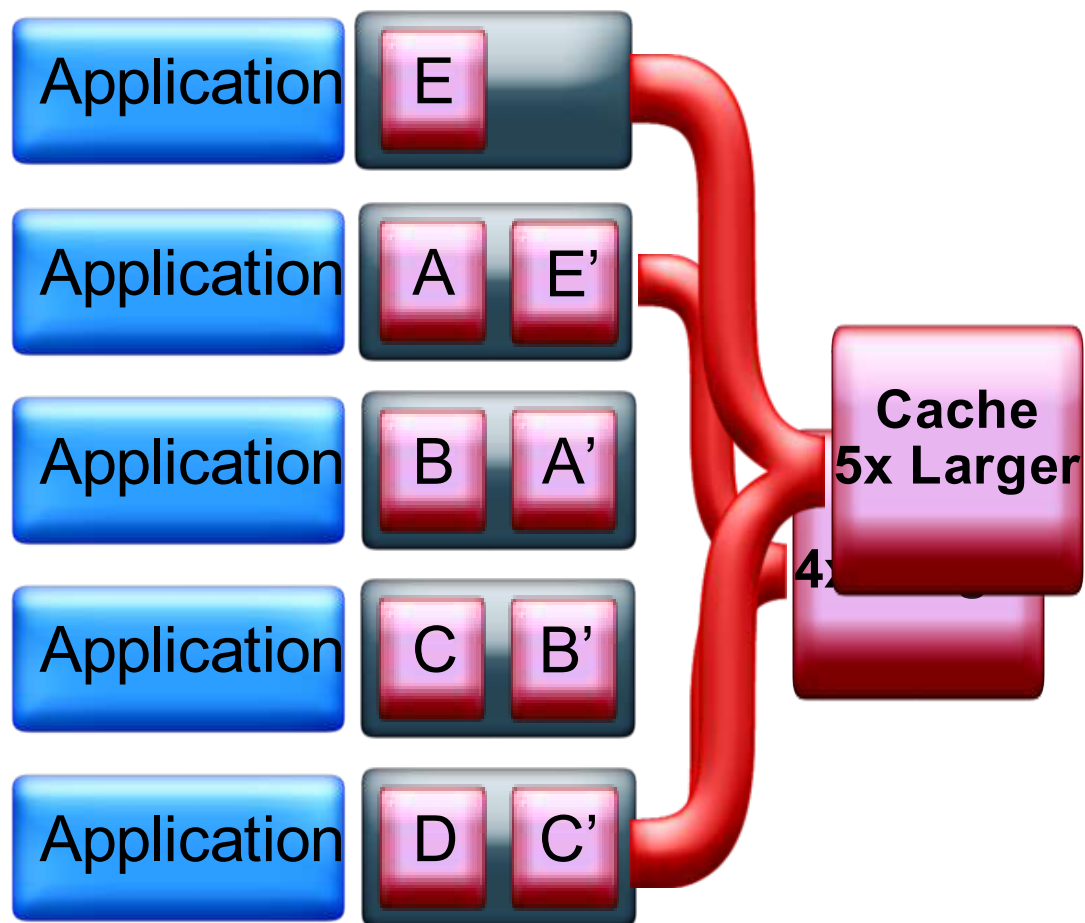
# Traditional Application Cache Operation

- Cache capacity determined by individual JVM Size.
  - Size of each cache =  $M$
  - # JVMs =  $N$
  - **Total cache =  $M$**
- Invalidation load per server increases as cluster grows.
- Cold start servers hit the database.



# WebSphere eXtreme Scale, XC10 Cache Operation

- Cache capacity determined by total cluster size
  - Size of each cache =  $M$
  - # JVMs =  $N$
  - Total Cache =  $M \times N$
- No invalidation chatter
- Linearly scalable
- Less load on database and no cold start spikes



# Modern Application Infrastructure Topology

*Web Server Tier*



**WebSphere. software**

**IBM HTTP Server**

*App Server Tier*



**WebSphere  
Application Server**

*Elastic Data Grid*

**New!**

1



DataPower XC10 for simple data oriented scenarios:

- HTTP Session Replication
- Elastic Dynacache
- Web Side Cache

2

**WebSphere. software**

eXtreme Scale for maximum flexibility covering data and application oriented scenarios

*Database Tier*



**Information Management**

**DB2 UDB**

# Innovative Elastic Caching Solutions



## DataPower XC10 Appliance

- Drop-in cache solution optimized and hardened for data oriented scenarios
- High density, low footprint improves datacenter efficiency

*“Data Oriented”*

- Session management
- Elastic DynaCache
- Web side cache
- Worldwide cache
- Data buffer
- Event Processing
- Petabyte analytics
- In-memory OLTP
- In-memory SOA

*“Application Oriented”*



## eXtreme Scale

- Ultimate flexibility across a broad range of caching scenarios
- In-memory capabilities for application oriented scenarios

*Elastic caching for linear scalability*  
*High availability data replication*  
*Simplified management, monitoring and administration*

## Customer Diversity



# Data Grids are not just for investment banks!

- Retailers
- Insurance Companies
- Health care
- Retail Banking
- Hotel
- Travel Agencies
- Investment Banking
- Telco
- Government
- Utilities
- Trucking Companies



# What is WebSphere eXtreme Scale?

*A flexible framework for realizing high performance, scalable and data-intensive applications*



It can be used as a **very powerful cache** that scales from simple in-process topologies to powerful distributed topologies.

It can be used as a **form of in memory database** to manage application state (and it scales to 1000's of servers).

It can be used as a **platform for building powerful XTP/ Data Grid applications.**

## WXS Installation Options

- Standalone Install
  - Packaged as a single JAR, 15MB total size
  - Requires J2SE, no other requirements
  - Main JARs: ogclient.jar, objectgrid.jar
- Integrate the eXtreme Scale with WebSphere Application Server:
  - Integrates with WAS V6.0, V6.1 and V7.0
  - Main JARs: wsogclient.jar, wsobjectgrid.jar

## Native Integration with WebSphere Application Server

- **eXtreme Scale was designed for native & deep integration with WAS**
- **eXtreme Scale works with ALL WAS 6.X and 7.X versions**
- **eXtreme Scale extends the value of WAS deployments by providing:**
  - Session management plug-in for multi data centre support
  - Dynamic cache plug-in to turbo-charge existing caching environments
  - JPA / Hibernate side-cache to accelerate existing database queries

Session Management  
Plug-in

Dynamic cache  
service Plug-in

SIP Support

JPA / Hibernate  
Side-cache

Global Transaction Integrity

Deep  
Management  
Integration

Programming  
Models

Serviceability

**WebSphere eXtreme Scale**

**WebSphere Application Server**

## Ease of Use/Exploitation

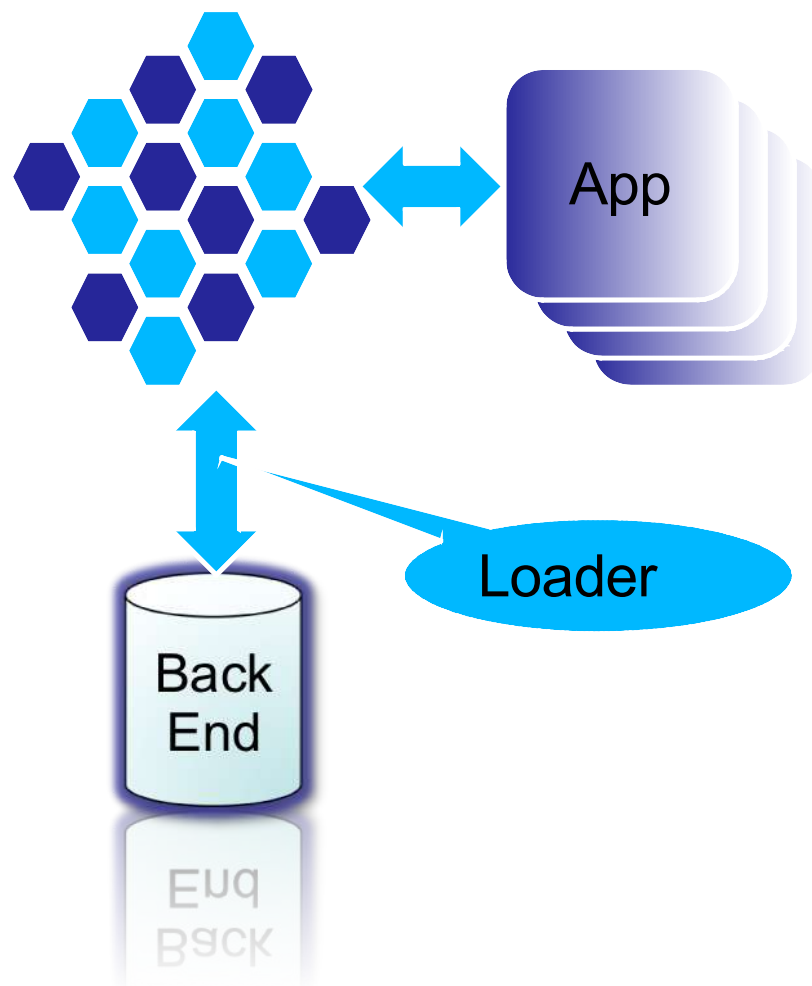
- Dynamic Cache Service plugin (read-only side cache pattern)
  - Configuration only
- Hibernate / OpenJPA L2 cache plugin (read-only side cache pattern)
  - Configuration only
- HTTP Session Management
  - Configuration, run tool on WAR files
- RYO Side cache for:
  - ESB cache mediation
  - SOA result cache
  - WebService/EIS/.../any result cache
  - Insert some code at key points to cache “expensive” results/data

# Ease of Use/Exploitation

- In-line cache, system of record
  - Change app DB access to grid access
  - Also need partitionable (Constrained Tree Schema) data model
- Advanced grid exploitation
  - Move processing to the data
  - Map/Reduce programming model
  - Add advanced APIs to your code

## Common set of patterns

- **Inline backend cache**
  - Loaders used to integrate with an existing data service
  - Read through cache
  - Write through cache
- **System of Record Data Store**
  - Cache is used as the system of record
  - Write behind technology pushes changes asynchronously to the backend.



# Introducing IBM WebSphere DataPower XC10 Appliance

**New!**

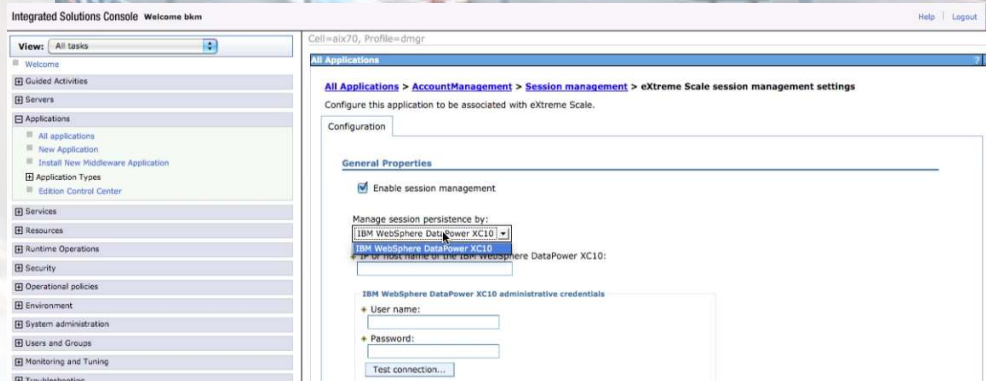
- Easy drop in use for common scenarios
  - Support for **data-oriented caching scenarios** without rip & replace
- Scale out with ease
  - **Large, elastic cache** allows you to scale more economically while providing high Quality of Service
- Fault tolerance
  - Lower risk of data loss while providing **continuous availability**
- Flexible and simple user management
  - Simple solution for **real world management and monitoring**





# XC10: Easy drop-in use for common scenarios

- HTTP Session Management
- WAS Dynamic Cache Service (“Dynacache”) support
- Web Side Cache



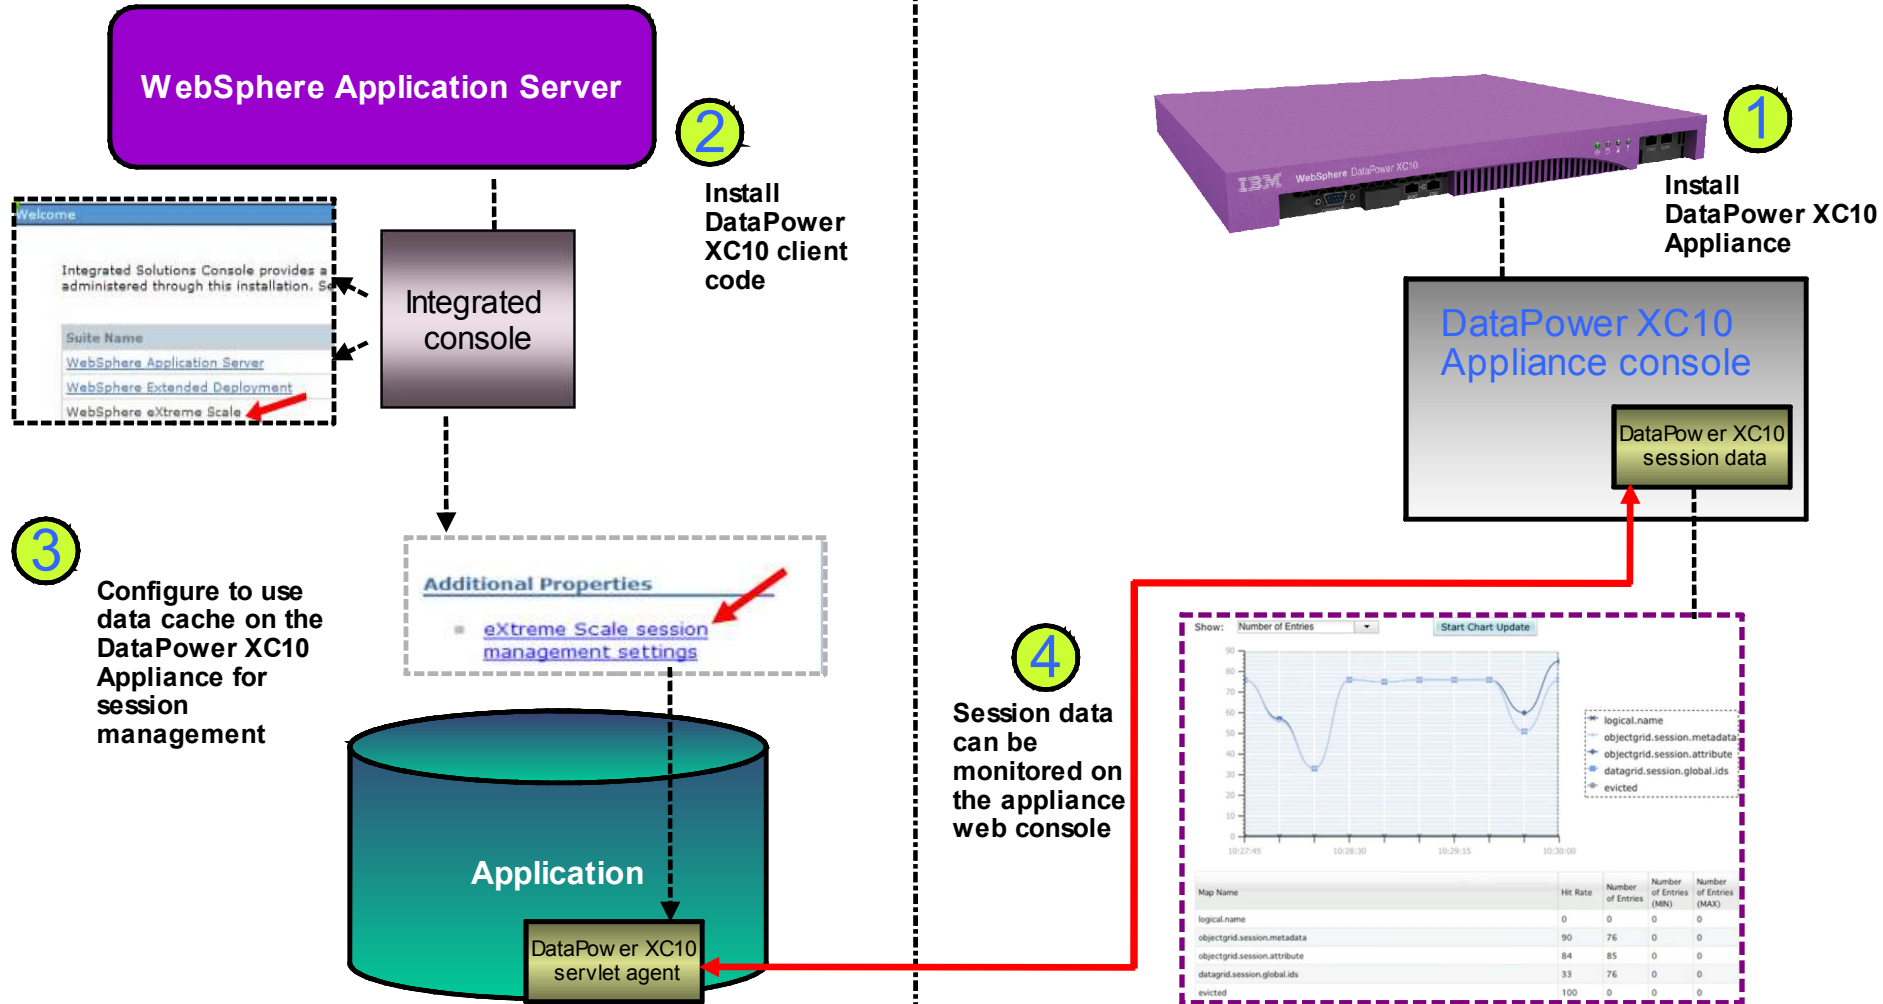
**Little or no code changes required!**



# End-to-end Session Management Scenario

## WebSphere Application Server side

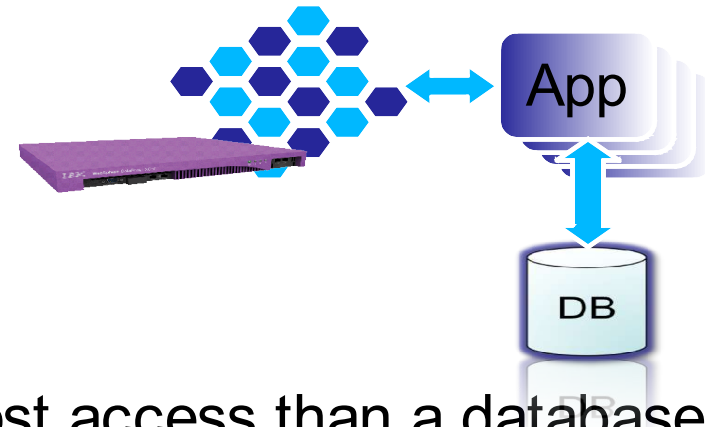
## DataPower XC10 appliance side



## XC10: Dynamic Cache service support

- DataPower XC10 provides client code and a plug-in for WebSphere Application Server applications to support DynaCache API
- Allows applications deployed to WebSphere servers to use DataPower XC10 as a “drop-in” cache, instead of storing cache data in local memory or multiple instances of a disk cache

## XC10: Web Side Cache



- Used to store data for fast, lower-cost access than a database
- Uses ObjectMap APIs from WebSphere eXtreme Scale
- Every time data is needed, the web side cache on the DataPower XC10 Appliance is checked first
- If the value is not found (cache miss), then the data is retrieved from the backend database and inserted into the cache
- Client can run in a standard Java EE compliant server environment or in any Java Virtual Machine compliant with Java SE V1.4 or beyond

# WXS and XC10 Focus on Data-Oriented Distributed Caching Scenarios

