WebSphere Application Server (WAS) for z/OS

Combining Industry Leading SOA Runtime and z/OS to Deliver Superior Customer Value

- **WAS is a cross platform product**

  The Industry's leading application server for building, running & managing business-critical application services

  – All application interface functions and features are available on all platforms

  – Benefits include common programming model, common administrative functions, etc.

- **The *difference* is in the WAS and z/OS platform integration**

  – WAS specifically leverages z/OS and Sysplex capabilities

  – … In an application-transparent fashion

IBM WebSphere Application Server Family
Your Choice of Innovative Performance Based Foundations

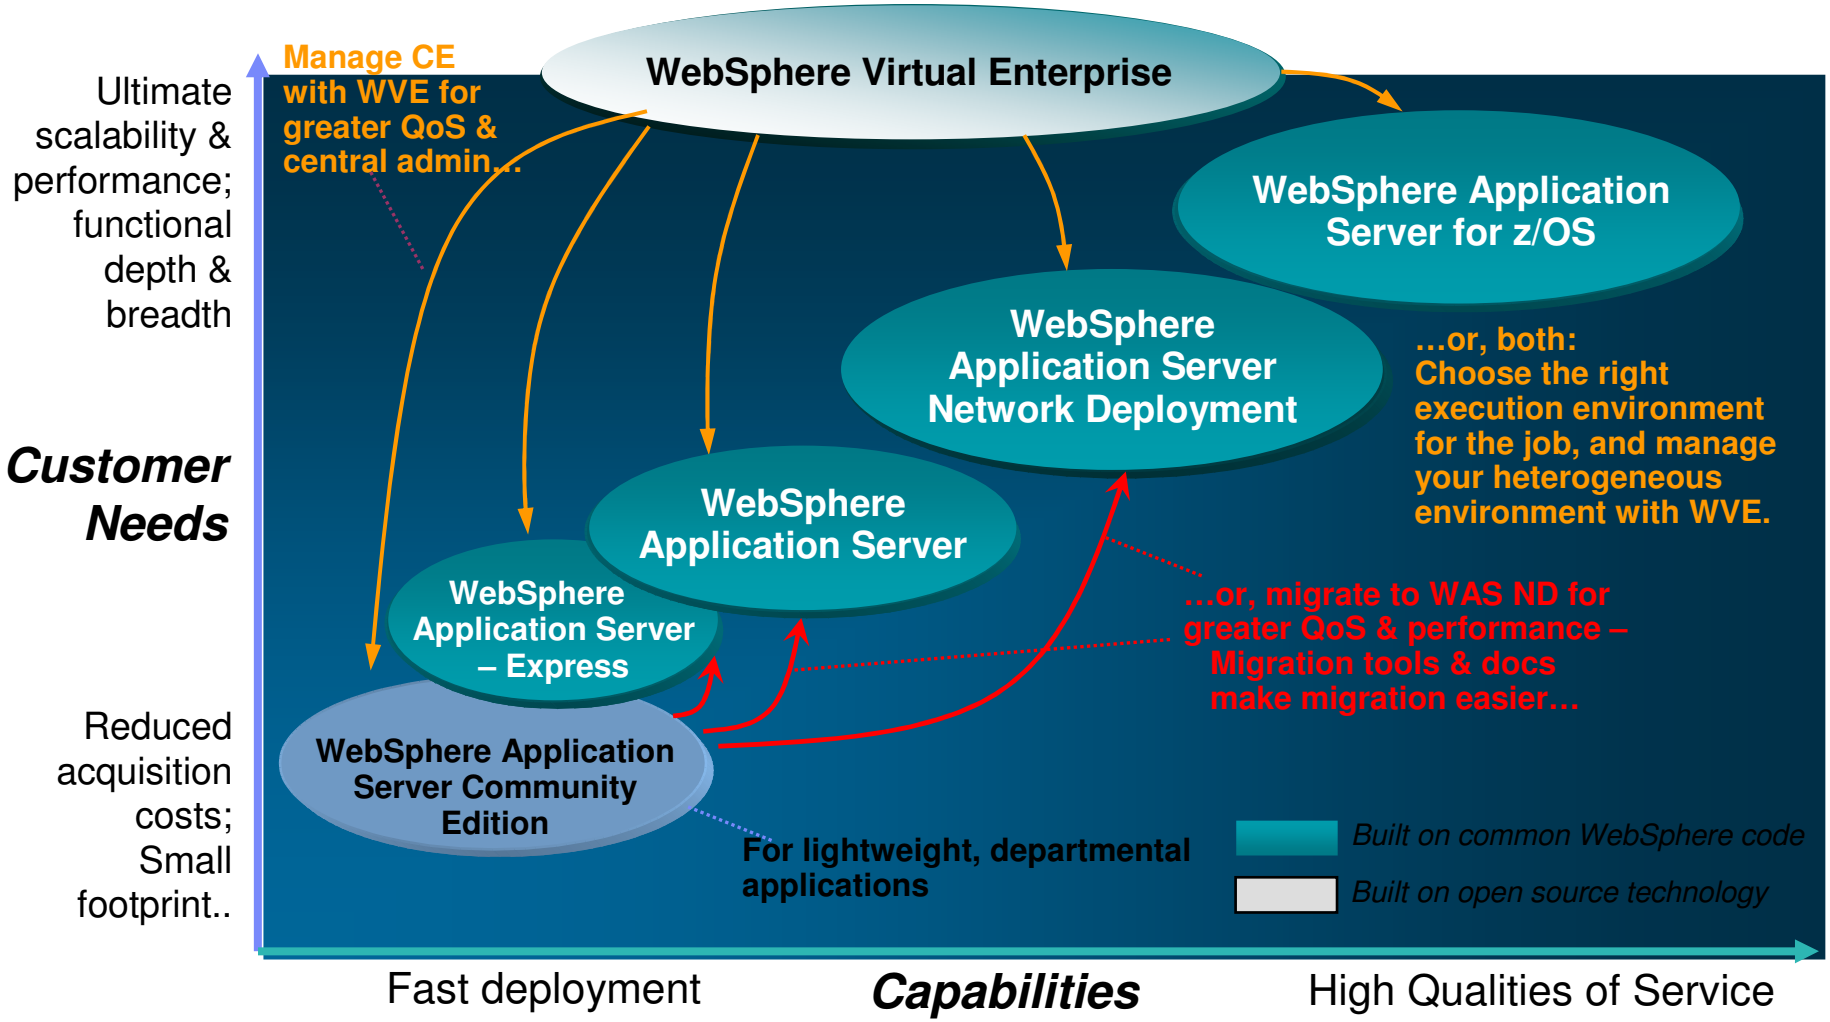

WebSphere Application Server: Cross Platform Common Elements

- Common Programming Model
- Application Touch Interface
- Security Facilities

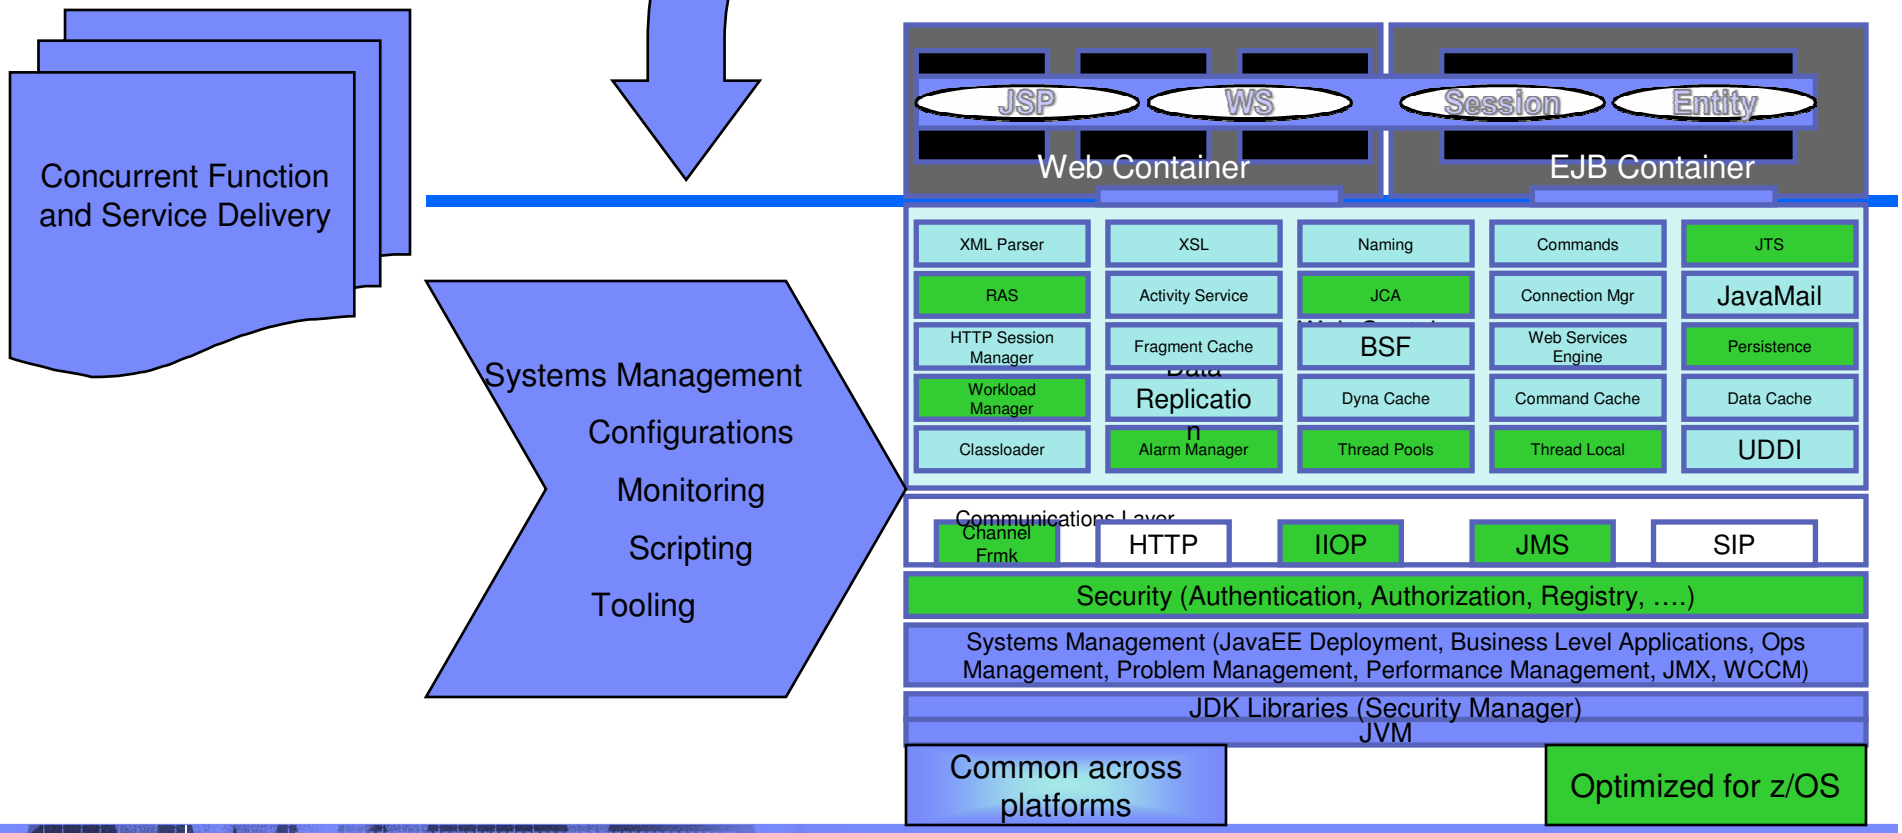

WebSphere Application Server for z/OS: Unique Elements

Integration with z/OS that maintains application transparency

- **Server Architecture**
  - Control/Servant Region Split
- **Workload Management**
  - Leverages Workload Manager
- **Security**
  - Use of the Security Authorization Facility
- **Transaction Management**
  - Leverages Resource Recovery Services
- **Connectors**
  - Leverages available local (Type 2) connectors
- **Thread Management**
  - OS level threads for monitoring and control
- **Scalability**
  - Multiple Servant Region
- **Communications layer**
  - True Asynchio model
- **Recovery**
  - Leverages Automatic Restart Manager
- **Reporting**
  - System Management Facility

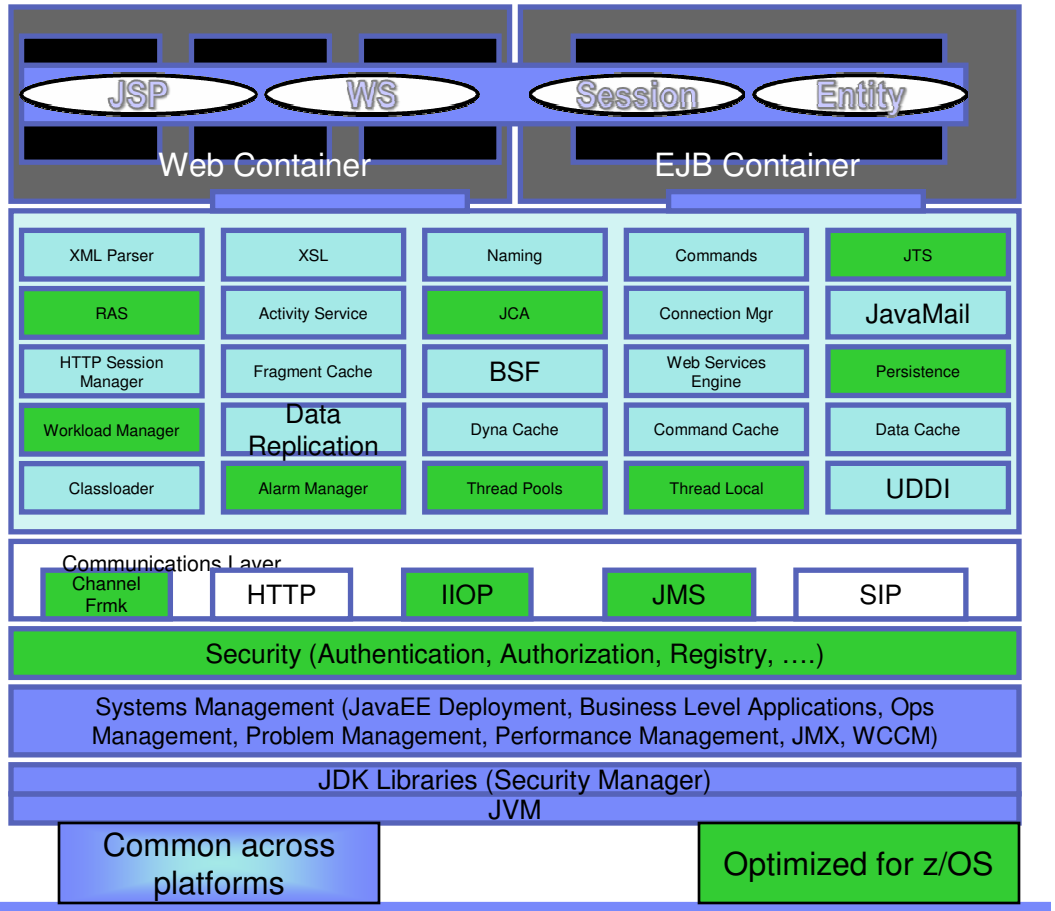

WebSphere Application Server for z/OS V6.1

Integrated into the fabric of the z/OS operating environment

- Replicated Server cluster architecture leveraging shared data for scale and availability
- Optimal access to existing assets through Workload Manager
- No single points of failure, integrated with z/OS recovery mechanisms
- Integrated with local SAF security, application isolation for additional integrity
- Integrated with z/OS automation capabilities for superior manageability
- Comprehensive and recoverable transaction management leveraging RRS and System Logger
- SMF-based comprehensive reporting for capacity chargeback and diagnosis

Hardware, operating system, and middleware working together to bring true 99.999% application availability to your business critical services.

Evolution of WAS

Feature Packs enable you to selectively take advantage of new standards and features while maintaining a more stable internal release cycle.

Delivering ease-of-use improvements in each Feature Pack now and in the future.

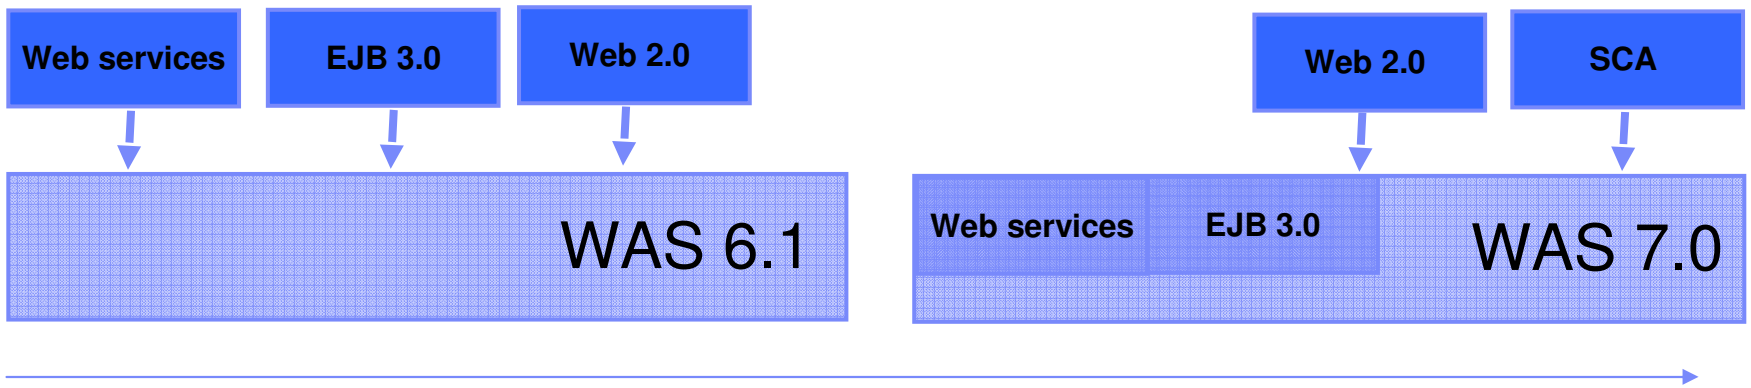

1. Choose the application server technology you need.
2. Install additional functionality on core WAS 7
3. Build the Application Server you want without waiting for new releases.

WebSphere Application Server v7.0 (3Q 2008)

Standards Currency

- Standards currency with Java EE 5, including EJB 3.0, enhances productivity and ease of use.
- New JDK 6 for improved performance and reduced footprint.
- Enhanced Web services standards.

Consumability, Simplicity and Performance Improvements

- Flexible systems management options.
- Enhanced diagnostic tools that help pinpoint problems.
- Security enhancements.
- New virtual image delivery option.
- Tight integration across the WebSphere family of products improves ROI .

New Application Types and Workloads

- Improved performance, scaling and reliability.
- Reduced cost of managing and administering large numbers of individual servers.

Open Beta is now available from WebSphere Software Early Programs website at: www14.software.ibm.com/iwm/web/cc/earlyprograms/websphere.shtml

z/OS Key Differentiation

Performance

- ✓ Improvements in response time for static and dynamic content with Fast Response Cache Acceleration first availability in z/OS 1.9.
- ✓ Increased application runtime performance with focused analysis and code path improvement effort for JEE, Web Services and Connectors.

High Availability and Reliability

- ✓ High Availability Manager based on Cross-System Coupling Facility (XCF).
- ✓ Thread Hang Recovery improves server reliability and performance.

Consumability and Usability

- ✓ Redesigned data collection facility to improve chargeback capabilities.
- ✓ More unified install and configuration tasks (load modules in HFS).

WebSphere AppServer v7.0 Performance Improvement

- DayTrader 1.2 Performance Improvements from WAS v6.1 to v7.0
  - Clear performance improvement by moving to WAS v7.0 for Legacy JEE applications
    - zWAS v7.0 performance is up 42% from v6.1
      - JDK improvements
      - Servant/Controller communication optimizations
      - Codepath improvements throughout WebSphere v7.0
- DayTrader 2.0 EJB3 Performance Improvements from WAS v6.1 FeP to v7.0
  - Large performance improvements by moving to v7.0 from v6.1 FeP for next generation JEE applications
    - WAS v7.0 is 70% faster than v6.1 EJB3 FeP
      - Runtime improvements above compounded with massive improvements in EJB3/OpenJPA code base

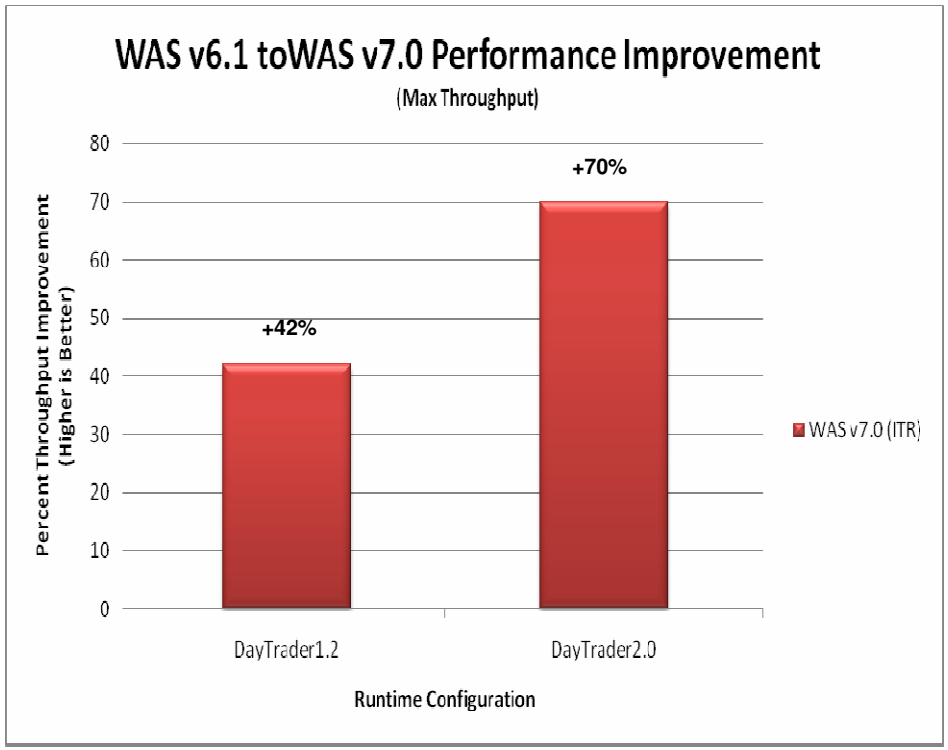

*** Results are on a z9 system
*** Your applications results may vary

**All data is pre-GA WAS v7.0 Info
GA information will be distributed with Performance Report

Differentiation Design Principles

- **z/OS differentiation does not impact programming model**
  - Consistent, full cross platform function available
  - Improved application behavior without application change
- **New functions are shipped for forward compatibility**
  - New functions delivered disabled to minimize deployment impact
  - Almost all new functions are compatible in a mixed cell environment
  - Each new function can be enabled individually using configuration switches
  - Where practical, the function can be enabled/disabled dynamically using an operator command

Leverage unique platform capabilities and qualities of service to deliver additional business value

Function	Business Value	Technical Benefit
Fast Response Cache Acceleration (FRCA) support	**Improved** performance and reduced response time for requests that have response cached.	Transport layer cache used to hold static pages and dynamic content such as servlets and JSPs.
High Availability Manager (HAM) based on Cross-System Coupling Facility (XCF)	**Reduced** overhead for using High Availability Manager on z/OS.	When WebSphere is idle (i.e. not processing transactions) both DCS failure detection and discovery causes unacceptable overhead. XCF is a sysplex wide notification service that will be used in place of TCP/IP heart beat messages.
Thread Hang Recovery	**Improved** recovery, and reliability.	Option for the server to recover a thread that appears to be hung.
New data collection records	**Improved** reporting of zAAP utilization. Consolidated all data needed to provide chargeback function. **Reduced** overhead associated with collecting data	Data collected includes information on server, z/OS, request, network, classification, security and optional CPU usage. Basic information collected when option enabled, additional information can be collected with additional options.
More unified install and configuration tasks	**Improved** consumability	Makes it easier to install and manage. Eliminates a common source of errors. .

Function	Business Value	Technical Benefit
Fast Response Cache Acceleration (FRCA) support	**Improved performance and reduced response time for requests that have response cached.**	**Transport layer cache used to hold static pages and dynamic content such as servlets and JSPs.**
High Availability Manager (HAM) based on Cross-System Coupling Facility (XCF)	Reduced overhead for using High Availability Manager on z/OS.	When WebSphere is idle (i.e. not processing transactions) both DCS failure detection and discovery causes unacceptable overhead. XCF is a sysplex wide notification service that will be used in place of TCP/IP heart beat messages.
Thread Hang Recovery	Improved recovery, and reliability.	Option for the server to recover a thread that appears to be hung.
New data collection records	Improved reporting of zAAP utilization. Consolidated all data needed to provide chargeback function. Reduced overhead associated with collecting data	Data collected includes information on server, z/OS, request, network, classification, security and optional CPU usage. Basic information collected when option enabled, additional information can be collected with additional options.
More unified install and configuration tasks	**Improved** consumability	Makes it easier to install and manage. Eliminates a common source of errors. .

Business Impact of Server Scalability

- **Scalability can be viewed in terms of the end user**
  - How many clients can connect and use the server?
  - What is the response time?
  - What is the cost to support each client?
- **Server scalability has its limits**
  - Each client request holds and uses resources within the server for the duration of its processing
  - As the number of concurrent client requests increases, the server continues to use more resources until it reaches a maximum the server can handle
  - Depending on the server and the nature of the application, response time can slow down due to contention for shared resources, and the throughput of the server can suffer
- **Possible ways to extend server scalability**
  - Replicate the server and split the workload
    - May solve the number of connections and the response time, but it increases the per unit cost to support each client by maintaining multiple servers
  - Reduce the time it takes to process a request in the server
    - For some workloads, caching a response and reusing it rather than building the response from scratch significantly reduces time to process requests

Caching

- **What is caching, and why do I care?**
  - Caching is a way to use memory to save CPU or processing time
  - Studies have shown that use of memory as a cache can have a big impact on throughput and response time of a server for e-business workloads
    - http://www.ibm.com/developerworks/websphere/techjournal/0405_hines/0405_hines.html
- **Caching is used today in multiple places within a configuration.**
  - WebSphere Application Server
    - Dynacache provides ability to cache Servlets/JSPs, commands, and Web Services responses
    - Dynacache supported on all platforms
  - IBM HTTP Server
    - Caches static files
    - Supported on all platforms

Fast Response Cache Acceleration (FRCA) Support

- **What It is?**
  - FRCA is integrated caching technology in the z/OS TCP/IP stack
  - FRCA is supported by the IBM HTTP Server, but only for static content
  - WAS 7.0 supports direct exploitation of FRCA in conjunction with existing Dynacache capability
    - FRCA holds static pages and dynamic content such as servlets and JSPs
- **Customer Value**
  - Improved response time on web requests
  - Reduced CPU load
    - Initial prototype shows significant improvement over dynamic cache
  - Simplified administration

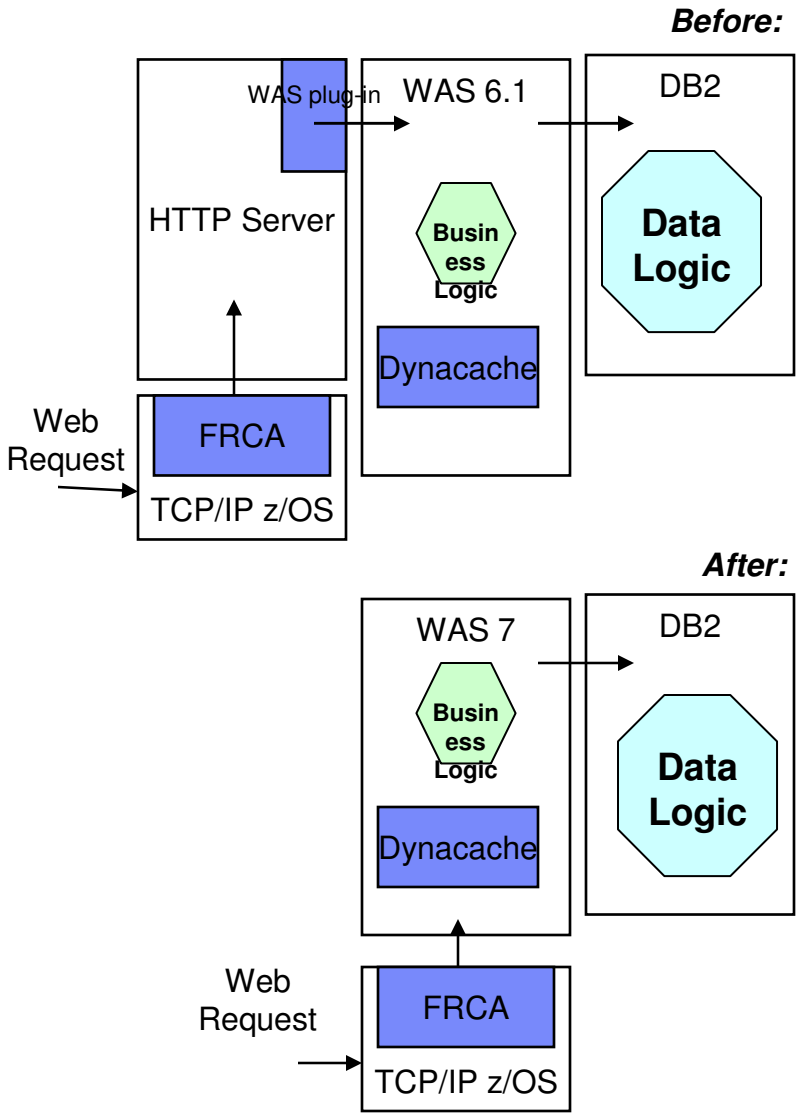

Function	Business Value	Technical Benefit
Fast Response Cache Acceleration (FRCA) support	Improved performance and reduced response time for requests that have response cached.	Transport layer cache used to hold static pages and dynamic content such as servlets and JSPs.
High Availability Manager (HAM) based on Cross-System Coupling Facility (XCF)	**Reduced overhead for using High Availability Manager on z/OS.**	**When WebSphere is idle (i.e. not processing transactions) both DCS failure detection and discovery causes unacceptable overhead.** **XCF is a sysplex wide notification service that will be used in place of TCP/IP heart beat messages.**
Thread Hang Recovery	Improved recovery, and reliability.	Option for the server to recover a thread that appears to be hung.
New data collection records	Improved reporting of zAAP utilization. Consolidated all data needed to provide chargeback function. Reduced overhead associated with collecting data	Data collected includes information on server, z/OS, request, network, classification, security and optional CPU usage. Basic information collected when option enabled, additional information can be collected with additional options.
More unified install and configuration tasks	Improved Consumability	Makes it easier to install and manage. Eliminates a common source of errors. .

Business Impact of Inefficient Processing

- **What are "idle MSUs", and why do I care?**
  - When a server is active but not processing any work,
    - it should consume little or no processing resources
    - What is consumed while waiting for work are "idle MSUs"
  - Within a highly managed environment, processing resources can and should be fully utilized
    - "Idle MSUs" detract from the full and effective utilization of the resources available
    - "Idle MSUs" have a measurable cost or dollar value

"Idle MSUs" are not being used to support your business!

High Availability Manager Based on Heartbeats

- **High Availability Manager (HAM)**
  - A function within the WebSphere Application Server.
  - It detects when a member within the cluster has failed, and takes recovery action.
- **High Availability Manager was originally implemented using TCP/IP heartbeats**
  - For effective detection of a failure, heartbeats need to be regular & frequent
    - the longer the time between heartbeats, the longer it takes to detect a failure
  - Heartbeats consume resources even when the server is idle
  - Heartbeats consume more resources as the topologies scale up
- **Reducing the impact of the High Availability Manager**
  - Turn off HAM when not needed,
    - However, certain Application Server functions do require that HAM be active, e.g. Data Replication Services
  - Offer a more efficient failure detection methodology
    - Use the z/OS Sysplex service called Cross System Coupling Facility (XCF)
    - Uses notification rather than heartbeats which requires less overhead
    - XCF sends notification to group members when there is a change in member status

High Availability Manager Based on XCF

- **What It Is**
  - WAS Distribution and Consistency Services (DCS)
    - Distribution of information among a set of members who belong to groups
    - Failure detection of said members/groups
    - Forms the infrastructure utilized by the High Availability Manager
    - DCS can optionally use XCF or Heartbeats
  - High Availability Manager on z/OS starting in Version 7
    - Can continue to use Heartbeats
    - Can optionally use XCF when all members are at Version 7
- **Customer Value**
  - XCF also provides superior solution for recovery
  - Reduced CPU consumption

Function	Business Value	Technical Benefit
Fast Response Cache Acceleration (FRCA) support	Improved performance and reduced response time for requests that have response cached.	Transport layer cache used to hold static pages and dynamic content such as servlets and JSPs.
High Availability Manager (HAM) based on Cross-System Coupling Facility (XCF)	Reduced overhead for using High Availability Manager on z/OS.	When WebSphere is idle (i.e. not processing transactions) both DCS failure detection and discovery causes unacceptable overhead. XCF is a sysplex wide notification service that will be used in place of TCP/IP heart beat messages.
Thread Hang Recovery	**Improved recovery, and reliability.**	**Option for the server to recover a thread that appears to be hung.**
New data collection records	Improved reporting of zAAP utilization. Consolidated all data needed to provide chargeback function. Reduced overhead associated with collecting data	Data collected includes information on server, z/OS, request, network, classification, security and optional CPU usage. Basic information collected when option enabled, additional information can be collected with additional options.
More unified install and configuration tasks	Improved Consumability	Makes it easier to install and manage. Eliminates a common source of errors. .

Business Impact of Server Availability

- **Application issues and software errors are inevitable**
  - No single process can claim 100% availability
- **Strategy for high availability: Run multiple processes**
  - When a process becomes unavailable, work is routed to the another process
  - Cost in time and resources to recover original process
  - Additional load on remaining processes potentially impacts response time and throughput
- **What happens to the work in the unavailable process?**
  - All active transactions fail
    - Failures may be visible to users and may result in loss of business
  - All in-flight work needs to be “un-done”

Thread Hang Recovery

- **What It Is**
  - Mechanism to interrupt errant applications and return thread to take new work from queue
  - Other work continues
- **Customer Value**
  - Improved Failover/ recovery
  - Improved Reliability
  - Improved performance
    - avoid starting replacement process

WebSphere Application Server for z/OS

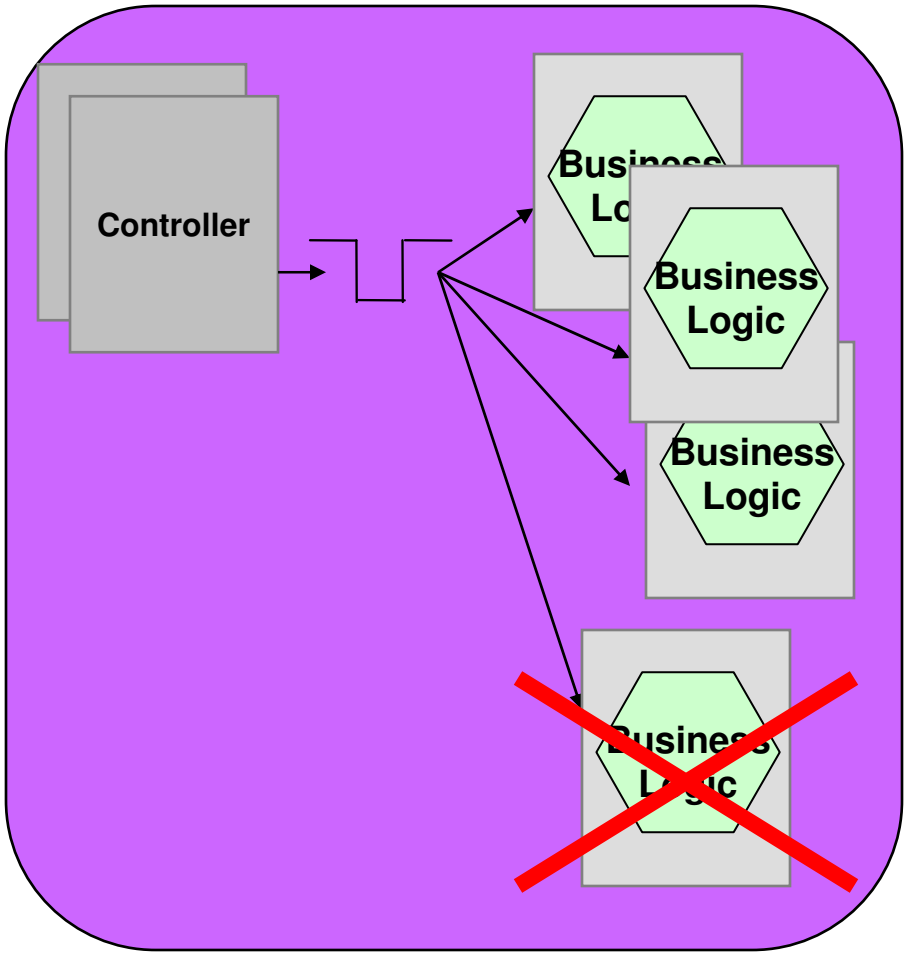

Thread Hangs

- **What is a thread hang, and why do I care?**
  - A thread hang occurs when a request is waiting for an external event to happen and it never does or takes a *really* long time
    - Example: call to off platform via TCP/IP doesn't return due to network problems, or problems with down level system
  - Hung threads hold onto resources, so accumulating them results in degraded performance that escalates over time
  - Recovery is to recycle the process with the hung thread
    - In a single server process this means recycling the entire server
    - On z/OS, this means recycling a servant process. Since the server is made up of multiple processes, a recycle of a servant does not result in the server becoming unavailable.

Thread Hang Recovery on z/OS

- **What is thread hang recovery, and why do I care?**
  - Thread hang recovery is a new feature of WAS on z/OS that provides an infrastructure to “shake loose” a thread that is hung
  - This means that server resources will not be drained, and a server recycle can be avoided.
    - Server continues to operate normally
    - Other “innocent” work in the server is not affected
  - Thread hang recovery also includes a mechanism to attempt to stop a “runaway” thread
    - “Run-away” threads are worse than “hung” threads in that they not only hold resources, but eat up CPU cycles as well.
    - The same recovery actions apply to these as well.

Function	Business Value	Technical Benefit
Fast Response Cache Acceleration (FRCA) support	Improved performance and reduced response time for requests that have response cached.	Transport layer cache used to hold static pages and dynamic content such as servlets and JSPs.
High Availability Manager (HAM) based on Cross-System Coupling Facility (XCF)	Reduced overhead for using High Availability Manager on z/OS.	When WebSphere is idle (i.e. not processing transactions) both DCS failure detection and discovery causes unacceptable overhead. XCF is a sysplex wide notification service that will be used in place of TCP/IP heart beat messages.
Thread Hang Recovery	Improved recovery, and reliability.	Option for the server to recover a thread that appears to be hung.
New data collection records	**Improved reporting of zAAP utilization. Consolidated all data needed to provide chargeback function.** **Reduced overhead associated with collecting data**	**Data collected includes information on server, z/OS, request, network, classification, security and optional CPU usage.** **Basic information collected when option enabled, additional information can be collected with additional options.**
More unified install and configuration tasks	Improved Consumability	Makes it easier to install and manage. Eliminates a common source of errors. .

Charge Backs

- **What are charge backs and why do you care?**
  - Businesses are usually organized as units each with a focus on different aspects of the business, and as such, run different workloads
  - Each workload associated with a business unit uses IT resources which impacts the IT budget
  - As SOA principles and virtualization techniques come into play, resources are less and less dedicated to specific workloads and business units
  - In these shared multiple workload system environments, the ability to understand what workload is using which resources is key to calculating the cost to run the business unit.
  - The cost of the resources used can be “charged back” to the business unit receiving the benefit to ensure an accurate business view
  - Being able to easily collect charge back data reduces the overhead of the basic processing required to run a large business

New data collection records

- **What is it?**
  - Prior to v7:
    - To provide effective charge backs 2 record types needed to be collected
    - Too much overhead / impact to system
    - Customers creating alternative methods or using very course grain data
  - New with v7:
    - New subtype record 9 of SMF 120 Records
      - More efficient collection mechanism using in-memory data structures
      - Basic information collected by default
      - Additional information optionally collected
        - Note: Incremental data collection may incur additional overhead
- **Customer value**
  - All data needed for charge backs collected in one record type provides more accurate accounting
  - Additional data provided (zAAP utilization)
  - Reduced overhead of collecting data

Function	Business Value	Technical Benefit
Fast Response Cache Acceleration (FRCA) support	Improved performance and reduced response time for requests that have response cached.	Transport layer cache used to hold static pages and dynamic content such as servlets and JSPs.
High Availability Manager (HAM) based on Cross-System Coupling Facility (XCF)	Reduced overhead for using High Availability Manager on z/OS.	When WebSphere is idle (i.e. not processing transactions) both DCS failure detection and discovery causes unacceptable overhead. XCF is a sysplex wide notification service that will be used in place of TCP/IP heart beat messages.
Thread Hang Recovery	Improved recovery, and reliability.	Option for the server to recover a thread that appears to be hung.
New data collection records	Improved reporting of zAAP utilization. Consolidated all data needed to provide chargeback function. Reduced overhead associated with collecting data	Data collected includes information on server, z/OS, request, network, classification, security and optional CPU usage. Basic information collected when option enabled, additional information can be collected with additional options.
More unified install and configuration tasks	**Improved Consumability**	**Makes it easier to install and manage. Eliminates a common source of errors.** **.**

More unified install and configuration tasks

- **What is it?**
  - Prior to Version 7:
    - WebSphere Application Server for z/OS ships as both HFS content and MVS datasets (different from the other platform packaging)
    - Requires the two sets of executables to be kept in sync
    - Potential for “difficult to debug” code-level-mismatch problems
  - New with Version 7:
    - Load modules will ship in the AppServer/lib directory with the rest of the runtime

- **Customer Value**
  - Reduced installation and management complexity
  - Elimination of a common source of errors

Summary Of Customer Value Differentiation for WAS for z/OS 7.0

Business Value
Improved* performance and reduced response time for requests that have response cached with support for FRCA.
Reduced* overhead for using High Availability Manager on z/OS.
Improved* failover, recovery, reliability and performance with capability for thread hang recovery.
Reduced* overhead associated with collecting SMF records & improved reporting of zAAP utilization.
Improved* consumability by a more unified install and elimination of common errors caused by 2 sets of executables.

WebSphere Application Server (WAS) for z/OS

Combining Industry Leading SOA Runtime and z/OS to Deliver Superior Customer Value

- **WAS is a cross platform product**

  The Industry's leading application server for building, running & managing business-critical application services

  – All application interface functions and features are available on all platforms

  – Benefits include common programming model, common administrative functions, etc.

- **The *difference* is in the WAS and z/OS platform integration**

  – WAS specifically leverages z/OS and Sysplex capabilities

  – … In an application-transparent fashion

Thank you

Legal Disclaimer

- THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

- REFERENCES IN THIS PRESENTATION TO IBM PRODUCTS, PROGRAMS, OR SERVICES DO NOT IMPLY THAT THEY WILL BE AVAILABLE IN ALL COUNTRIES IN WHICH IBM OPERATES. PRODUCT RELEASE DATES AND/OR CAPABILITIES REFERENCED IN THIS PRESENTATION MAY CHANGE AT ANY TIME AT IBM'S SOLE DISCRETION BASED ON MARKET OPPORTUNITIES OR OTHER FACTORS, AND ARE NOT INTENDED TO BE A COMMITMENT TO FUTURE PRODUCT OR FEATURE AVAILABILITY IN ANY WAY. NOTHING CONTAINED IN THESE MATERIALS IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, STATING OR IMPLYING THAT ANY ACTIVITIES UNDERTAKEN BY YOU WILL RESULT IN ANY SPECIFIC SALES, REVENUE GROWTH OR OTHER RESULTS. PERFORMANCE IS BASED ON MEASUREMENTS AND PROJECTIONS USING STANDARD IBM BENCHMARKS IN A CONTROLLED ENVIRONMENT. THE ACTUAL THROUGHPUT OR PERFORMANCE THAT ANY USER WILL EXPERIENCE WILL VARY DEPENDING UPON MANY FACTORS, INCLUDING CONSIDERATIONS SUCH AS THE AMOUNT OF MULTIPROGRAMMING IN THE USER'S JOB STREAM, THE I/O CONFIGURATION, THE STORAGE CONFIGURATION, AND THE WORKLOAD PROCESSED. THEREFORE, NO ASSURANCE CAN BE GIVEN THAT AN INDIVIDUAL USER WILL ACHIEVE RESULTS SIMILAR TO THOSE STATED HERE.

End of Presentation