

System z Premier Executive Event



New Innovative High Availability Solutions

Dave Petersen
IBM Distinguished Engineer
IBM Systems and Technology Group



Disclaimer


- **All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these statements of general direction is at the relying party's sole risk and will not create liability or obligation for IBM.**
- **This presentation is being given to you under the condition that its contents are not disclosed to further parties.**

Active/Active Sites Overview - Agenda

- **Level Set**
- **Active/Active Sites Overview**
- **Components**
- **Roadmap**




System z
environments





Active/Active Sites Overview



Level Set

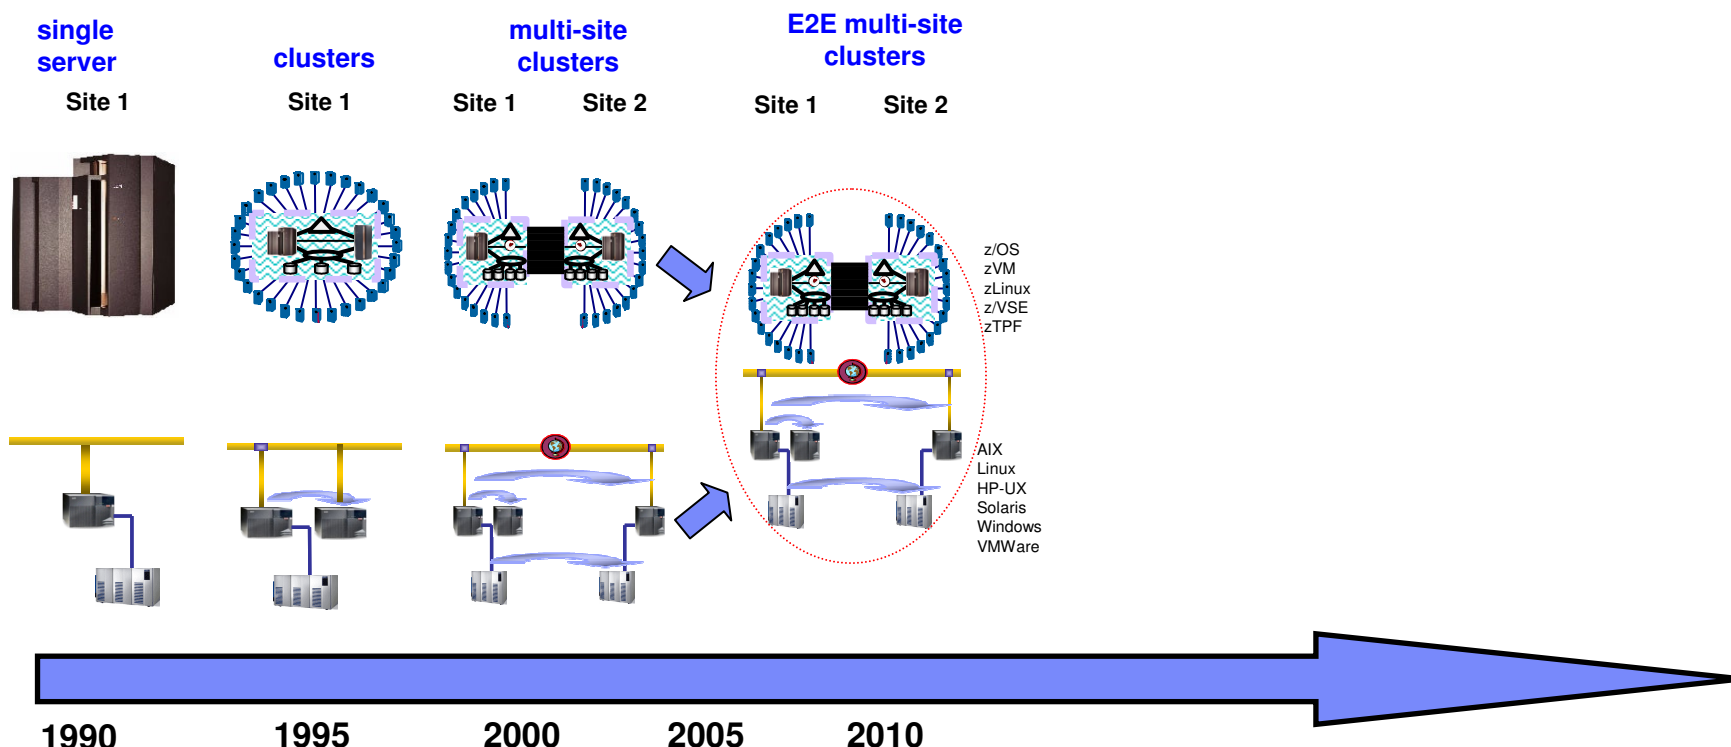
Active/Active Sites Overview

Components

Roadmap



Enterprise Wide BC Solution Direction



- Built In Redundancy
- Capacity Backup
- Hot Pluggable I/O
- Sys z RPO 48-72H / RTO 48-72H
- Provides CA / DR
- Sync/Asynch Data Mirroring
- Eliminates Tape/Disk SPoF
- Application Independent
- Sys z RPO 0-3 sec / RTO < 1H
- Planned/unplanned HW/SW outages
- Flexible, non-disruptive growth
- Dynamic Workload and Resource Management
- Sys z RPO 48-72H / RTO 48-72H
- Provides E2E CA / DR
- Automated failover/switchover
- Browser-based Administration
- E2E RPO 0-3 sec / RTO < 1H

Business Continuity

▪ The top Causes of Business Interruption:

- Planned Maintenance
 - System and Software Upgrades or Reconfiguration
 - Database Administration
- Component Failure
 - Caused by Operator Errors, Software defects, Disk Failure, Subsystems, Hardware, Power Grid
 - Data is recoverable.
 - But, changes might be stranded until component is restored
- Disaster
 - Flood, Earthquake, Fire, ..., Loss of a site
 - Data is not recoverable

Establishing the Objectives:

- *Recovery Time Objective (RTO)* = How much time is needed to restore business operations?
- *Recovery Point Objective (RPO)* = How much data could we afford to *lose*?

How Much Interruption can your Business Tolerate?

Ensuring Business Continuity:

Standby

- **Disaster Recovery**

- Restore business after an unplanned outage

- **High-Availability**

- Meet Service Availability objectives e.g., 99.9% availability or 8.8 hours of down-time a year

- **Continuous Availability**

- No downtime (planned or not)

Active/Active

Global Enterprises that operate across time-zones no longer have any 'off-hours' window. Continuous Availability is required.

What is the cost of 1 hour of downtime during core business hours?

Cost of Downtime by Industry

Industry Sector	Loss per Hour
Financial	\$8,213,470
Telecommunications	\$4,611,604
Information Technology	\$3,316,058
Insurance	\$2,582,382
Pharmaceuticals	\$2,058,710
Energy	\$1,468,798
Transportation	\$1,463,128
Banking	\$1,145,129
Chemicals	\$1,071,404
Consumer Products	\$989,795

Source: Robert Frances Group 2006, "Picking up the value of PKI: Leveraging z/OS for Improving Manageability, Reliability, and Total Cost of Ownership of PKI and Digital Certificates."

Disruptions affect more than the bottom line...

September 9, 2008

**London Stock Exchange
Paralyzed by Glitch**

THE WALL STREET JOURNAL.

June 6, 2008

**Amazon website hit by technical failure –
shares fall 4.1% by mid-afternoon trade**

REUTERS

April 28, 2008

CBSA responds to recent system outages

Canada Border Services Agency

August 4, 2010

**Singapore Censures DBS Bank For
System Outage On July, 5 2010**

GOVMONITOR
Public Sector News & Information

... with enormous impact on the business

- Downtime costs can equal up to 16 percent of revenue ¹
- 4 hours of downtime severely damaging for 32 percent of organizations, ²
- Data is growing at explosive rates – growing from 161EB in 2007 to 988EB in 2010³
- Some industries fine for downtime and inability to meet regulatory compliance
- Downtime ranges from 300–1,200 hours per year, depending on industry¹

¹ Infonetics Research, *The Costs of Enterprise Downtime: North American Vertical Markets 2005*, Rob Dearborn and others, January 2005.

² Continuity Central, "Business Continuity Unwrapped," 2006, <http://www.continuitycentral.com/feature0358.htm>

³ The Expanding Digital Universe: A Forecast of Worldwide Information Growth Through 2010, IDC white paper #206171, March 2007

What are customers doing for Business Continuity today ?



Continuous Availability of Data within a Data Center


Continuous Availability / Disaster Recovery within a Metropolitan Region

Disaster Recovery at Extended Distance

Continuous Availability Regionally and Disaster Recovery Extended Distance

**Single Data Center
Applications remain active**


Continuous access to data in the event of a storage subsystem outage



**GDPS/HyperSwap Mgr
RPO=0 & RTO=0**

**Two Data Centers
Systems remain active**


Multi-site workloads can withstand site and/or storage failures



**GDPS/PPRC
RPO=0 & RTO<1 hr**

**Two Data Centers
Rapid Systems Disaster Recovery with "seconds" of Data Loss**


Disaster recovery for out of region interruptions



**GDPS/GM & GDPS/XRC
RPO secs & RTO <1 hr**

**Three Data Centers
High availability for site disasters**

Disaster recovery for regional disasters



GDPS/MGM & GDPS/MzGM

Customer Requirements

- **Want to shift focus from a failover model to a nearly-continuous availability model (RTO near zero)**
- **Access data from any site (unlimited distance between sites)**
- **Multi-sysplex, multi-platform solution**
 - “Recover my business rather than my platform technology”
- **Ensure successful recovery via automated processes (similar to GDPS technology today).**
 - Can be handled by less-skilled operators
- **Provide workload distribution between sites (route around failed sites, dynamically select sites based on ability of site to handle additional workload).**
- **Provide application level granularity**
 - Some workloads may require immediate access from every site, other workloads may only need to update other sites every 24 hours (less critical data).
 - Current solutions employ an all-or-nothing approach (complete disk mirroring, requiring extra network capacity).

What are GDPS/PPRC customers doing today?



- **GDPS/PPRC, based upon a multi-site Parallel Sysplex and synchronous disk replication, is a metro area Continuous Availability (CA), Disaster Recovery solution (DR)**
- **GDPS/PPRC supports two configurations:**
 - Active/standby
 - Active/active
- **Some customers have deployed GDPS/PPRC active/active configurations**
 - All critical data must be PPRCed and HyperSwap enabled
 - All critical CF structures must be duplexed
 - Applications must be parallel sysplex enabled
 - Signal latency will impact OLTP thru-put and batch duration resulting in the sites being separated by no more than a couple tens of KM (fiber)
- ***Issue:* the GDPS/PPRC active/active configuration does not provide enough site separation for some enterprises**

What are GDPS/XRC & GDPS/GM customers doing today ?



- **GDPS/XRC and GDPS/GM, based upon asynchronous disk replication, are unlimited distance DR solutions**
- **The current GDPS async replication products require the failed site's workload to be restarted in the recovery site and this typically will take 30-60 min**
 - Power fail consistency
 - Transaction consistency
- **There are no identified extensions to the existing GDPS async replication products that will allow the RTO to be substantially reduced.**
- ***Issue:* GDPS/XRC and GDPS/GM will not achieve an RTO of seconds being requested by some enterprises**



Active/Active Sites Overview

Level Set



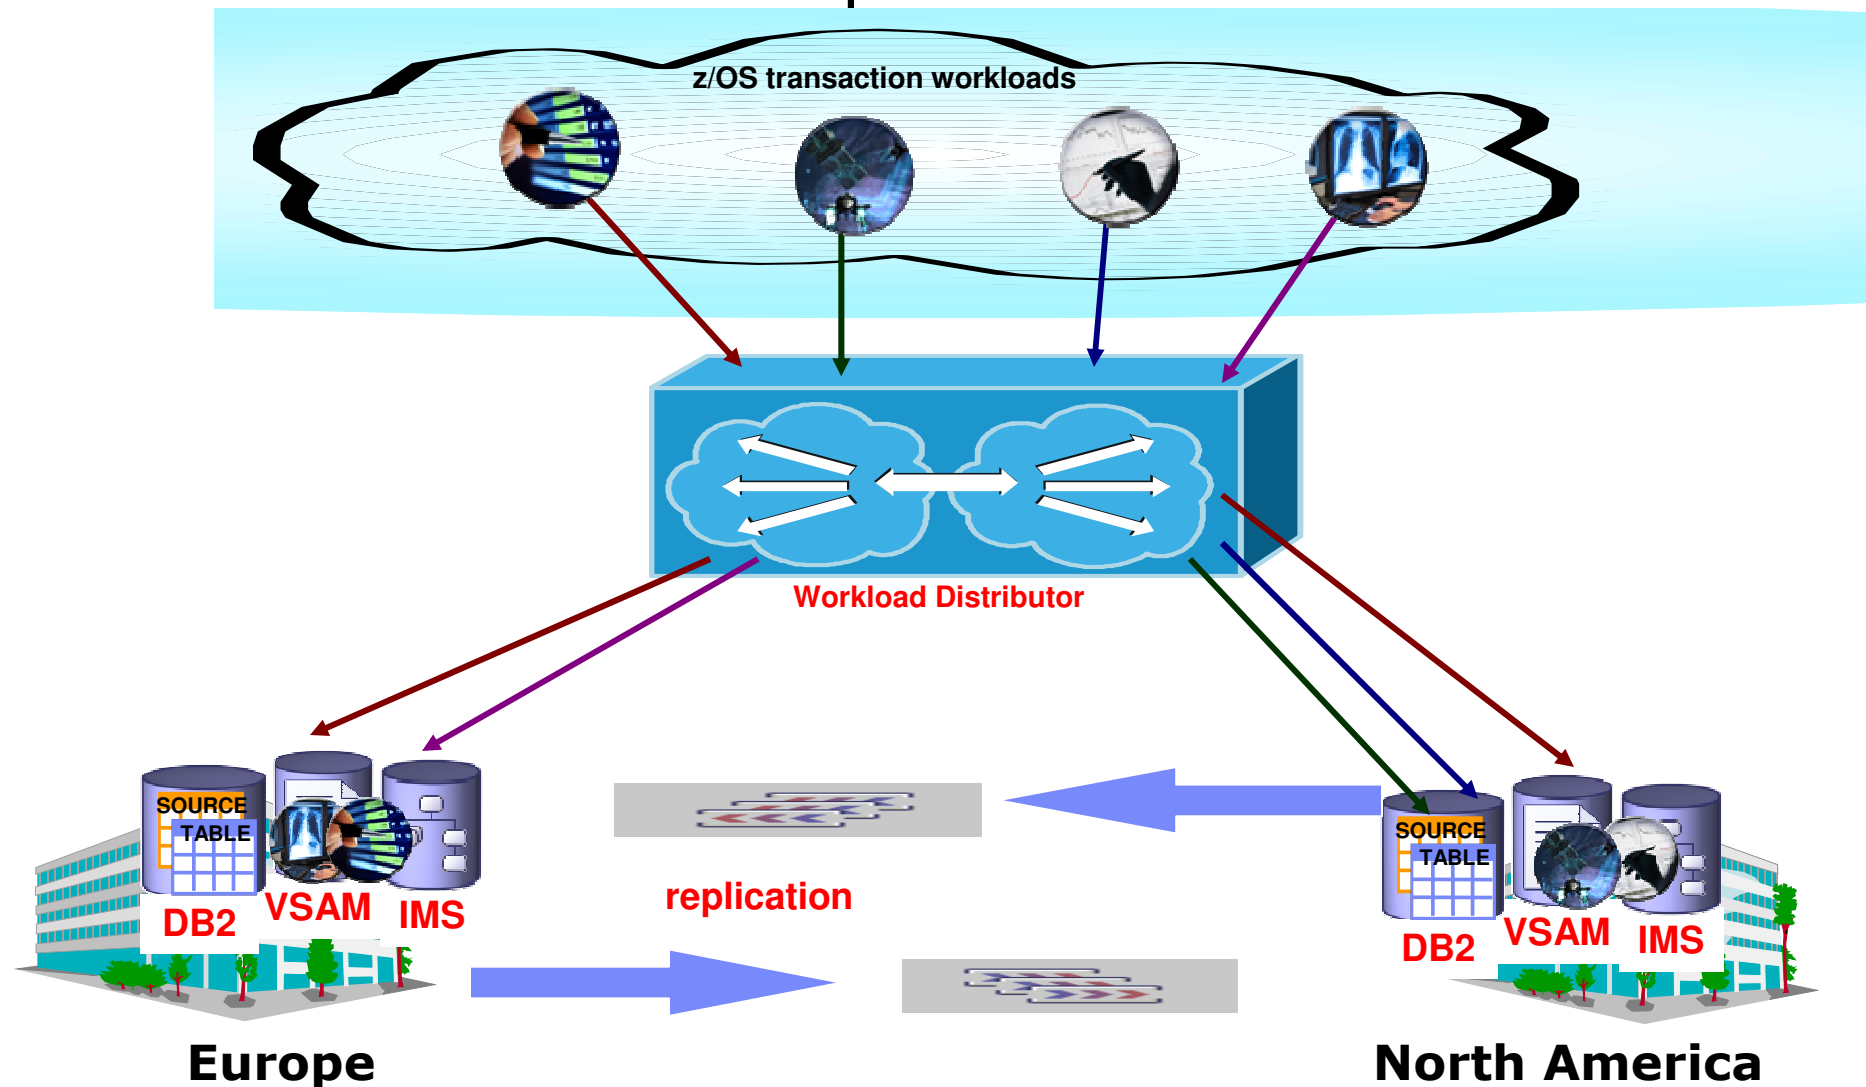
Active/Active Sites Overview

Components

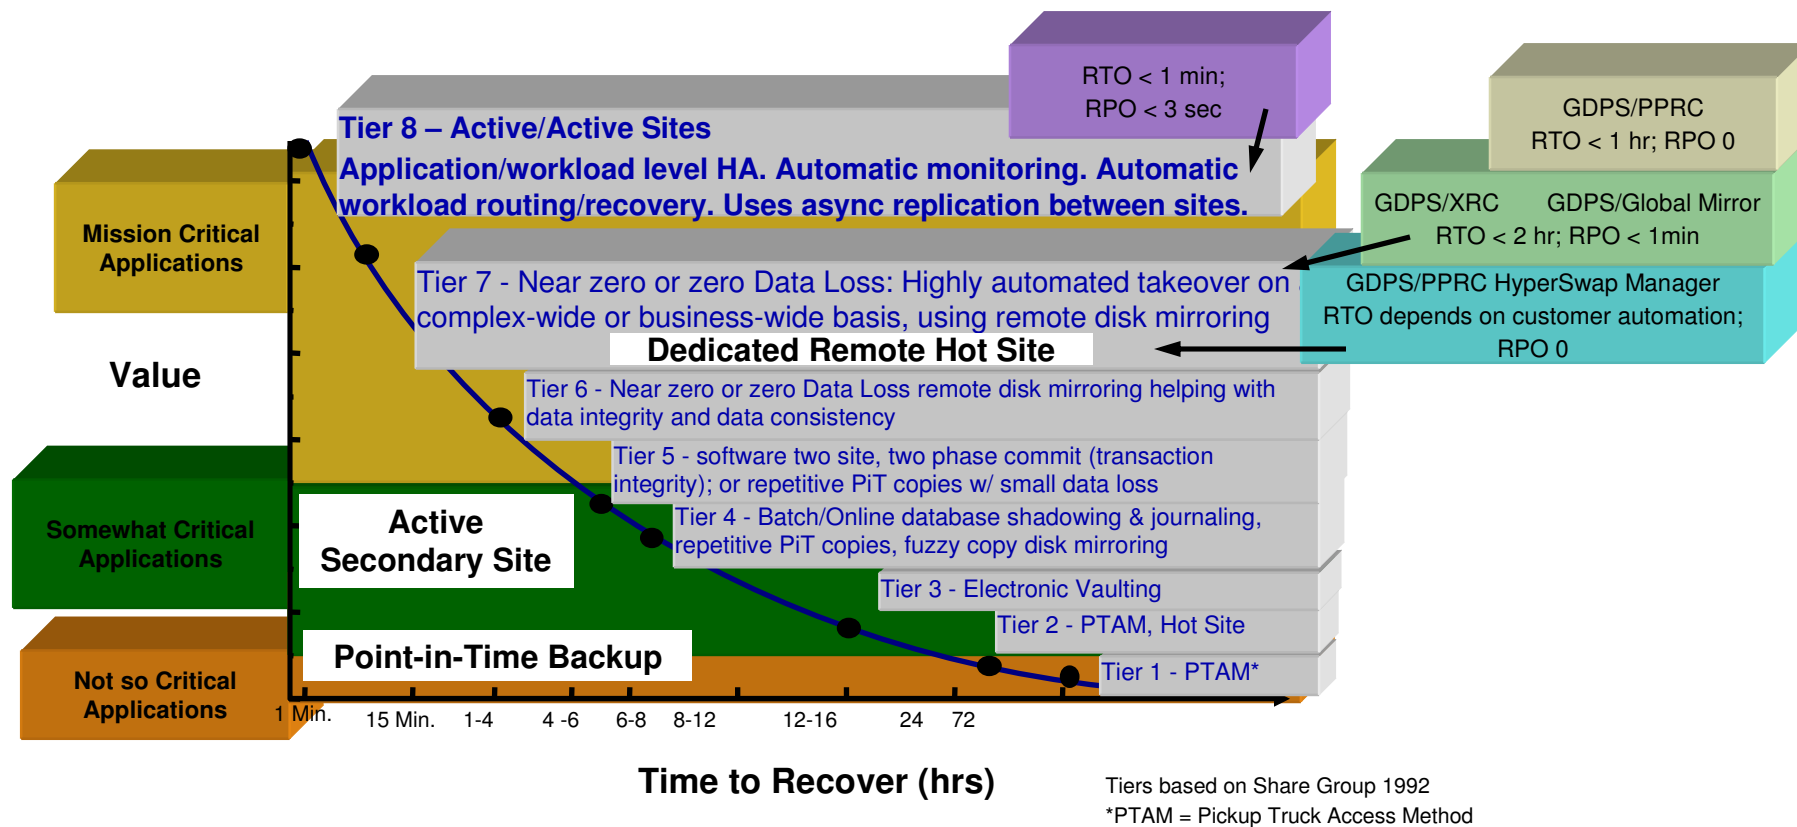
Roadmap



Active/Active Sites concept



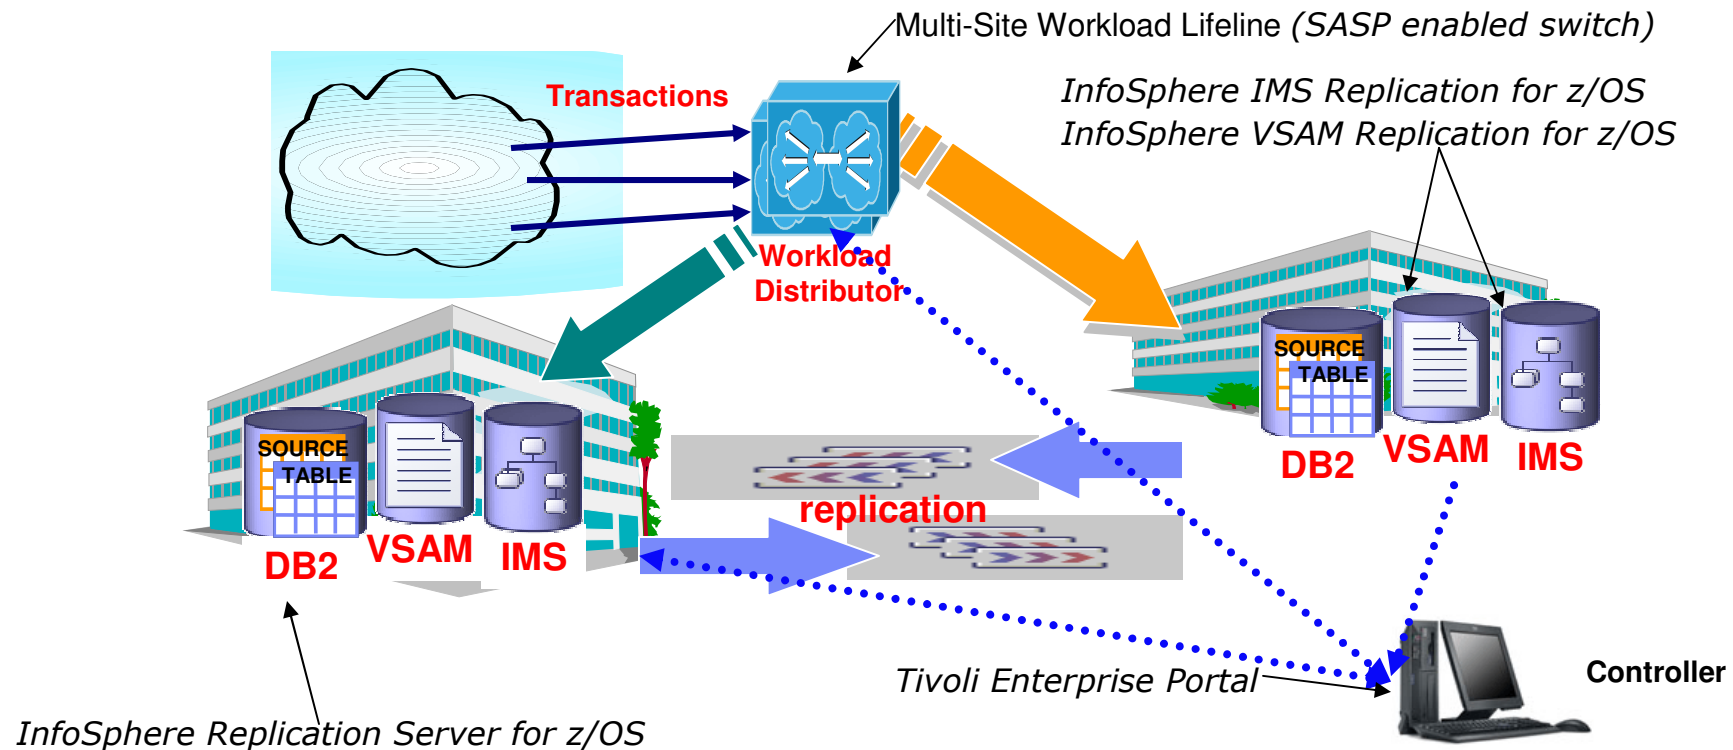
Tiers of Disaster Recovery: Level Setting Active/Active Sites



Best D/R practice is blend tiers of solutions in order to maximize application coverage at lowest possible cost . One size, one technology, or one methodology does not fit all applications

Active/Active Sites – What is it ?

- Two or more sites, separated by *unlimited* distances, running the same applications and having the same data to provide cross-site workload balancing and Continuous Availability / Disaster Recovery
- Paradigm shift: failover model => near continuous availability model



Active/Active Sites Configurations

- **Configurations**
 1. Active/Standby
 2. Active/Query
 3. Active/Active w/data partitioning
 4. Active/Active

- **A configuration is specified on a workload basis**

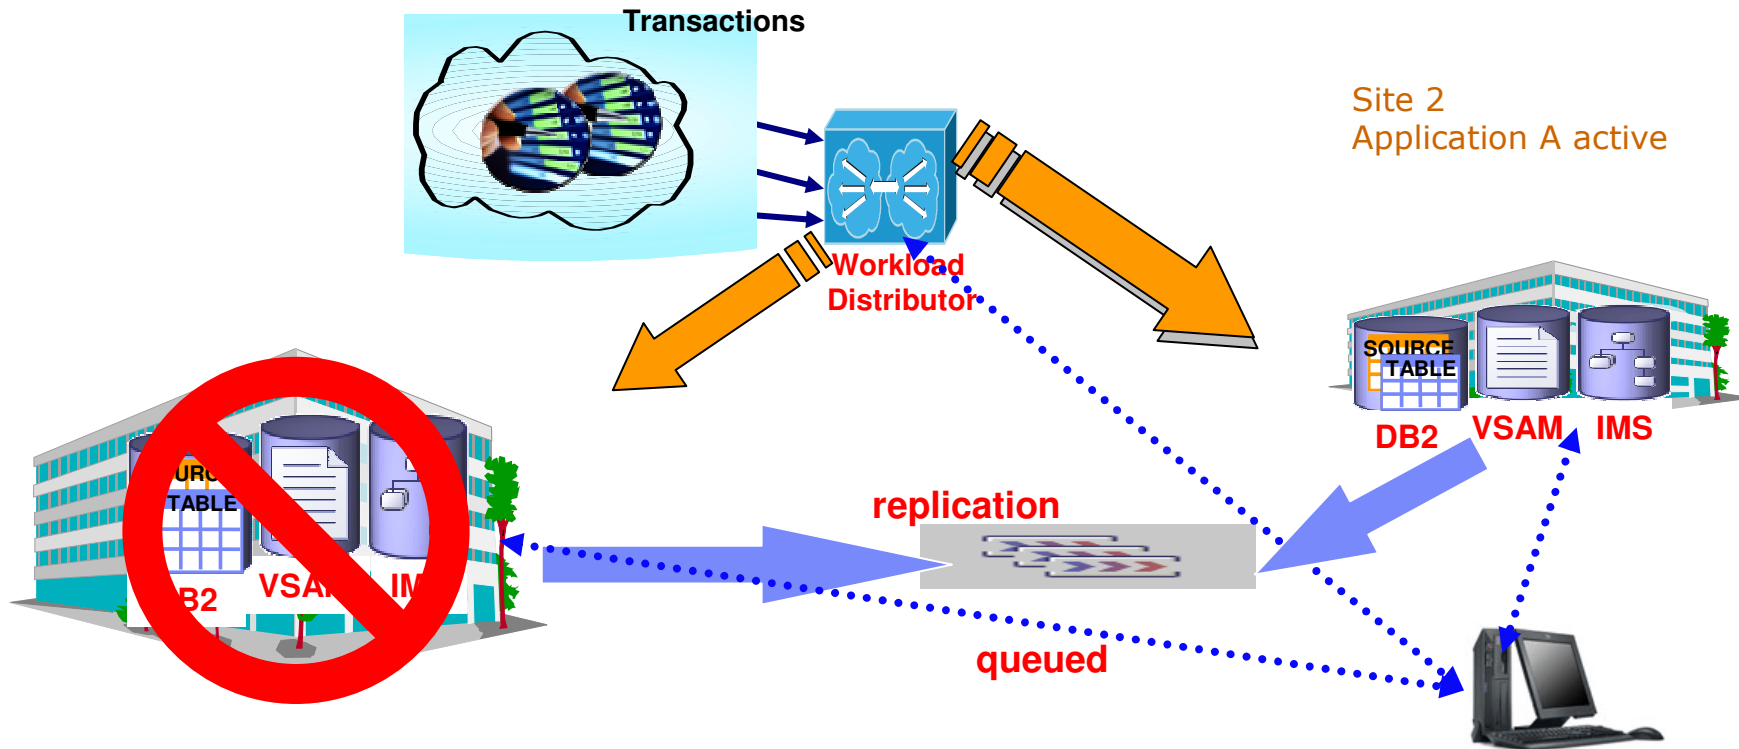
- **A workload is the aggregation of these components**
 - **Software:** user written applications (e.g., COBOL program) and the middleware run time environment (e.g., CICS region & DB2 subsystem)
 - **Data:** related set of objects that must preserve transactional consistency and optionally referential integrity constraints (e.g., DB2 Tables)
 - **Network connectivity:** one or more TCP/IP addresses & ports (e.g., 10.10.10.1:80)

1. Active/Standby Configuration

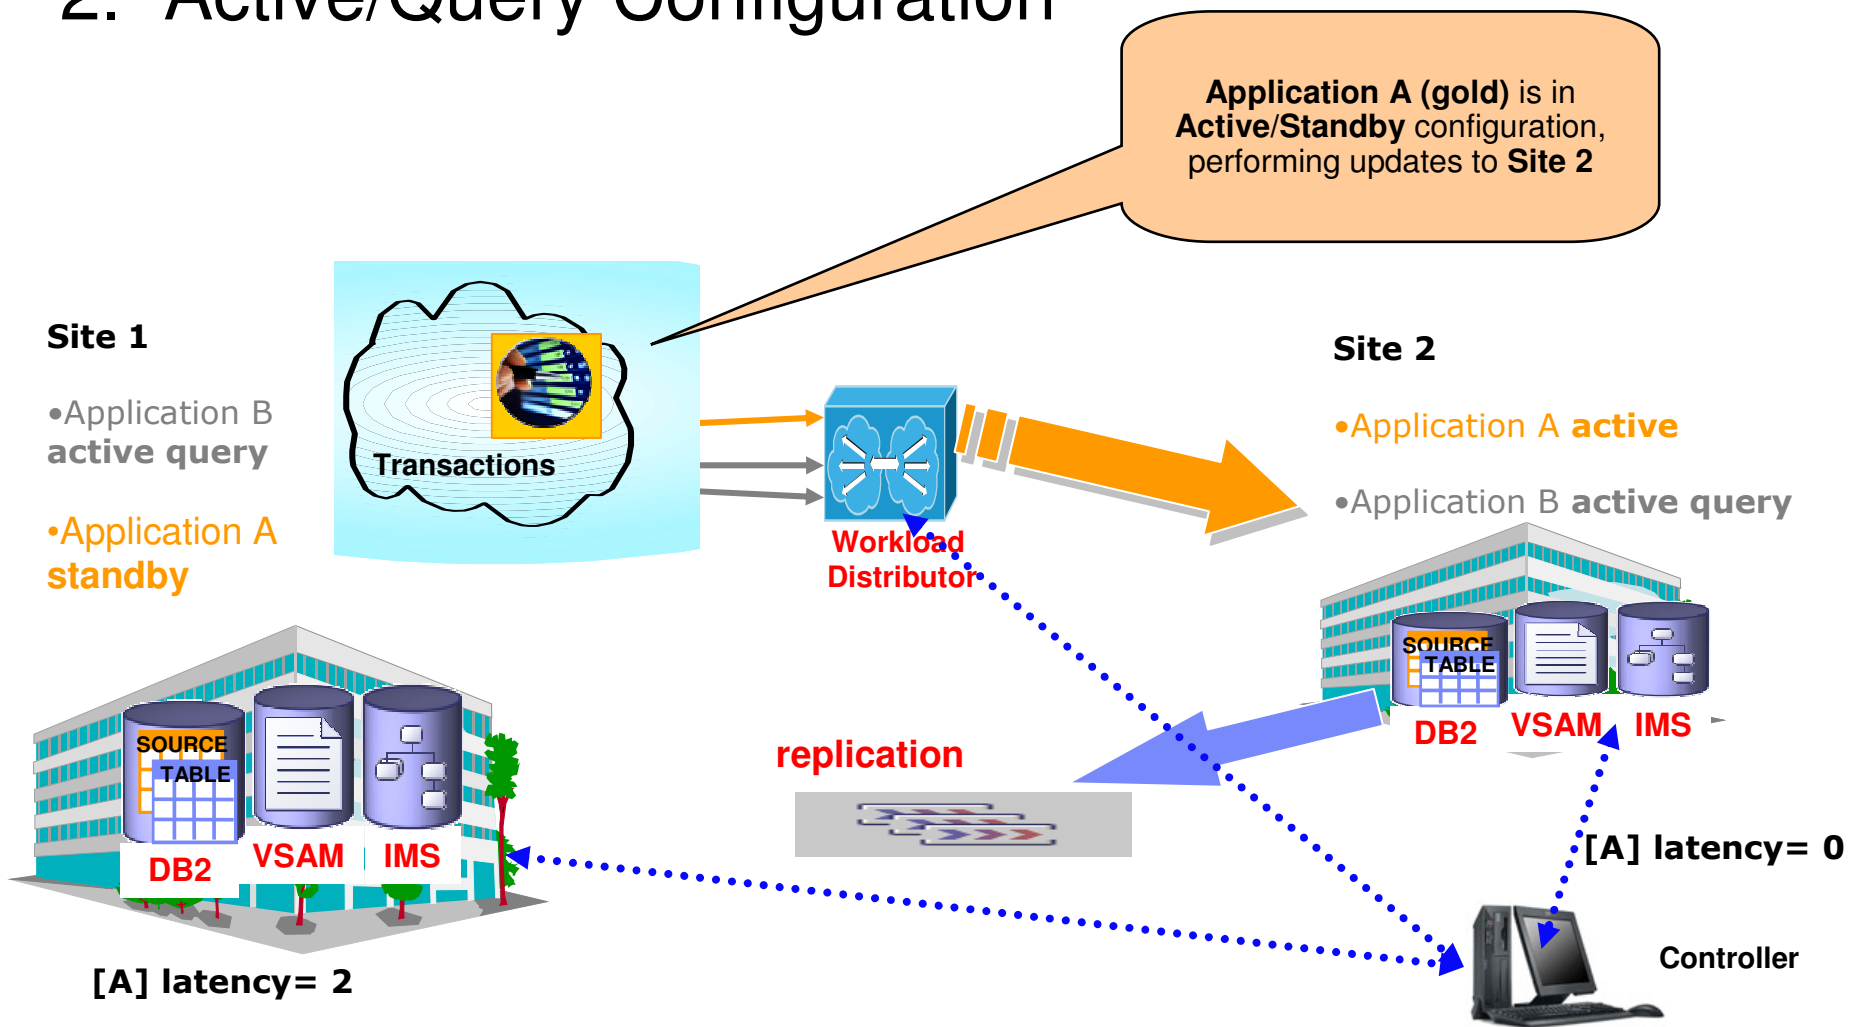
- Static routing
- Automatic failover

Site 1
Application A active

Site 2
Application A standby



2. Active/Query Configuration



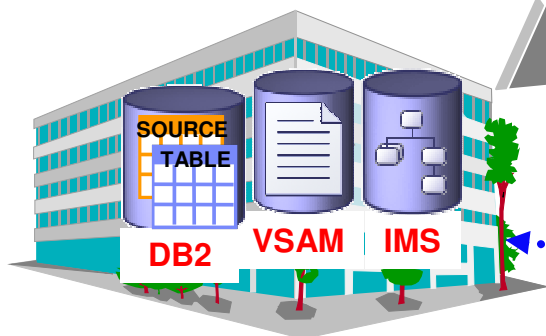
2. Active/Query Configuration

Application B (grey) is in **Active/Query** config

- Using same data as Application A
- Active to both **Site 1** and **Site 2**, but **favor Site 1**
- Routed according to replication latency policy
- Policy for query routing: *Max Latency 5, Reset Latency 3*

Site 1


- Application B **active query**
- Application A **standby**



Application A latency = 2

Site 2

- Application A **active**
- Application B **active query**

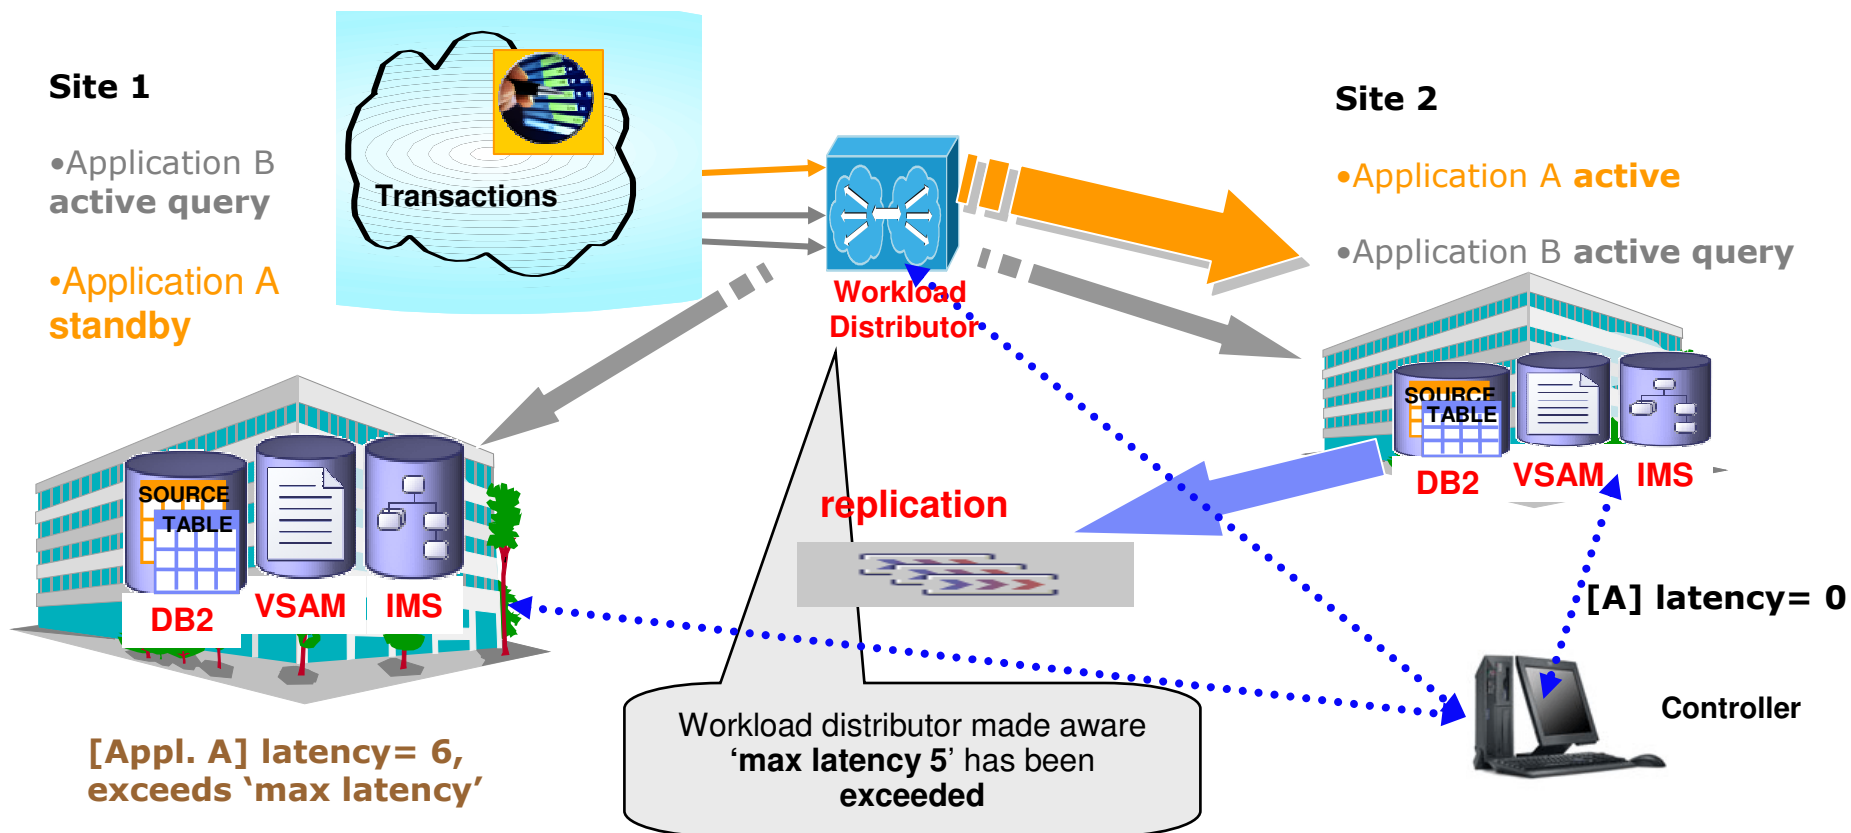


[A] latency = 0

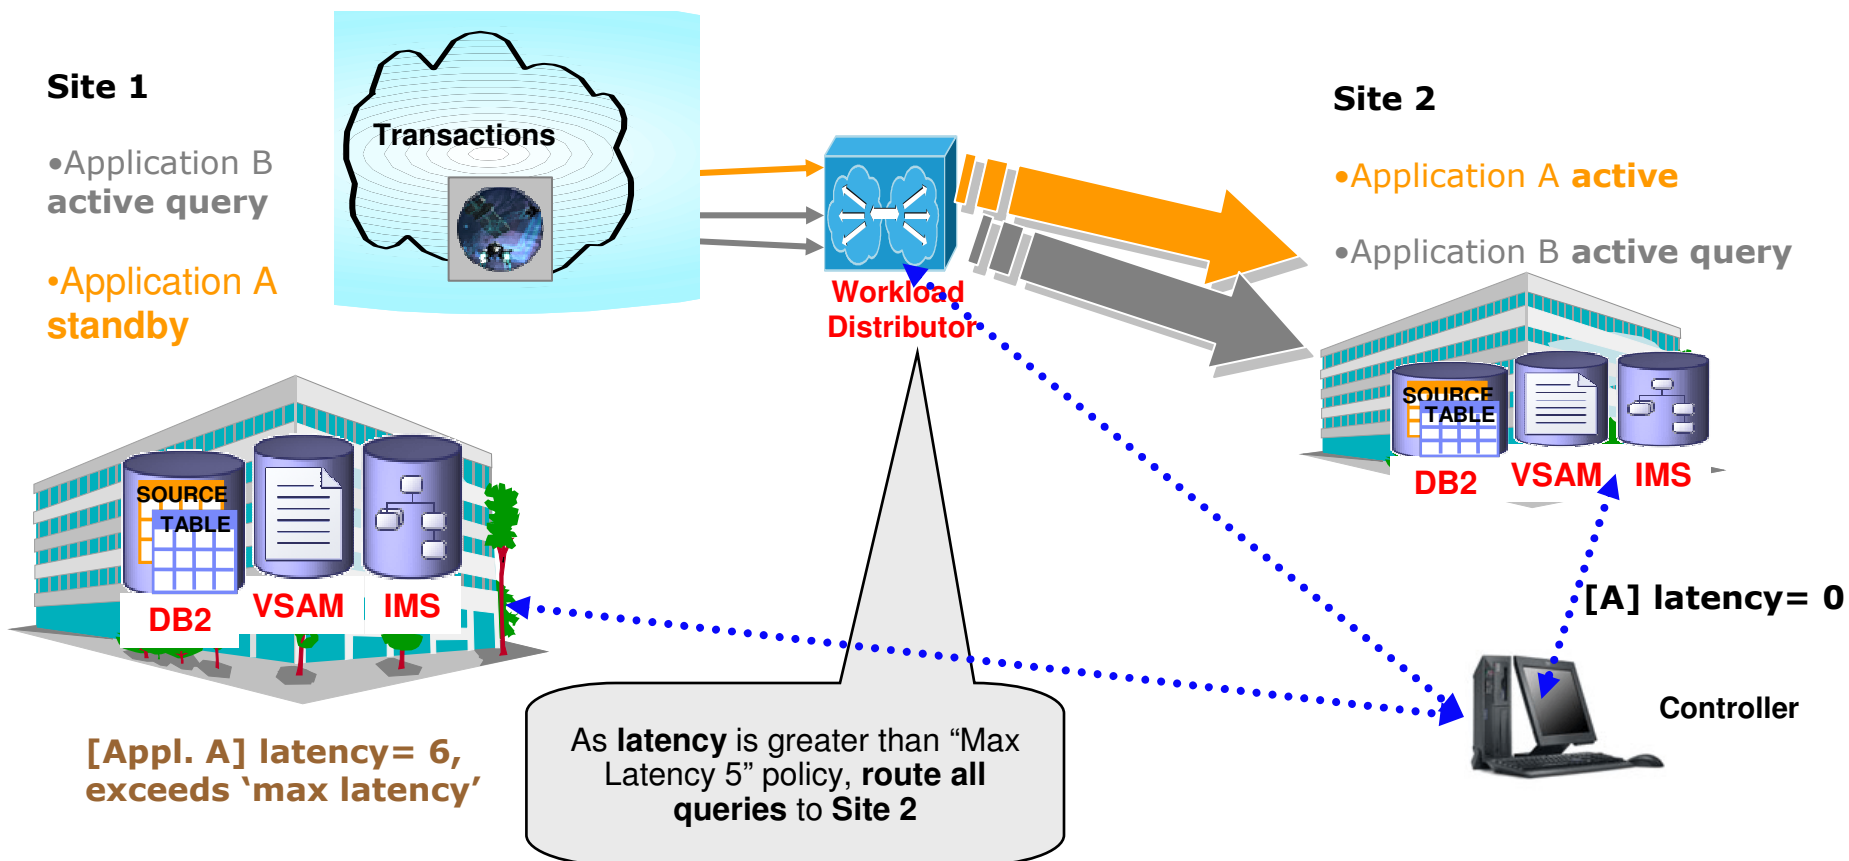
Controller

As **latency** is less than "Max Latency 5", follow policy to **skew queries to site 1**

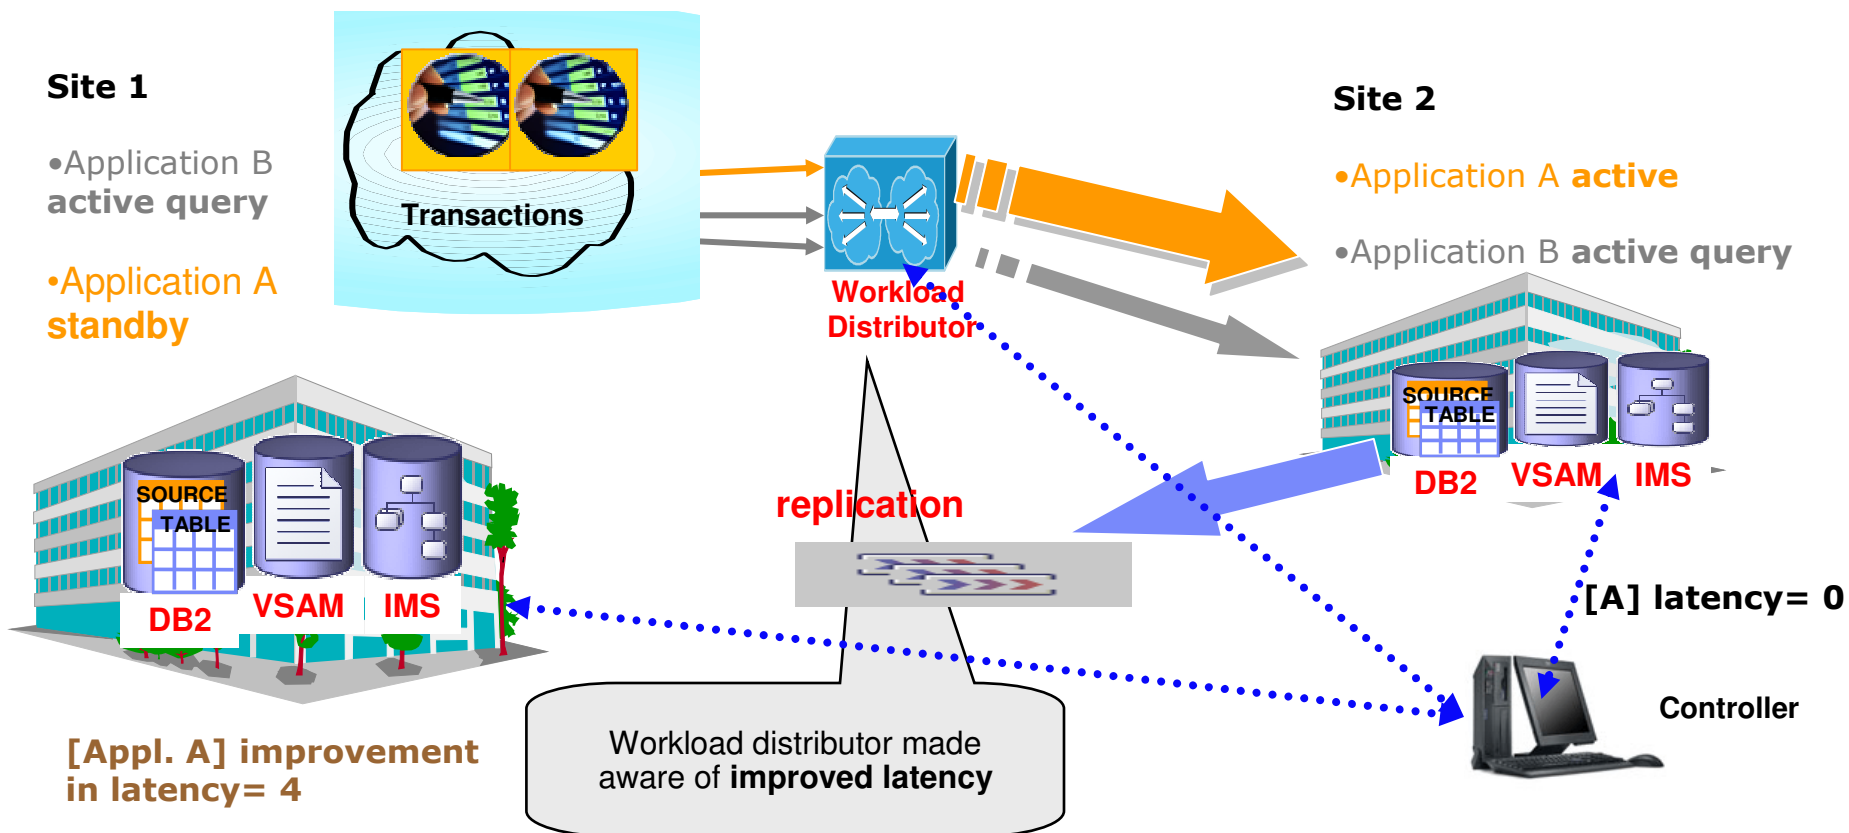
2. Active/Query Configuration



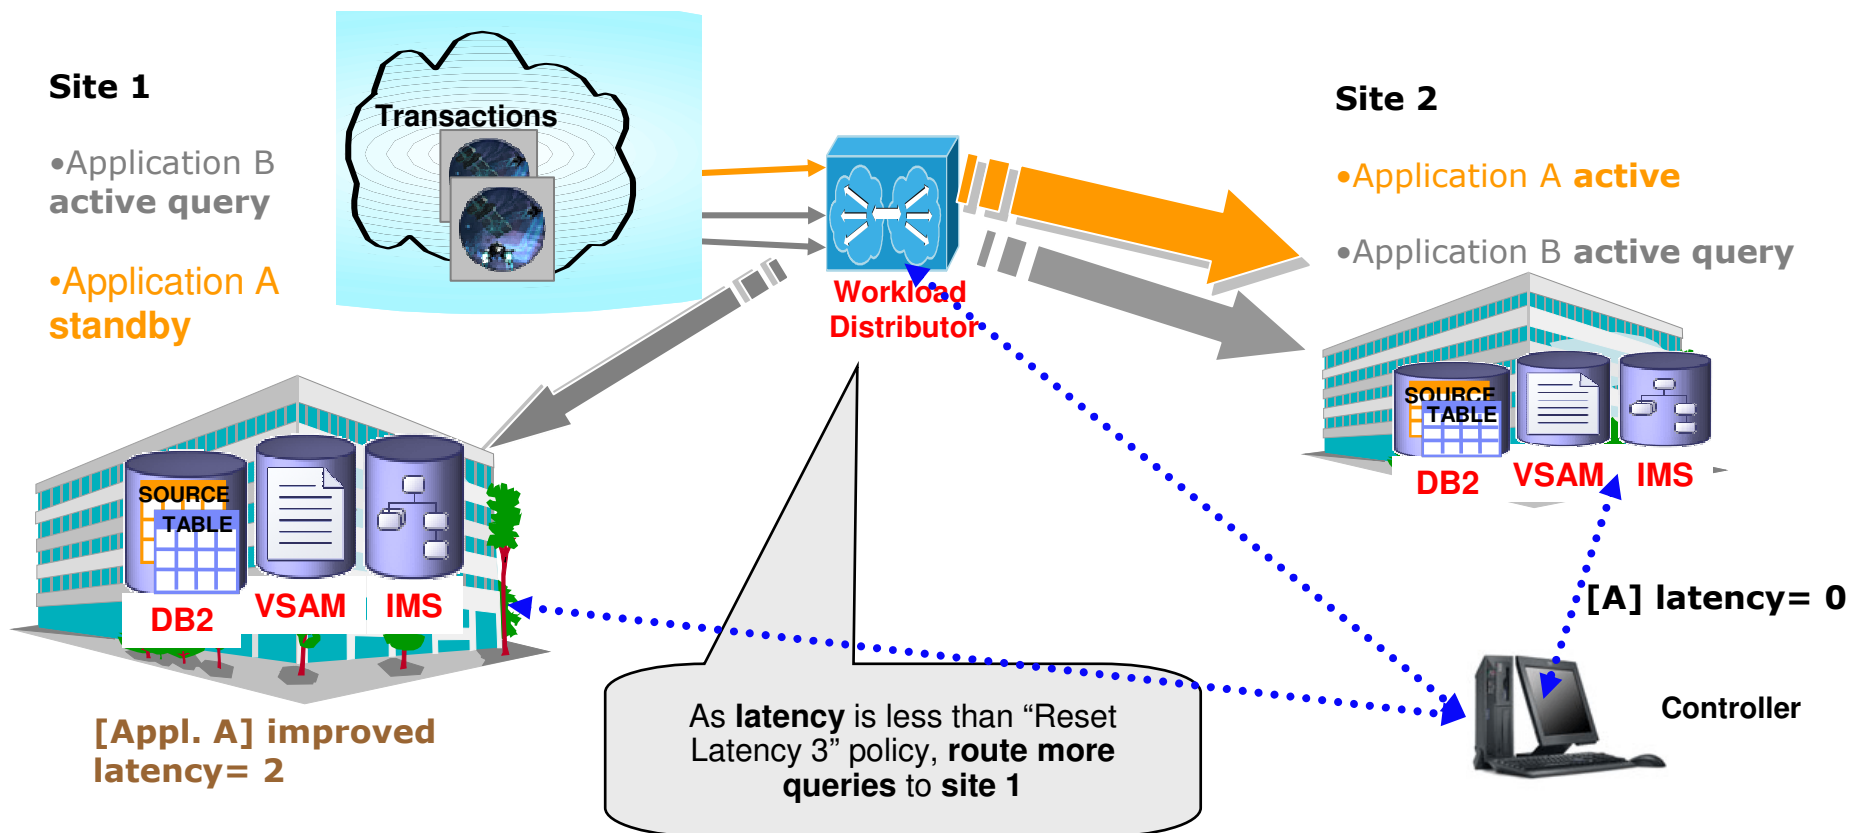
2. Active/Query Configuration



2. Active/Query Configuration



2. Active/Query Configuration



3. Active/Active w/Data Partitioning Configuration

- Context based routing based upon rules
- Automatic failover

Site 1

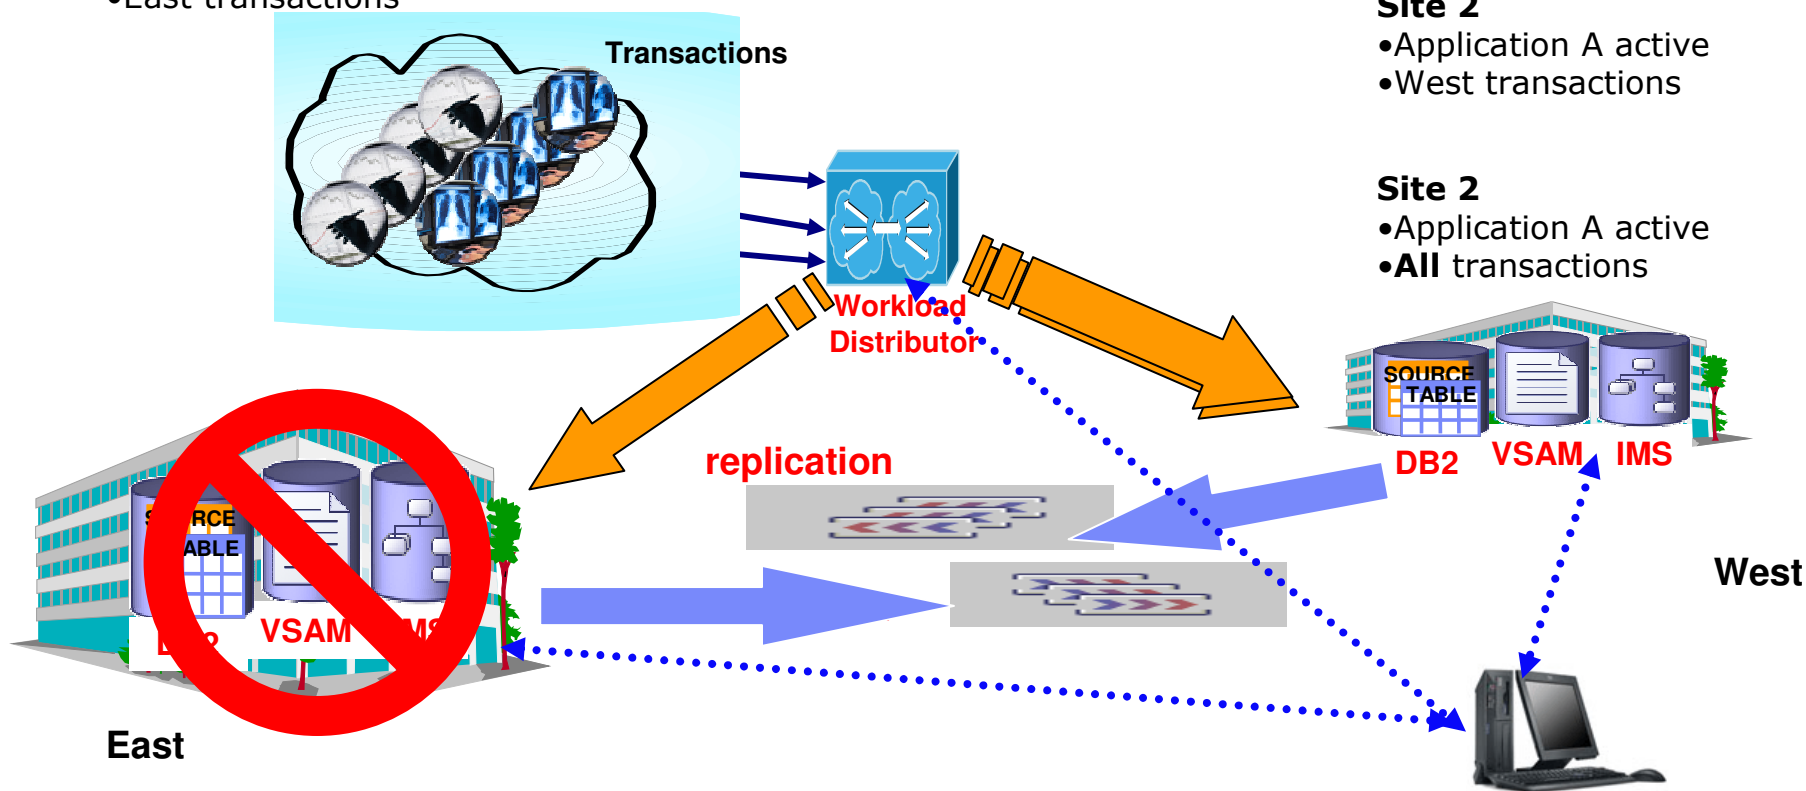
- Application A active
- East transactions

Site 2

- Application A active
- West transactions

Site 2

- Application A active
- All transactions



4. Active/Active Configuration

- Context based routing based upon SLA achievement
- Transactions with affinity routed to same location to avoid conflicts.
- Automatic failover

Site 1

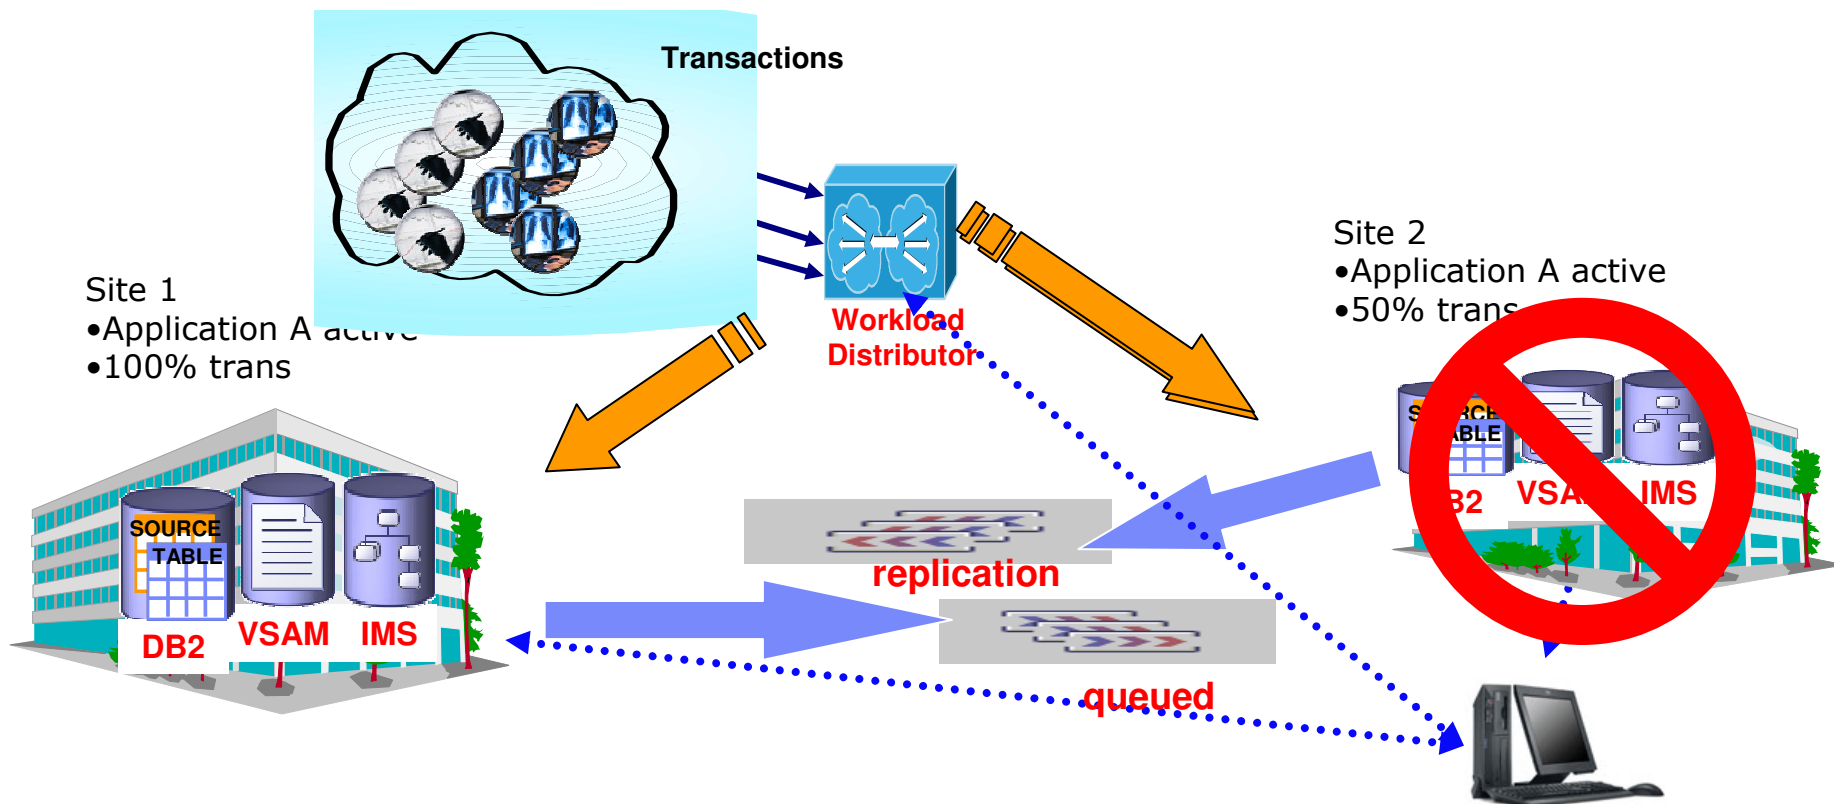
- Application A active
- 50% trans

Site 1

- Application A active
- 100% trans

Site 2

- Application A active
- 50% trans



Active/Active Sites overview



Level Set

Active/Active Sites Overview



Components

Roadmap



Assumptions

- **Version/Release numbers noted on these slides are the required minimum supported by Active/Active as of 7/30/2010**

Required/Supported products installed on z/OS production images

▪ Operating systems required/supported

- z/OS V1R11 (including Tivoli Directory Server for z/OS)

▪ Applications/MiddleWare

- DB2 for z/OS v9
- IMS v11
- WS MQ v7.0

▪ Replication

- InfoSphere Replication Server (DB2) Next
- InfoSphere IMS Replication for z/OS Next

▪ Management and monitoring

- NetView for z/OS Next
- System Automation for z/OS Next
- IBM Multi-site Workload Lifeline v1.1 (new product, announced and GA'd w/ Solution)
- IBM Tivoli Monitoring 6.2.2
- OMEGAMON products (required only if the customer wants to monitor the behavior of the respective products/resources that they deal with (DB2, CICS, storage, etc.)
 - OMEGAMON XE on z/OS v4.2.0
 - OMEGAMON XE for Mainframe Networks v4.2.0
 - OMEGAMON XE for Storage v4.2.0
 - OMEGAMON XE for DB2 Performance Expert (or Performance Monitor) on z/OS v4.2.0 (if DB2 is running)
 - OMEGAMON XE on CICS for z/OS v4.2.0 (if CICS is running)
 - OMEGAMON XE on IMS v4.2.0 (if IMS is running)
 - OMEGAMON XE for Messaging v7.0 (if MQ is running)

Required/supported products installed on z/OS controller images

- **Operating systems required/supported**
 - z/OS V1R11 (including Tivoli Directory Server for z/OS)

- **Management and monitoring**
 - GDPS/Active-Active 1.1 (new product, announced and Ga'd with Solution)
 - NetView for z/OS Next
 - System Automation for z/OS Next
 - IBM Multi-site Workload Lifeline v1.1 (new product, announced and GA'd with Solution)
 - IBM Tivoli Monitoring 6.2.2
 - OMEGAMON products (required only if the customer wants to monitor the behavior of the respective products/resources that they deal with (DB2, CICS, storage, etc.))
 - OMEGAMON XE on z/OS v4.2.0
 - OMEGAMON XE for Mainframe Networks v4.2.0
 - OMEGAMON XE for Storage v4.2.0

Active/Active Sites overview



Level Set

Active/Active Sites Overview


Components



Roadmap

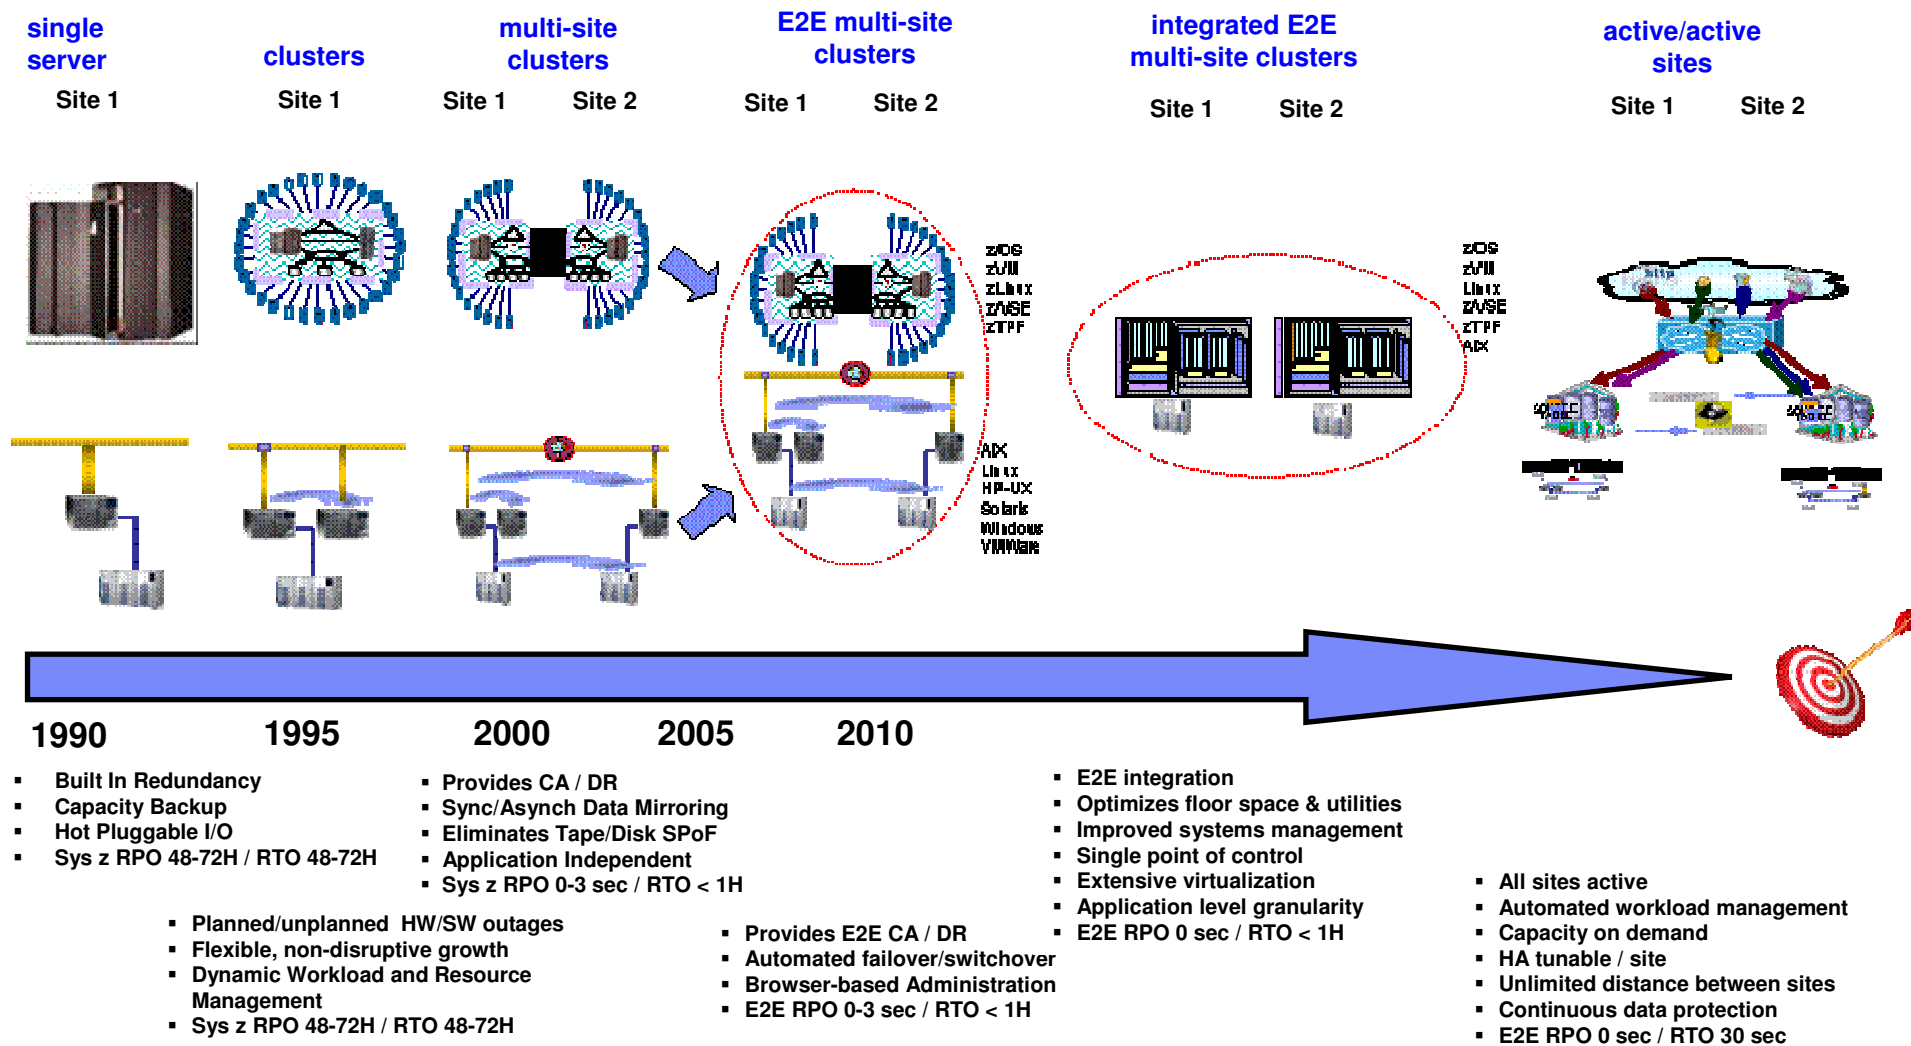


Active/Active Sites Roadmap



- **R1 – Mainframe servers**
 - 1a – **active/standby** – **Beta YE10**
 - Configuration – active/standby
 - Data sources – DB2 & IMS
 - Two sites
 - 1b – **active/query**
 - Configuration – active/standby & active/query
 - Data sources – DB2, IMS, & VSAM
 - Centralized policy repository
 - 1c – **data partitioning**
 - Configuration – active/standby, active/query, & data partitioning
 - Data sources – DB2, IMS, & VSAM
 - 1d – **active/active**
 - Configuration – active/standby, active/query, data partitioning, & active/active
 - Data sources – DB2, IMS, & VSAM
- **R2 – Mainframe & distributed servers**

Enterprise Wide BC Solution Direction





THANK
YOU