IBM Software Group – Enterprise Networking Solutions

Enabling Continuous Availability and Reducing Downtime with IBM Multi-site Workload Lifeline

SHARE Session 15197
March 12, 2014

Lin Overby – overbylh@us.ibm.com

Mark Figley – figley@us.ibm.com

Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchase decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Agenda

- What is Multi-site Workload Lifeline?
- Providing Continuous Availability as Part of GDPS Active-Active Sites Solution
- Reducing Downtime for Planned Outages
- Multi-site Workload Lifeline Commands

What is Multi-site Workload Lifeline?

Multi-site Workload Lifeline

Enabling very high availability with load balancing and workload rerouting

- **The IBM® Multi-site Workload Lifeline product**
  - Enables intelligent load balancing of TCP/IP workloads across two sites at unlimited distances for near continuous availability.
  - Facilitates planned outages by rerouting workloads from one site to another without disruption to users.

Why is Multi-site Workload Lifeline Important?

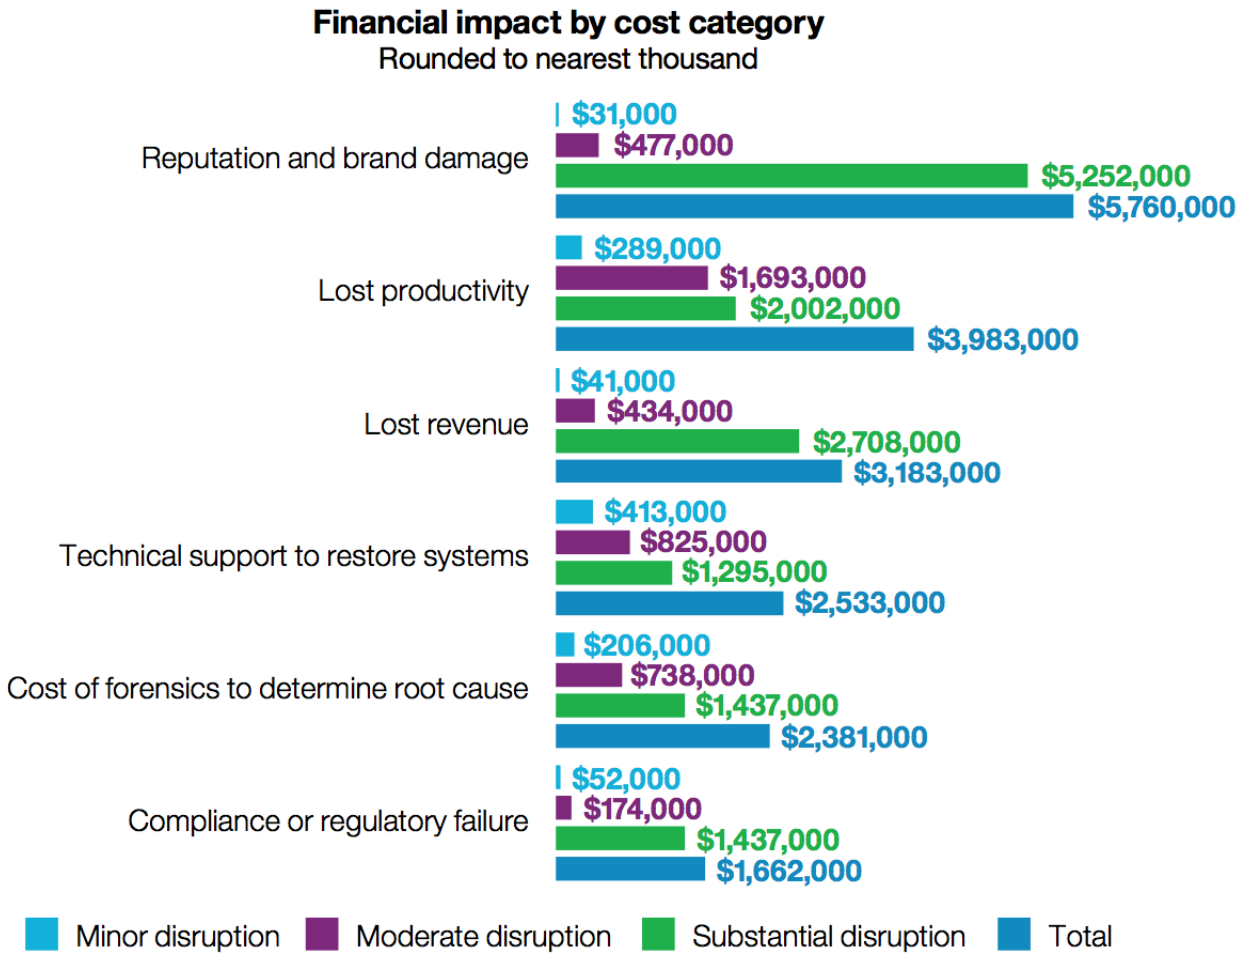

Why is Multi-site Workload Lifeline Important?

Corporate viability impact:

According to Coleman Parkes Research:

- 48% of companies believe downtime negatively impacts their brand and reputation.

- Over 34% of companies would be unable to meet compliance and regulatory commitments due to downtime.

- 44% of respondents believe IT downtime can damage staff morale and 35% report it can harm customer loyalty.

Why is Multi-site Workload Lifeline Relevant to Me?

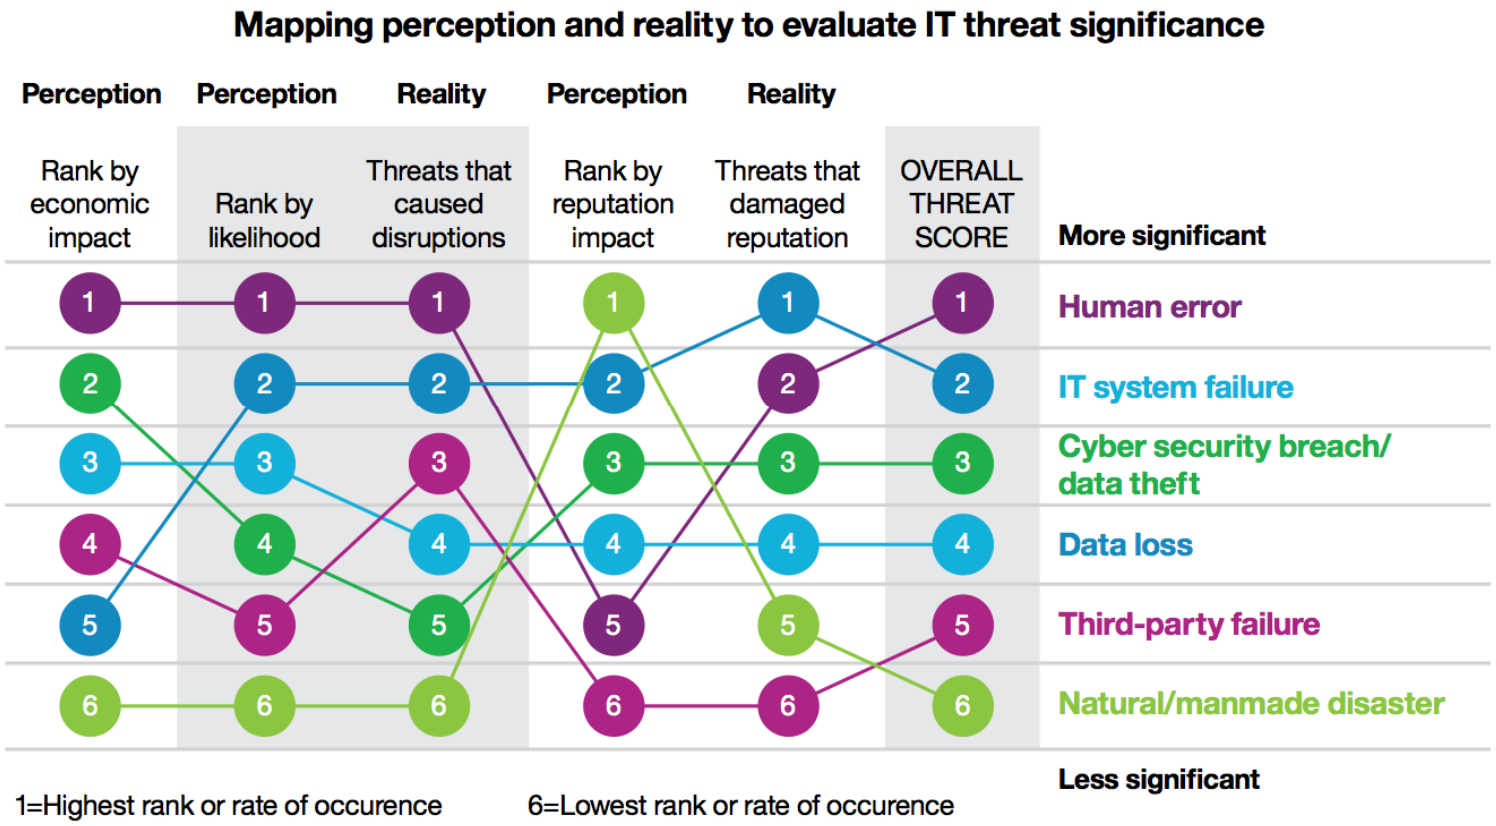

Why is Multi-site Workload Lifeline Relevant to Me?

- If a site outage would mean hours of downtime for your critical workloads while your secondary site starts up…
- If you are tired of your system being down for hours every time you need to do a Sysplex wide upgrade…
- If you are frustrated that your secondary site sits there cold instead of being up, active and useful…

… then Lifeline is relevant to you.

What does Multi-site Workload Lifeline do?

- Lifeline helps enable **near zero downtime** for critical workloads during unplanned workload or site level outages
- Lifeline can help reduce planned outage downtimes for critical workloads **from hours to minutes**
- For select workloads and usage scenarios, Lifeline can enable connections for a workload to be serviced on two sites at the same time with no geographic limitations.

What makes Multi-site Workload Lifeline different?

- **Lifeline is not an all-or-nothing solution.** Lifeline allows CIOs to pursue a gradual, incremental approach to HA/DR that focuses on the most critical workloads first.

- Lifeline supports routing read-only queries and workloads to your secondary site, reducing the strain on your primary transaction system and allowing organizations to **get more value from their secondary site investment**.

What makes Multi-site Workload Lifeline different?

- Lifeline doesn't use disk based replication or require the backup site to be "cold", allowing disparate workloads to be run on each site with failover to the other site
- Lifeline has no geographical distance limitations between sites
- Lifeline requires no configuration changes to:
  - Applications
  - Clients
  - Servers
  - The network topology

What capabilities does Multi-site Workload Lifeline enhance?

- **Increased performance:** Response time is reduced by ensuring new connections for a workload are distributed to the applications and systems most capable of handling them
- **Increased availability:** New connections for a workload can be routed to available applications even in the event of application, system, or site outages
- **Increased scalability:** Application instances can be added on demand
- **Analytic capability:** Network Management Interface (NMI) provides access to workload, application, and site status information
- **Improved recovery time:** Reduction of Recovery Time Objective from hours to minutes
- **Workload migration:** Ability to move workloads from one site to the other with minimal disruption
- **Increased availability:** Outages for maintenance updates or other planned events can be minimized
- **Verification of disaster recovery procedures:** Simpler, non-disruptive testing of disaster recovery procedures by validating workloads remain accessible on the recovery site without requiring a site outage on the production site

Providing continuous availability as part of GDPS Active-Active Sites solution

Background: Business Continuity Definitions

- Disaster recovery (DR)
  - The reconstruction of your physical production site in an alternate physical site, occurring after the loss of your primary data center.
  - The process of bringing up servers and applications, in priority order, to support the business from the alternate site.
- Continuous Availability (CA)
  - Application cannot undergo an *unplanned* outage for more than a few seconds/minutes at a time, but can do so as often as necessary, or may be down for a few hours for *scheduled* maintenance.
- Recovery Time Objective (RTO)
  - Duration of time and a service level within which a business process must be restored after a disaster (or disruption) in order to avoid unacceptable consequences associated with a break in business continuity.
- Recovery Point Objective (RPO)
  - Maximum tolerable period in which data might be lost from an IT service due to a major incident

Background: Data Replication

- What is data replication?
  - A solution for copying data between databases, typically residing in different sites
  - Emphasizes the copying of only changed data
    - An application makes updates to a database and these changes are captured locally and applied to a remote database
  - Replication scope
    - An entire database
    - A subset of the database (subset of tables or subset of columns or rows within a table)

- Data replication enables:
  - Continuous (High) Availability
    - Failover to replicated database after workload outage
  - Offload query workloads to replicated database
    - Read-only database provides near-real time reporting

Background: Data Replication Modes

- Synchronous data replication
  - Waits for replication to complete before a transaction is complete
  - Feasible for local to metro distances without significant transaction delay

- Asynchronous data replication
  - Transaction can complete before replication
  - Typically used for global distances where added latency of synchronous replication makes transaction time intolerable
  - Some data loss possible during unplanned failover

DR/CA solutions available prior to GDPS/Active-Active

Metro distance DR/CA solution

- GDPS/PPRC (Metro Mirror), based upon a multi-site Parallel Sysplex and synchronous disk replication, is a metro area Continuous Availability (CA), Disaster Recovery (DR) solution
  - Workloads can withstand site and/or storage failures
- Low recovery time and zero data loss
- **Issue: This GDPS product does not provide enough site separation for some enterprises**

Global distance DR solution

- GDPS/XRC and GDPS/GM, based upon separate Sysplexes and asynchronous disk replication, are unlimited distance Disaster Recovery (DR) solutions
  - Disaster recovery for metro-region interruptions
  - Longer recovery time with "seconds" of data loss
- These GDPS products require the failed site's workload to be restarted in the recovery site and this typically will take an hour or longer
- **Issue: These GDPS products will not achieve a recovery time of seconds being requested by some enterprises**

The GDPS Active-Active Sites Solution ….
…. With Multi-Site-Workload Lifeline

Metro/Global Distance DR/CA Solution

- GDPS Active-Active Sites enables 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
- Access data from any site (unlimited distance between sites)
- Provide workload distribution between sites
- Provide application level granularity
- Asynchronous data replication
- Low recovery time and near zero data loss
- Paradigm shift: failover model => near continuous availability model
  - For critical workloads requiring continuous availability
  - Not a replacement for disaster recovery of non-critical workloads

GDPS Active-Active Sites – What is it?

- GDPS Active-Active Sites is a bundle of products that include:
  - –GDPS scripting code
  - –NetView for system automation and monitoring
  - –Multi-site Workload Lifeline provides workload distribution between sites
  - –IBM InfoSphere Data Replication for DB2 database replication
  - –InfoSphere IMS Replication for z/OS
  - –InfoSphere VSAM Replication for z/OS
  - –IBM Tivoli Monitoring
- Similar to other GDPS solutions, GTS services to set up are available
- Requires a SASP-capable external load balancer

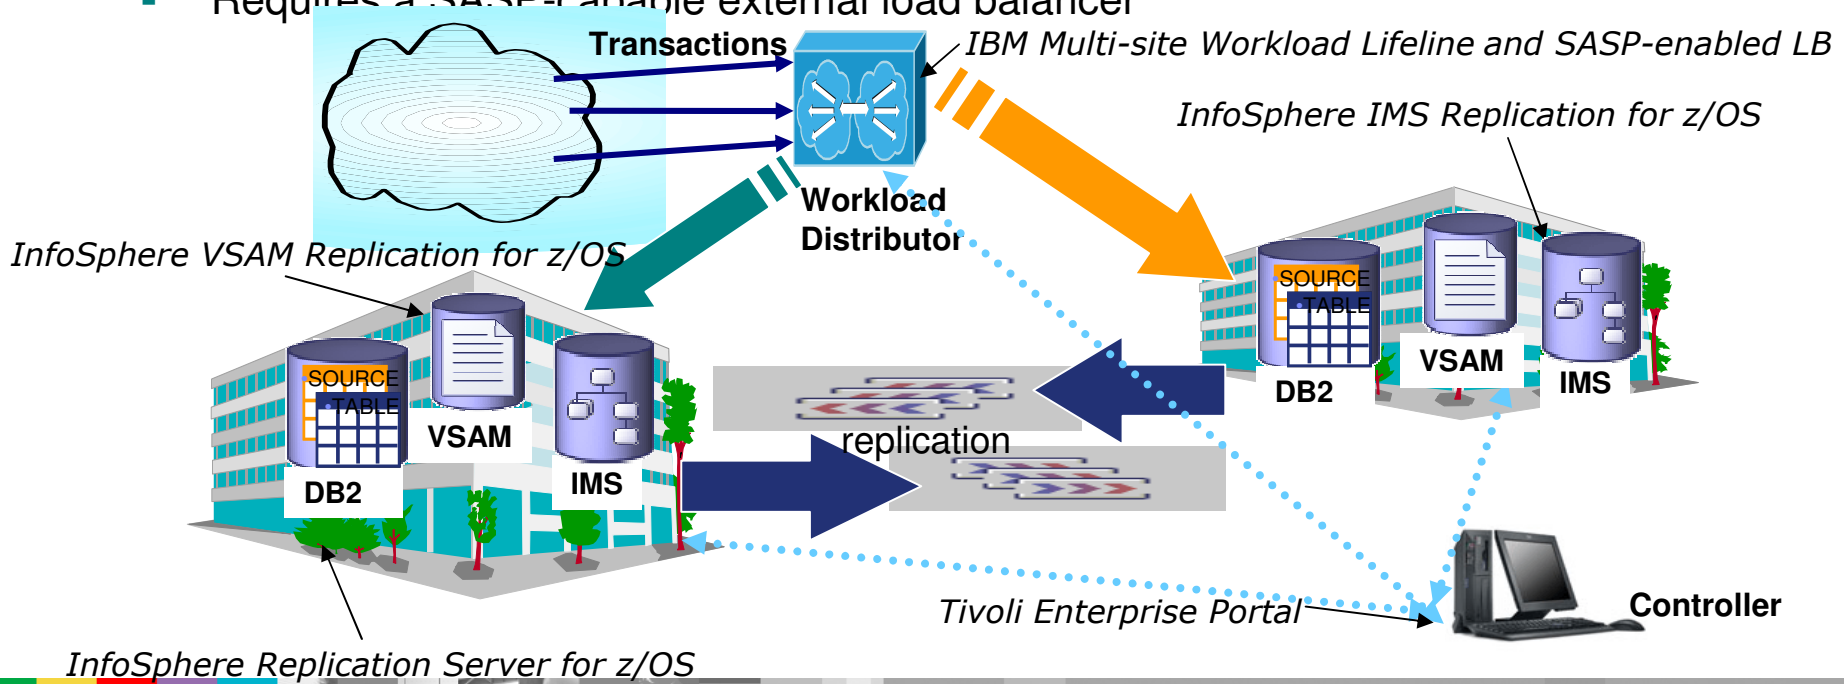

GDPS Active-Active Sites load balancing requirements

- Ability to distribute workloads between sites (and route around failed sites)
  - Based on capacity/health of sites and server application instances within a site
- Ability to detect workload or site failures
- Ability to switch workloads from one site to another site
  - Perform “graceful” failback following a workload or site disaster
- Ability to maintain workload configuration states in event of a workload manager failure
  - Keep a peer workload manager in sync with workload states
- Ability to dynamically add/modify workloads
- Ability to surface routing recommendations to network management agents

GDPS Active-Active Sites Configurations

- Configurations
  - Active/Standby (MSWL V1R1)
  - Active/Query (MSWL V2R0)

- Configuration is specified on a workload basis

- A workload is the aggregation of these components
  - ***Software:*** 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 (e.g., DB2 Tables)
  - ***Network connectivity*****:** one or more TCP/IP addresses & ports (e.g., 10.10.10.1:80)

Scenario 1: Active/Standby Configuration – Prior to workload outage

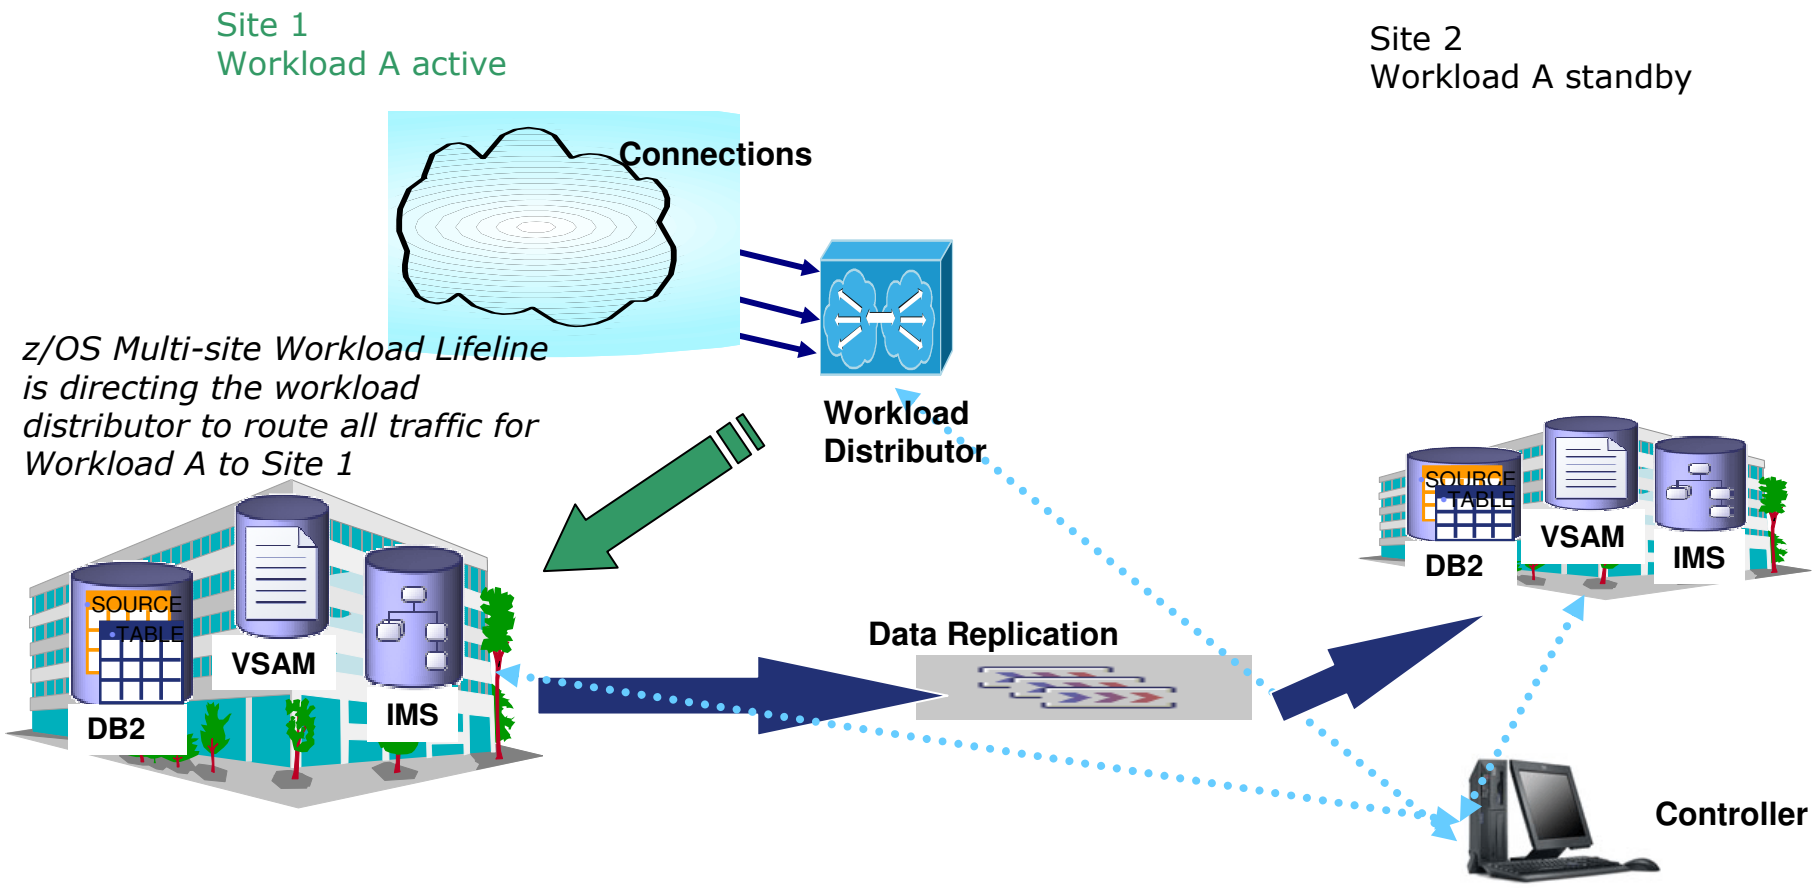

Scenario 1: Active/Standby Configuration – Workload outage / Site available

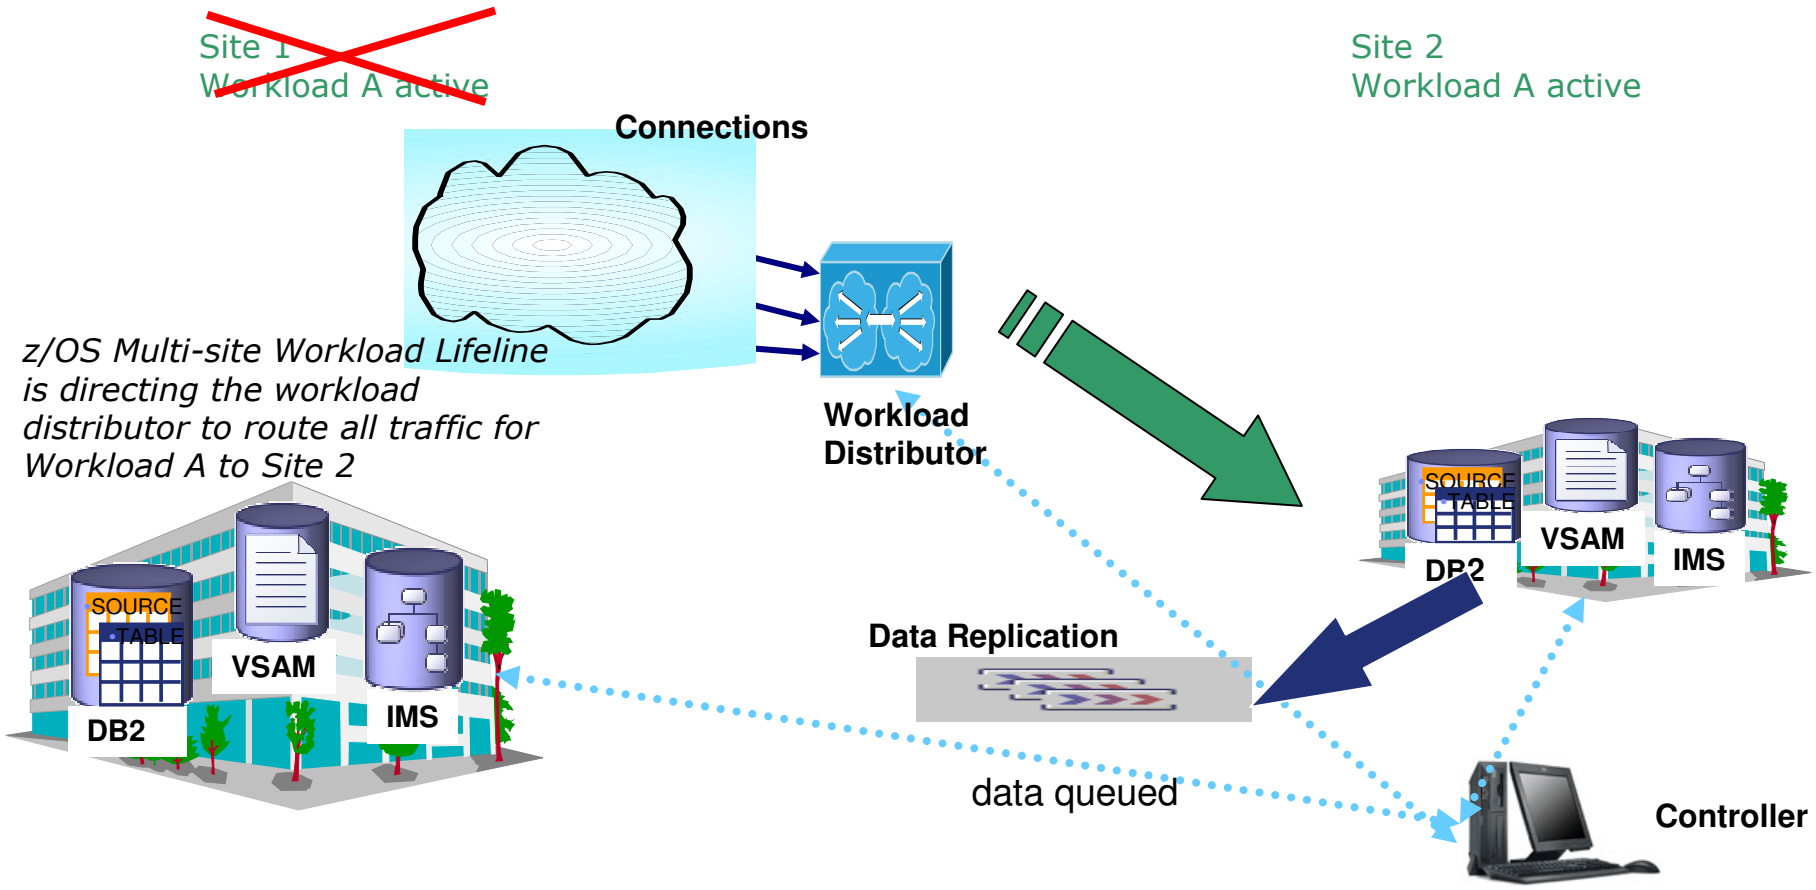

Scenario 1: Active/Standby Configuration – Workload outage / Site outage

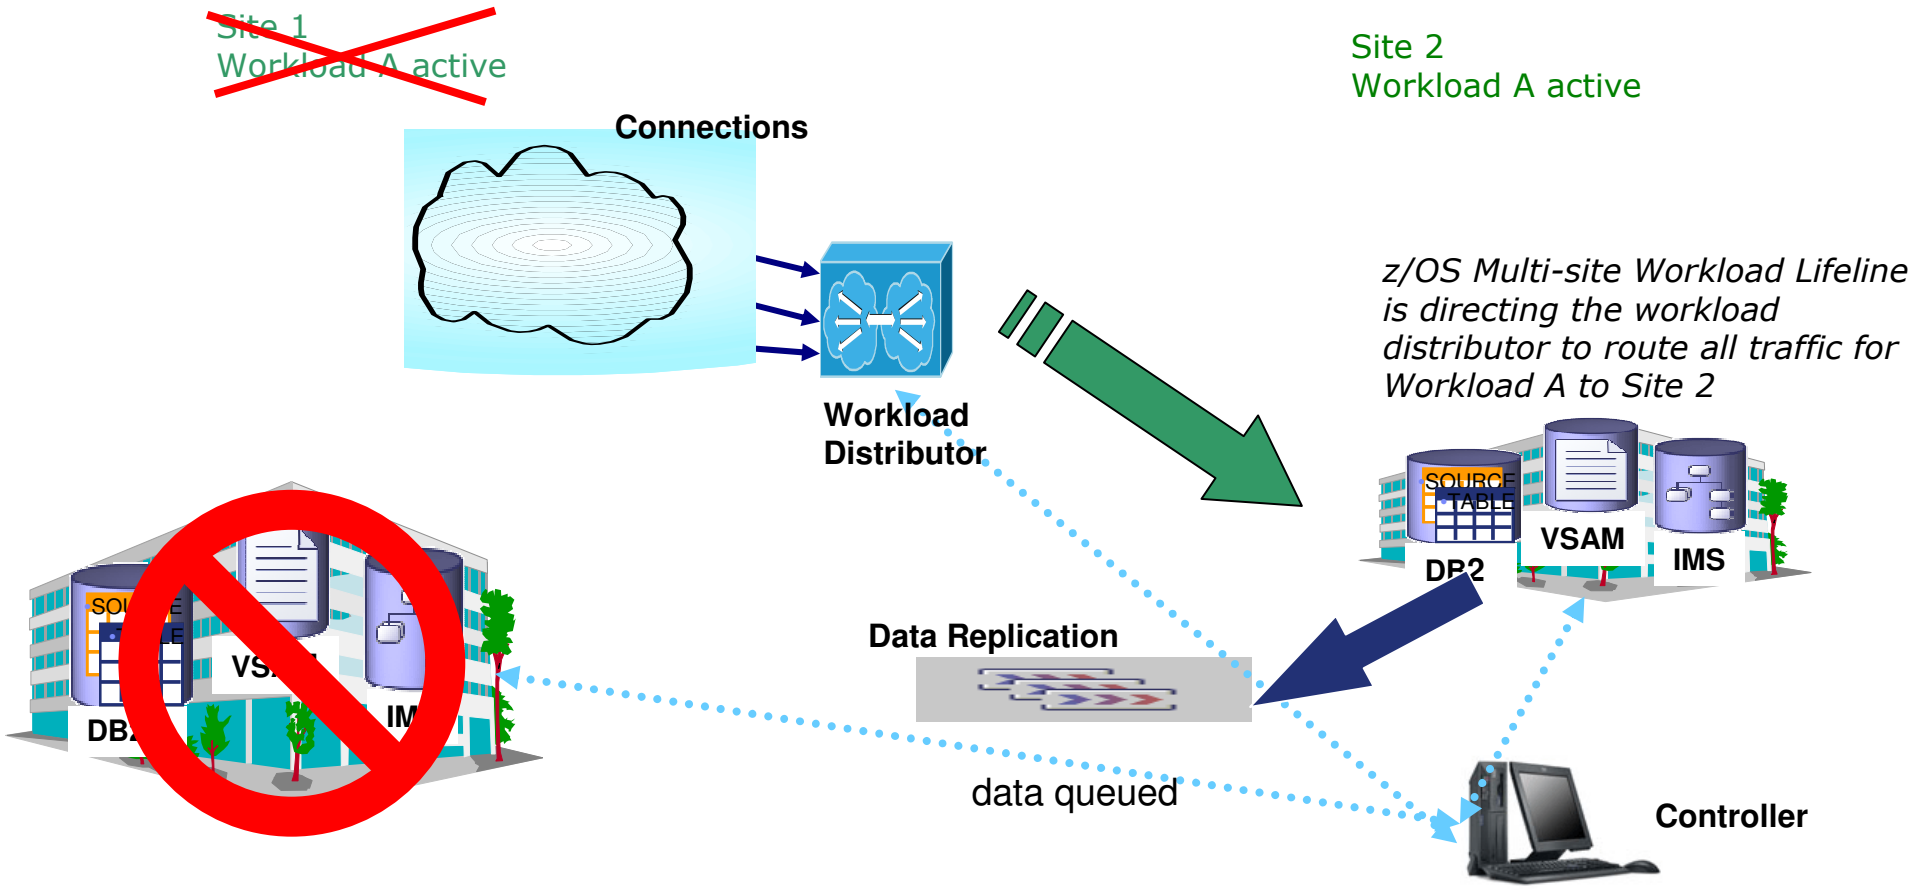

Scenario 2: Active/Standby Configuration - (multiple workloads – mutual continuous availability)

Site 1
Workload A active

Site 2
Workload A standby

Site 1
Workload B Standby

Site 2
Workload B active

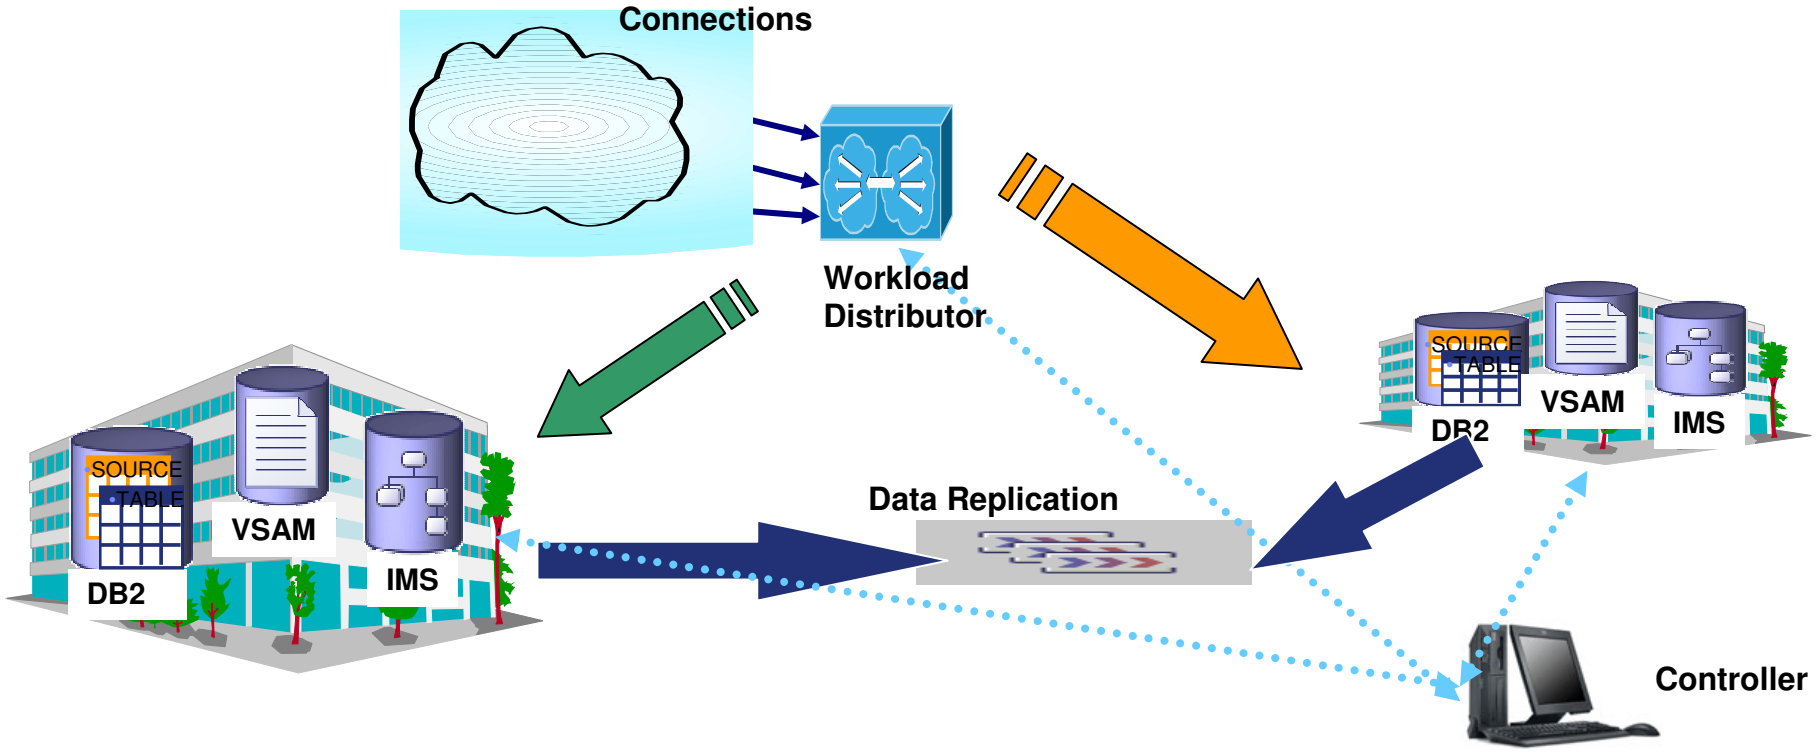

Active/Query Workloads

- An Active/Standby workload is active on only one site
  - Workload transactions update data on the active site
  - Database changes are replicated to the standby site
- An associated Active/Query workload can be active on both sites
  - Workload transactions access same data being updated by Active/Standby workload
  - Workload transactions only query data
- Active/Query workload connections are distributed to a site based on routing type and average replication latency
  - Dynamic workload routing – distribution between sites based on availability and health of server applications within each site
  - Static workload routing – distribution between sites based on a configured percentage
  - Replication latency – Average delay between when update transactions for a workload to the active site are replicated and applied to the standby site

Scenario 3: Active/Query Configuration – Distribution of query workload across two sites

Site 1
Update Workload A active
Replication Latency: 'Normal'

Site 1
Query Workload B active

Site 2
Update Workload A standby

Site 2
Query Workload B active

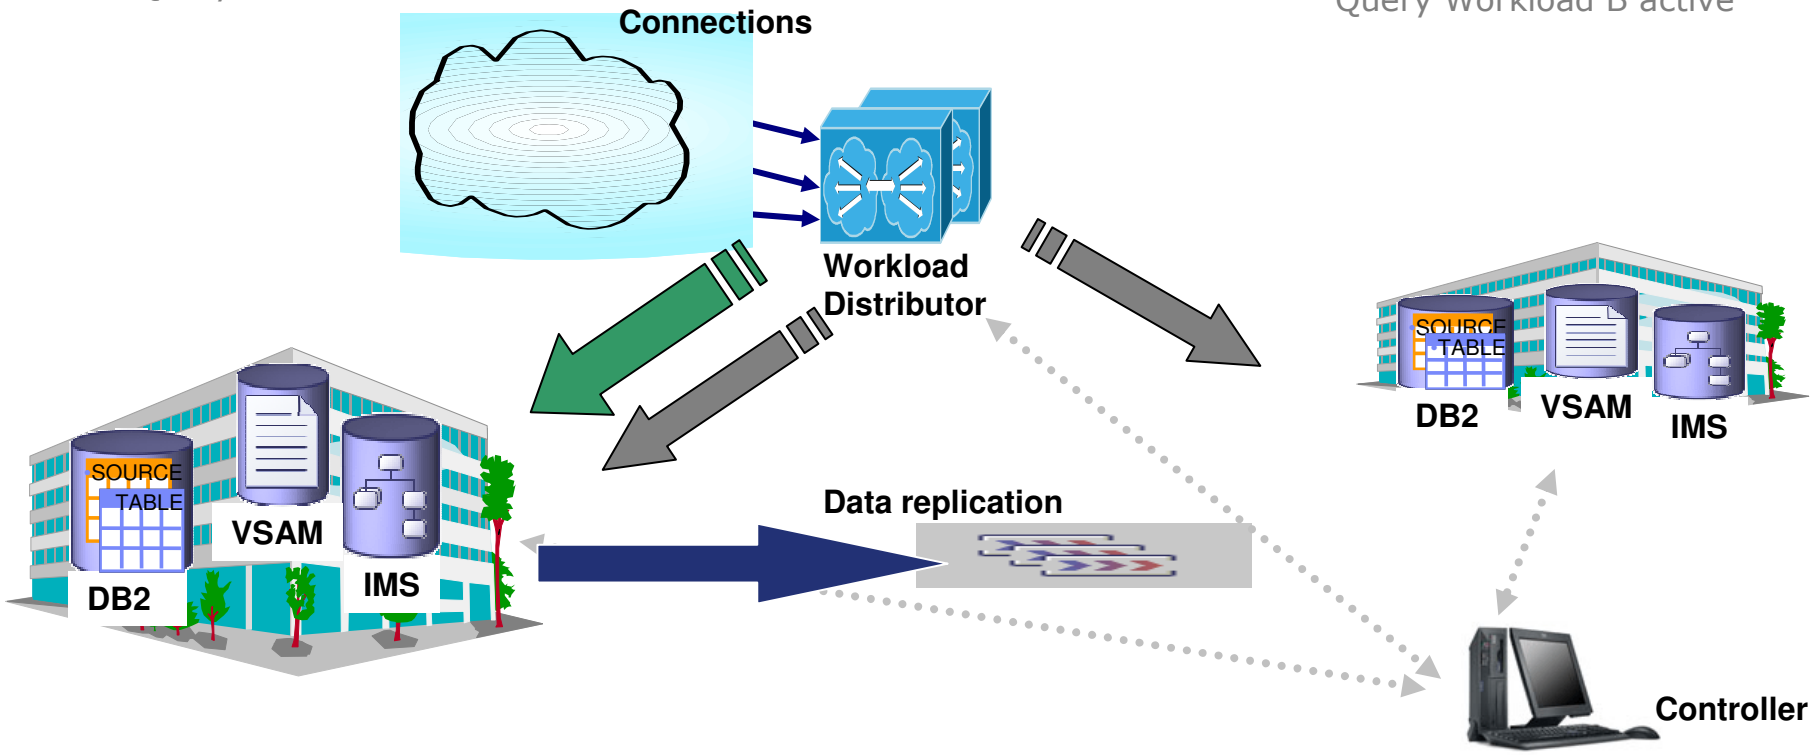

Replication latency, site capacity, and application health influence the distribution of query workload between sites.

Scenario 3: Active/Query Configuration –
Replication latency rises above maximum configured causing all query workload to be routed to Site 1

Site 1
Update Workload A active
Replication Latency: 'Acute'

Site 1
Query Workload B active

Site 2
Update Workload A standby

Site 2
Query Workload B active

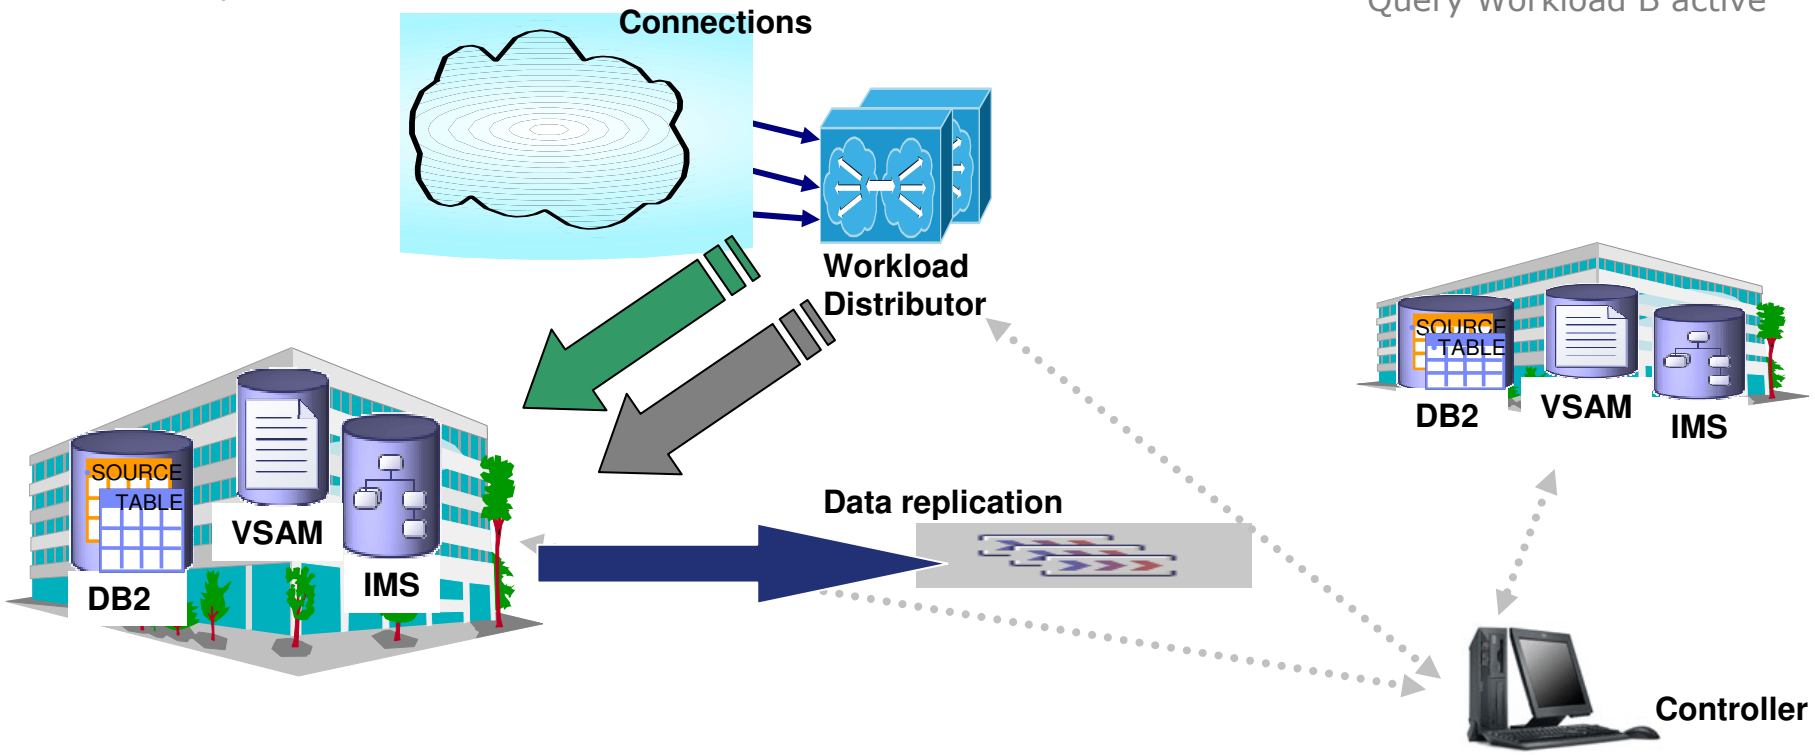

Replication latency, site capacity, and application health influence the distribution of query workload between sites.

Scenario 3: Active/Query Configuration –
Replication latency falls below minimum configured causing query workload to be distributed across two sites

Site 1
Update Workload A active
Replication Latency: 'Normal'

Site 1
Query Workload B active

Site 2
Update Workload A standby

Site 2
Query Workload B active

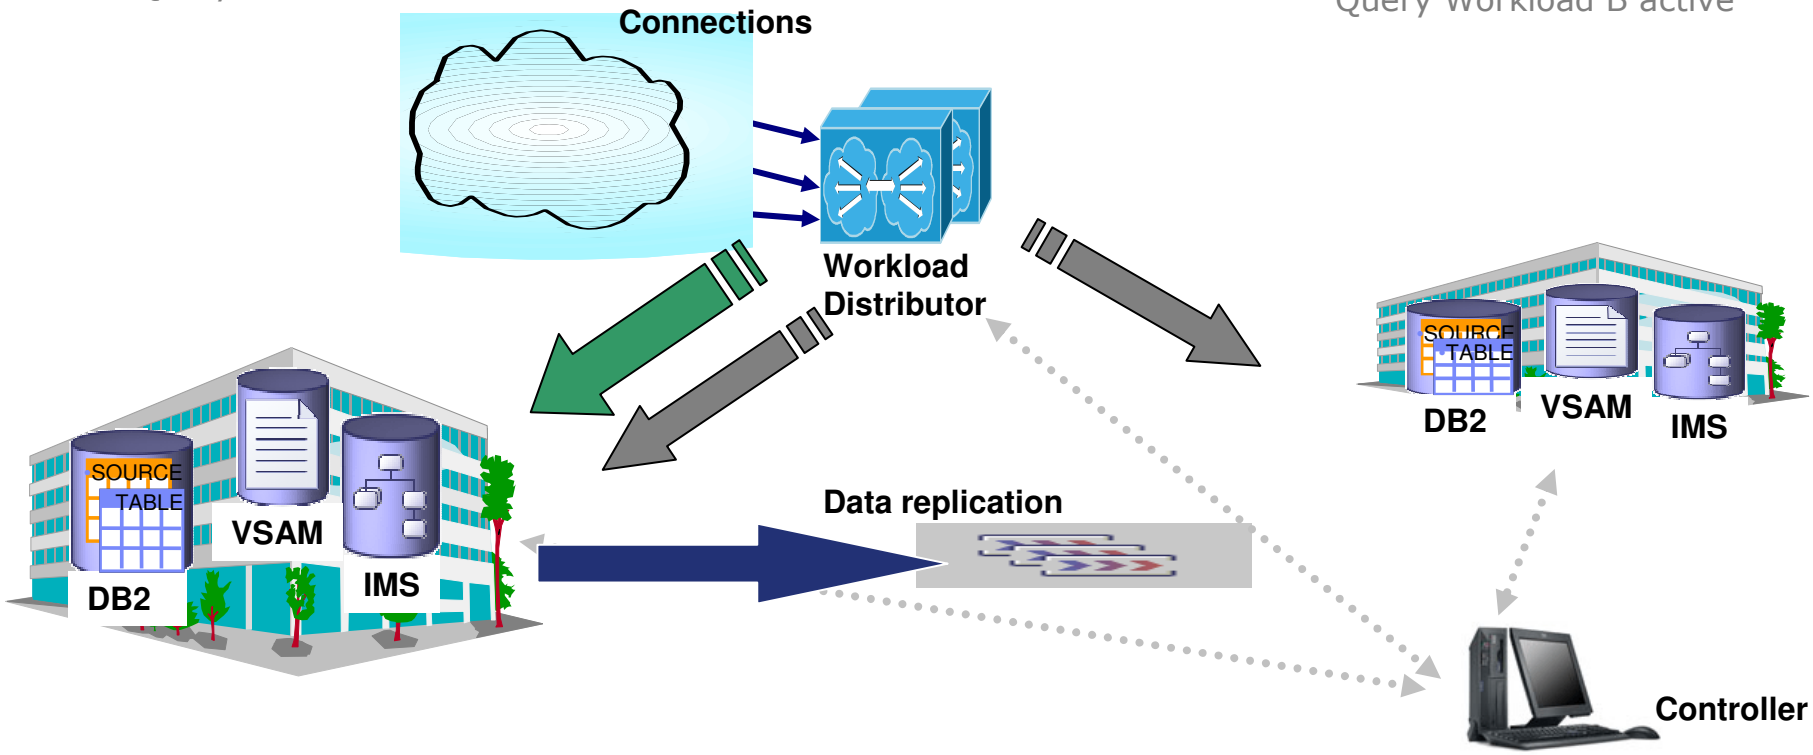

Replication latency, site capacity, and application health influence the distribution of query workload between sites.

Active/Active Sites Structure

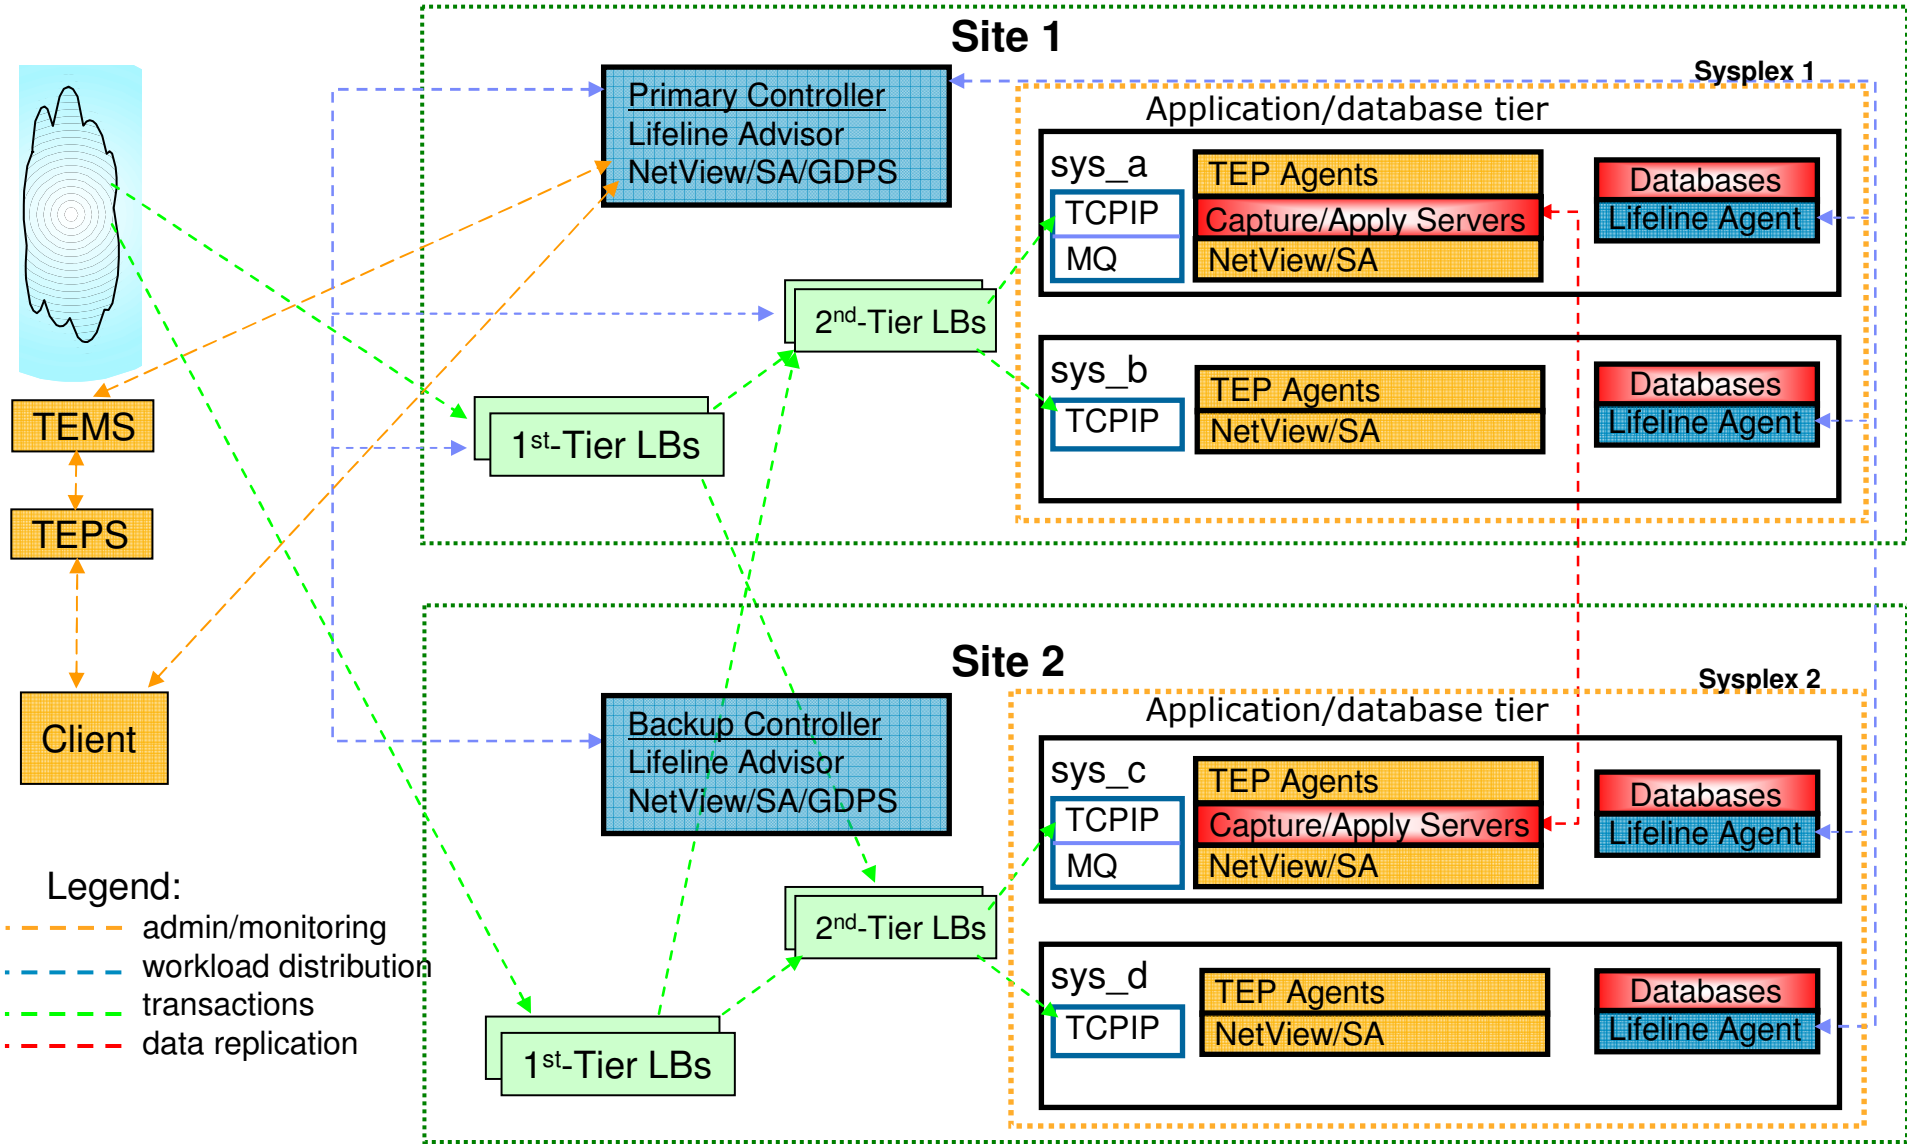

Workload Lifeline Providing Continuous Availability

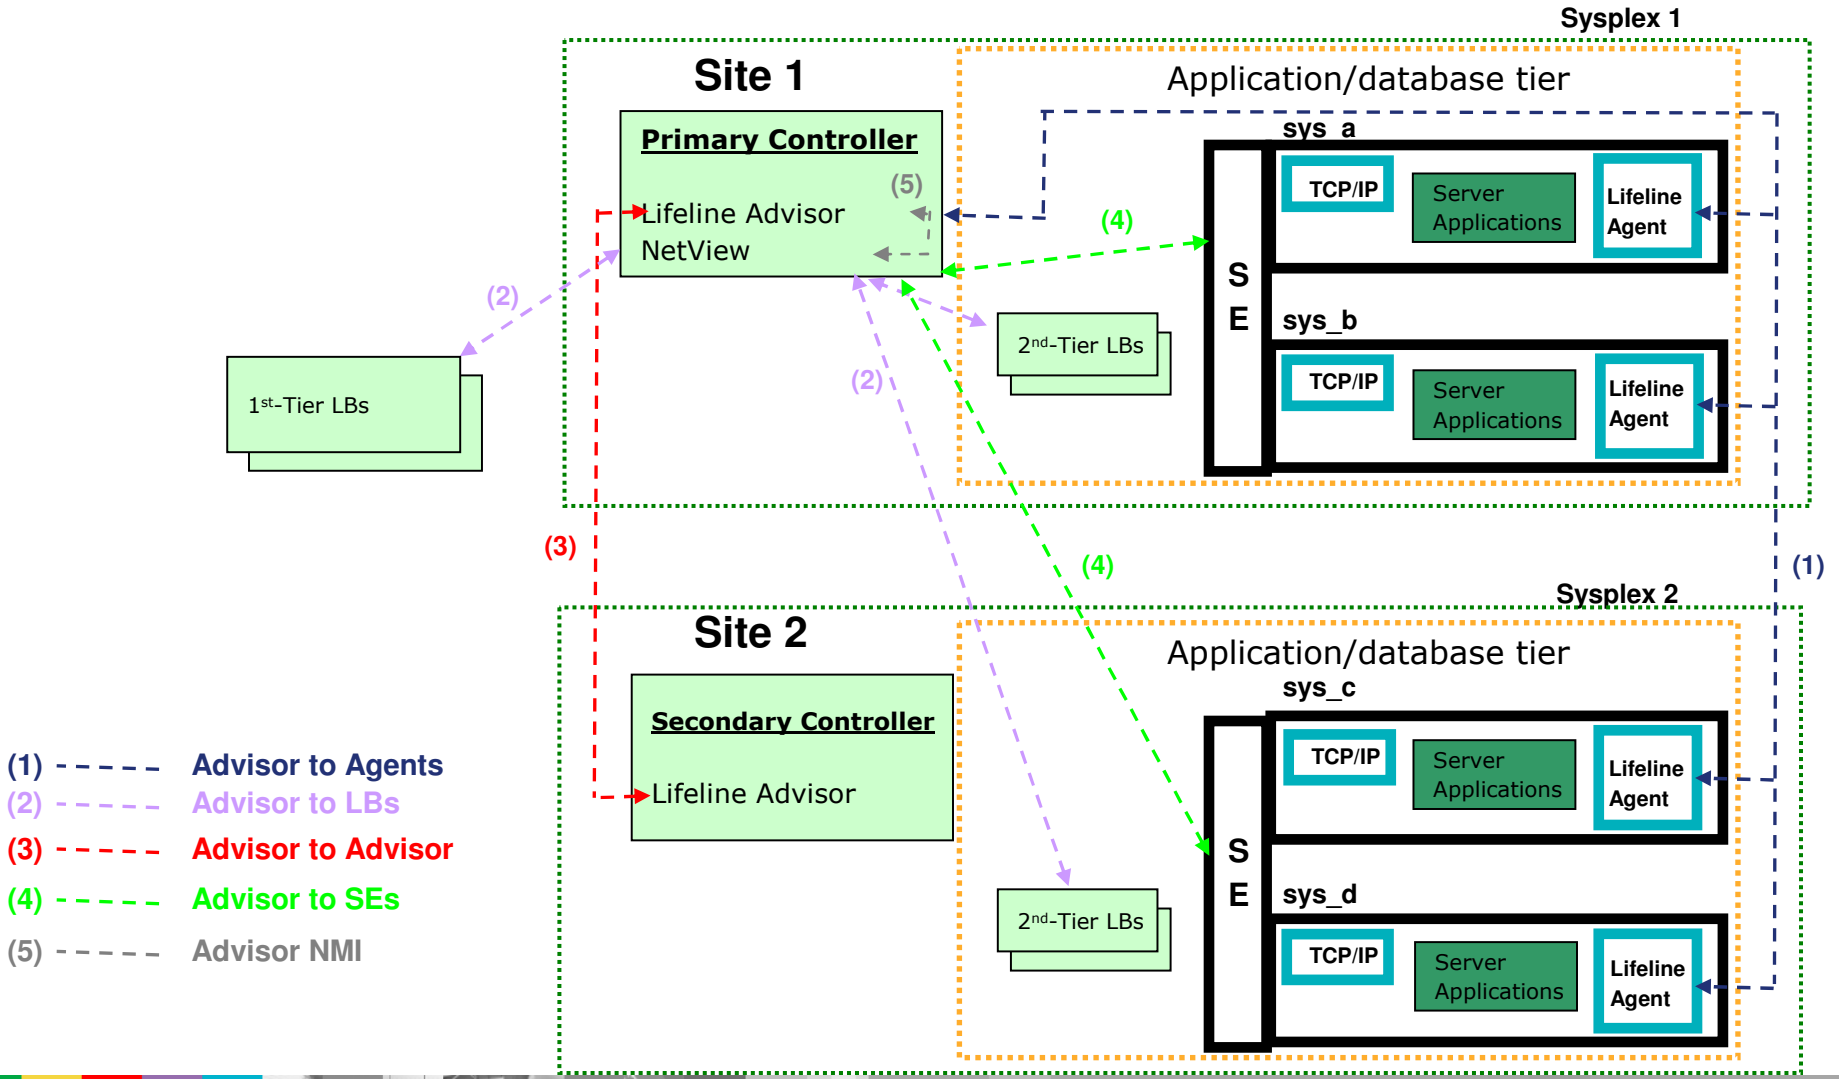

Reducing downtime for planned outages

Graceful workload movement load balancing requirements

- Ability to distribute workloads between sites
  - Based on customer-driven commands
- Ability to switch workloads from one site to another site
  - Perform “graceful” takeover for site maintenance
- Ability to maintain workload configuration states in event of a planned outage
  - Keep a peer workload manager in sync with workload states

Graceful workload movement z/OS product requirements

- Multi-site Workload Lifeline

- Depending on workload, data replication products that use Multi-site Workload Lifeline:
  - IBM InfoSphere Data Replication for DB2 database replication
  - InfoSphere IMS Replication for z/OS
  - InfoSphere VSAM Replication for z/OS

Multi-site Workload Lifeline providing graceful workload movement

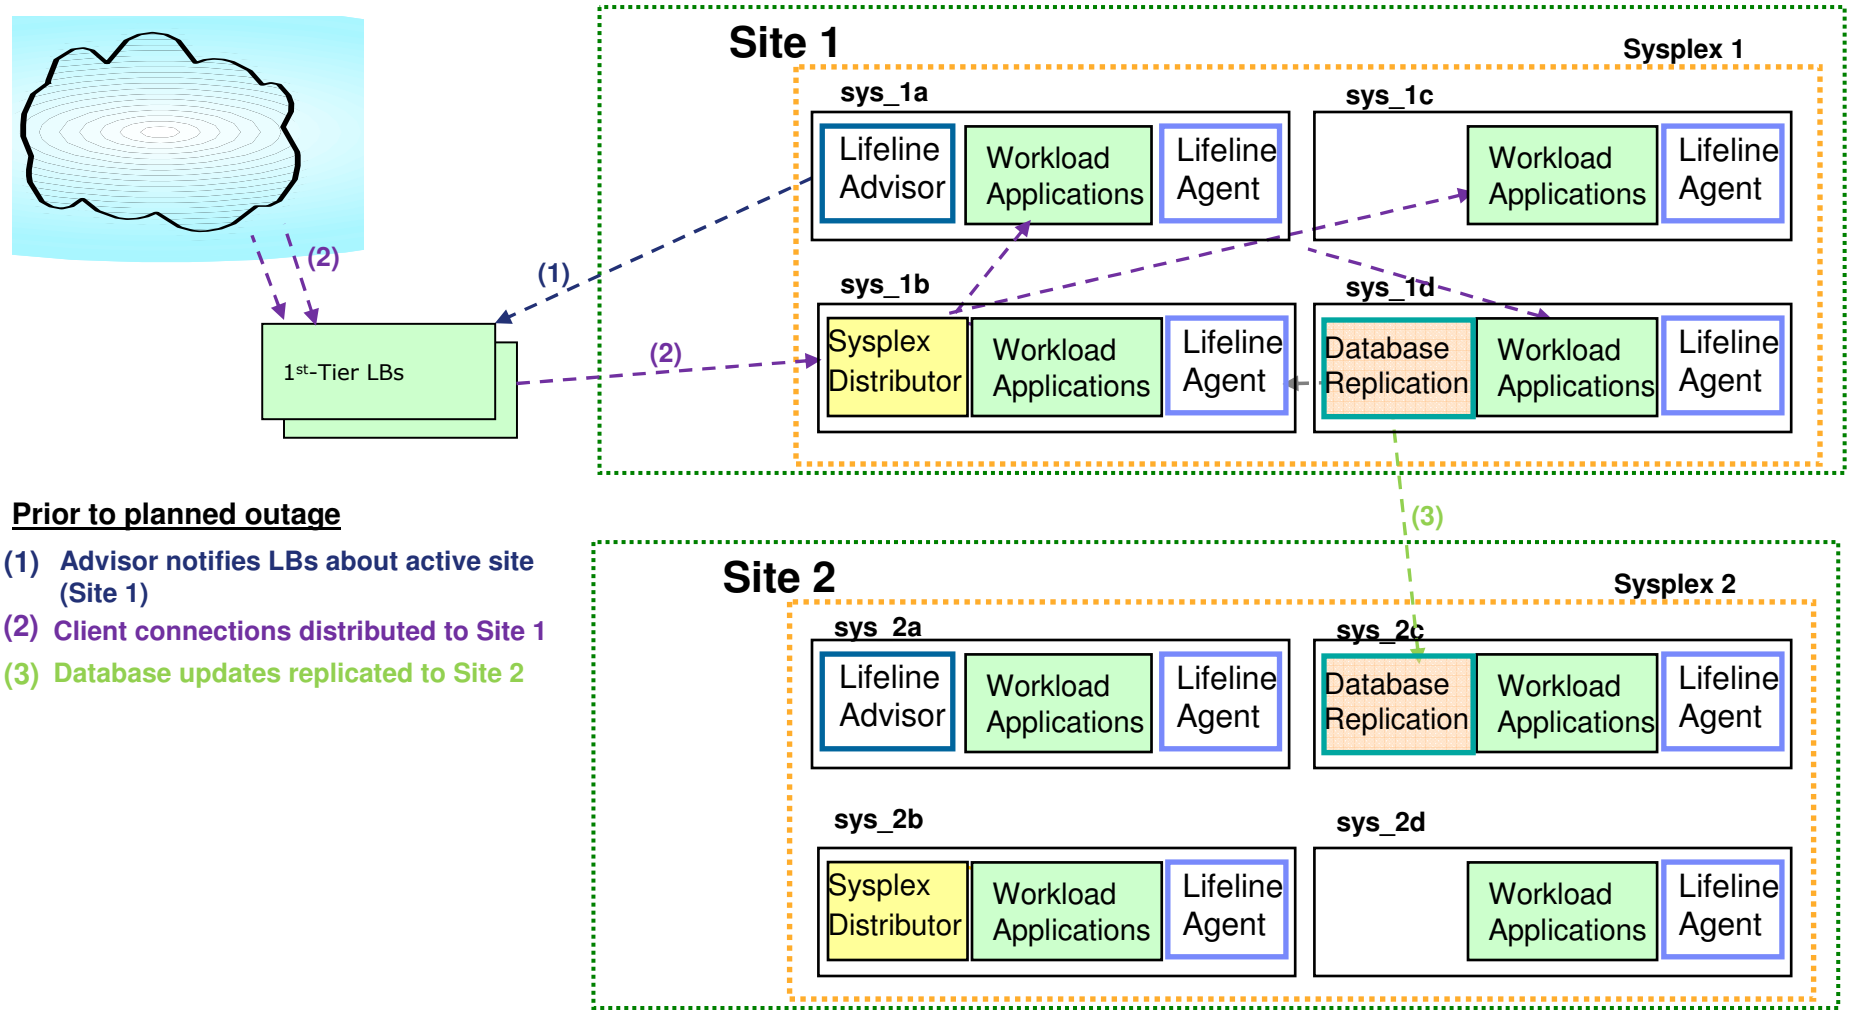

Prior to planned outage

(1) Advisor notifies LBs about active site (Site 1)

(2) Client connections distributed to Site 1

(3) Database updates replicated to Site 2

Multi-site workload Lifeline providing graceful workload movement…

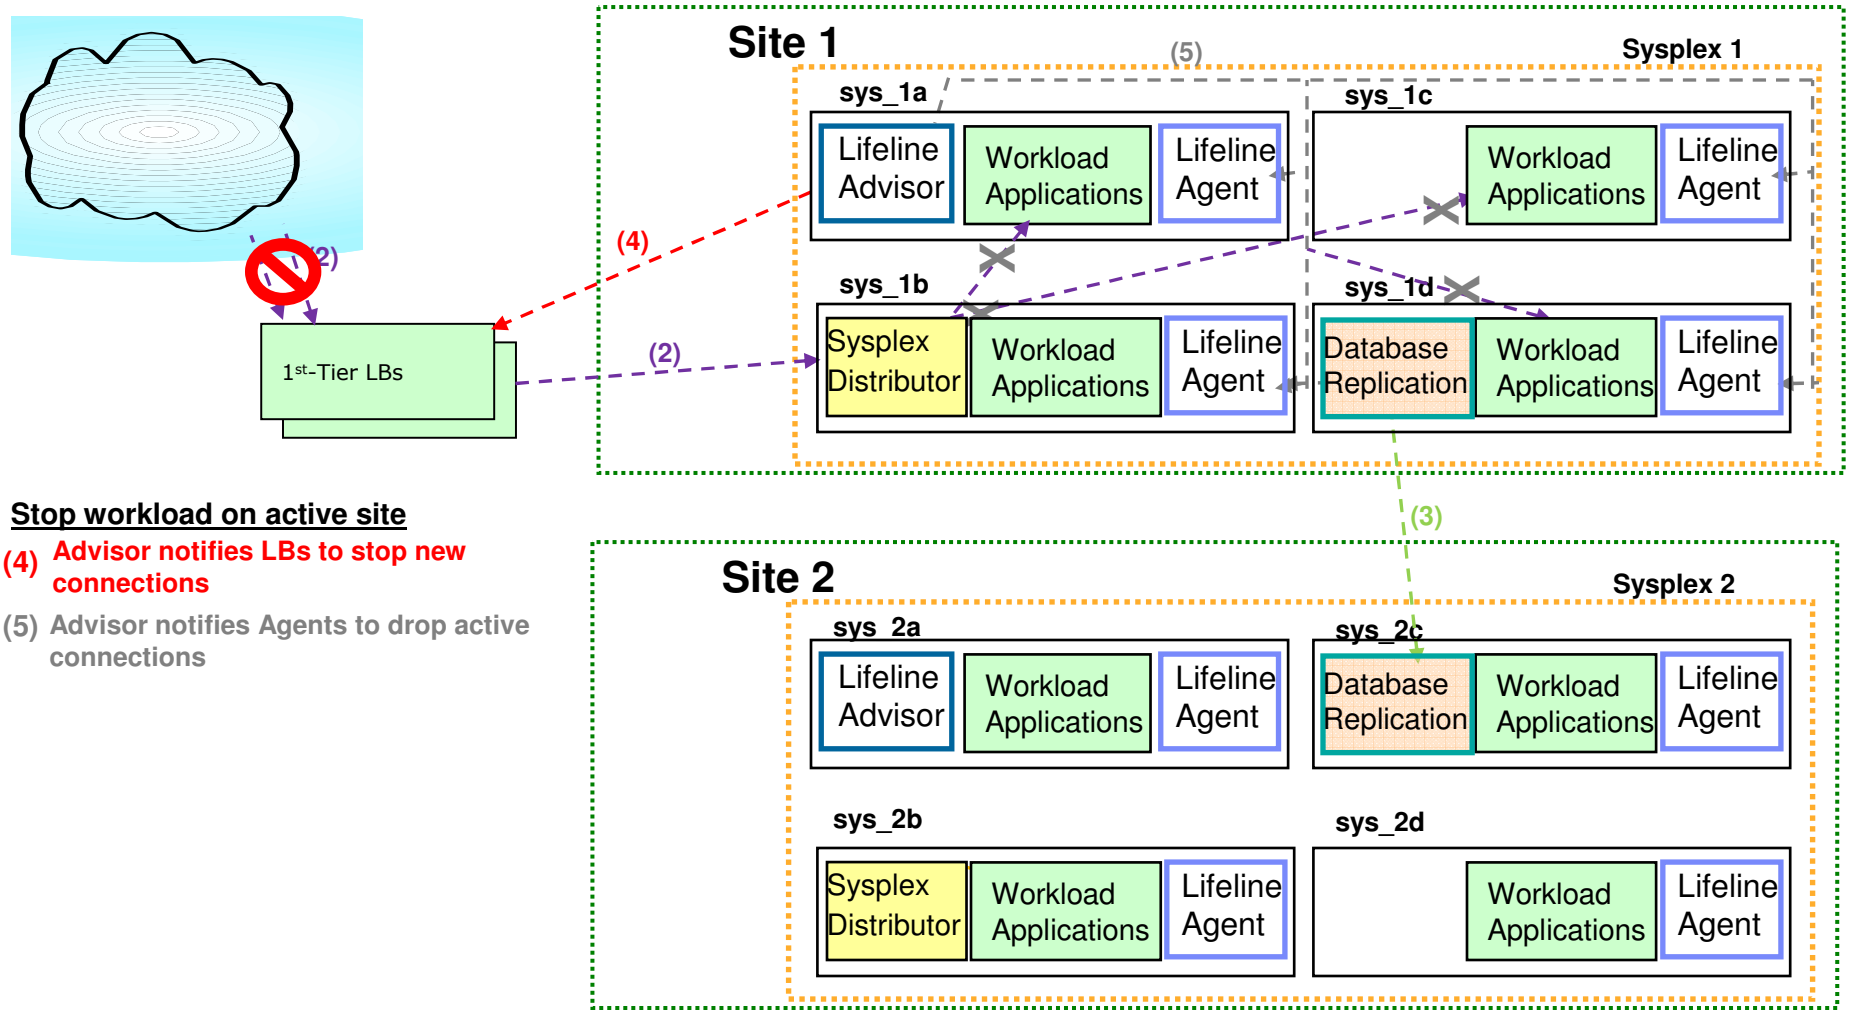

Stop workload on active site

(4) **Advisor notifies LBs to stop new connections**

(5) Advisor notifies Agents to drop active connections

Multi-site Workload Lifeline providing graceful workload movement…

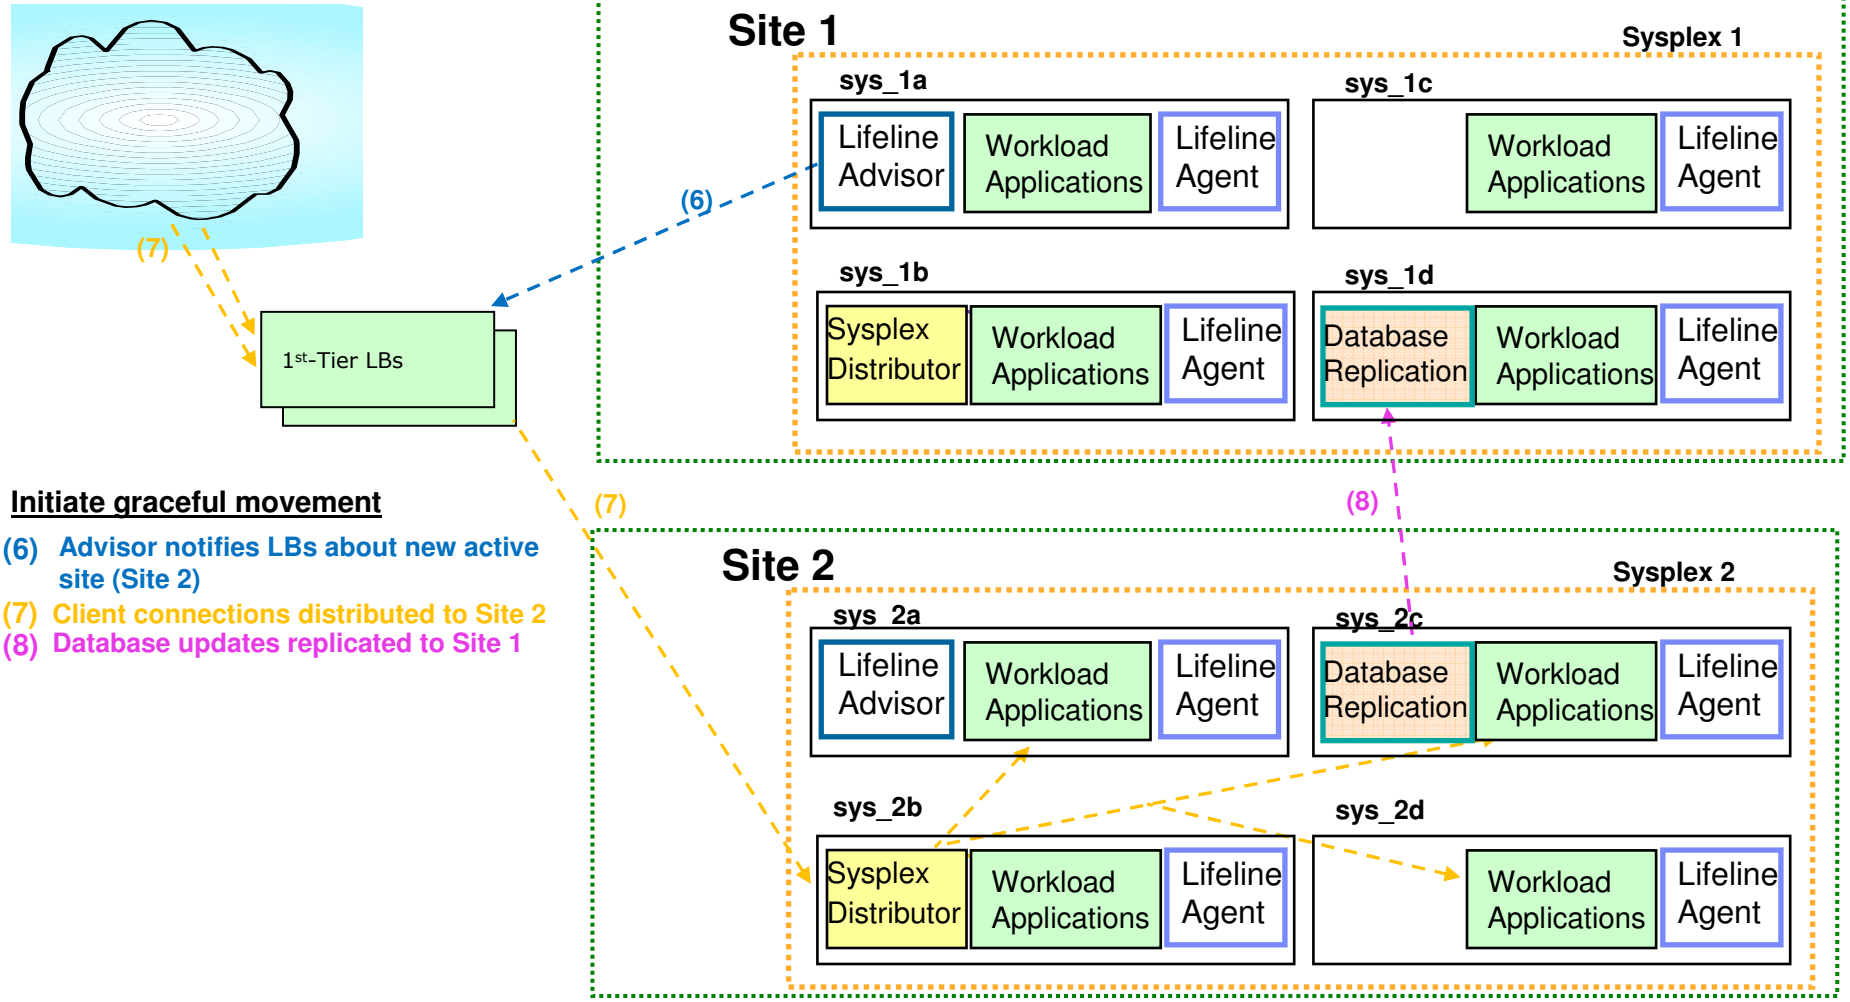

Initiate graceful movement

(6) Advisor notifies LBs about new active site (Site 2)

(7) Client connections distributed to Site 2

(8) Database updates replicated to Site 1

Multi-site Workload Lifeline commands

Key Advisor Display commands

- **MODIFY advproc,DISPLAY,ADVISOR,DETAIL**
  - When issued on the primary Advisor, displays the role of the Advisor, the connected load balancers (and whether it is a 1st-tier or 2nd-tier), the connected Agents (including system and site name where the Agents are active), and the connected peer Advisor (including the system name where the peer is active)
  - When issued on the peer Advisor, displays the role of the Advisor and the connected primary Advisor (including the system name where the primary is active)
- **MODIFY advproc,DISPLAY,CONFIG**
  - Displays the current active configuration for the Advisor
- **MODIFY advproc,DISPLAY,LB,DETAIL**
  - Displays the connected load balancers, including the list of groups registered by the load balancer, the members within each group, and the distribution recommendations provided for each member
- **MODIFY advproc,DISPLAY,WORKLOAD,DETAIL**
  - Displays the status of all defined workloads, including the status of all the server applications that make up the workload

Display Advisor information

F AQSADV,DISPLAY,ADVISOR,DETAIL
AQS0142I ADVISOR DETAILS
ADVISOR ROLE : PRIMARY
IPADDR : 192.10.1.1
LOAD BALANCERS:
IPADDR : 192.10.1.32
TIER : 1
IPADDR : 192.10.1.64
TIER : 2
AGENTS :
IPADDR : 192.10.110.1
SYSTEM NAME : SYS1 SITE : PLEX1
IPADDR : 192.10.110.2
SYSTEM NAME : SYS2 SITE : PLEX1
IPADDR : 192.20.110.1
SYSTEM NAME : SYS3 SITE : PLEX2
IPADDR : 192.20.110.2
SYSTEM NAME : SYS4 SITE : PLEX2
PEER ADVISOR :
IPADDR : 192.20.1.1
SYSTEM NAME : CNTL2

Display Workloads

F AQSADV,DISPLAY,WORKLOAD,DETAIL
AQS0146I WORKLOAD DETAILS
TYPE: ACTIVE/STANDBY
WORKLOAD NAME : WORKLOAD1
STATE : ACTIVE
SITE : PLEX2
SERVERS:
IPADDR..PORT : 192.10.110.1..5001
SYSTEM NAME : SYS1 SYSPLEX : PLEX1 STATUS : AVAIL
IPADDR..PORT : 10.20.1.1..5001
SYSTEM NAME : SYS3 SYSPLEX : PLEX2 STATUS : AVAIL
TYPE: ACTIVE/QUERY
WORKLOAD NAME : WORKLOAD2
STATE : ACTIVE
SITE : PLEX1
SITE : PLEX2
ASSOCIATED WORKLOAD : WORKLOAD1
REPLICATION STATE : NORMAL
SERVERS:
IPADDR..PORT : 192.10.111.1..8020
SYSTEM NAME : SYS1 SYSPLEX : PLEX1 STATUS : AVAIL
IPADDR..PORT : 10.21.1.1..8020
SYSTEM NAME : SYS3 SYSPLEX : PLEX2 STATUS : AVAIL
:

Key Advisor State Change commands

- **MODIFY advproc,ACTIVATE,WORKLOAD=…,SITE=…**
  - Signals the Advisor to direct 1st-tier load balancers to distribute new connections for the specified workload to the requested site
  - For Active/Query workloads, the SITE= keyword is optional, as the default is to activate the workload to both sites
- **MODIFY advproc,DEACTIVATE,WORKLOAD=…,SITE=…**
  - Signals the Advisor to direct Agents on the site where the specified workload was last active to reset any existing connections for this workload
  - For Active/Standby workloads, the SITE= keyword is not allowed, as the workload is deactivate only on the site where the workload was last activated
  - For Active/Query workloads, the SITE= keyword is optional, as the default is to deactivate the workload on both sites
- **MODIFY advproc,QUIESCE,WORKLOAD=…,SITE=…**
  - Signals the Advisor to direct 1st-tier load balancers to stop distributing new connections for the specified workload to any site
  - For Active/Standby workloads, the SITE= keyword is not allowed, as the workload is quiesced only on the site where the workload was activated
  - For Active/Query workloads, the SITE= keyword is optional, as the default is to quiesce the workload on both sites
- **MODIFY advproc,REFRESH**
  - Signals the Advisor to reread its configuration file and apply any updates to its active configuration
- **MODIFY advproc,TAKEOVER**
  - Signal the peer Advisor to take over primary Advisor responsibilities from the current primary Advisor

Key Agent Display commands

- **MODIFY ageproc,DISPLAY,CONFIG**
  - Displays the current active configuration for the Agent
- **MODIFY ageproc,DISPLAY,MEMBERS,DETAIL**
  - Displays information about each of the server applications this Agent was asked to monitor, including whether the server application exists, the jobname of the server application, and current state of the server application

Display Members information

F AQSAGE,DISPLAY,MEMBERS,DETAIL
AQS0115I MEMBER DETAILS
LB INDEX : 00 UUID : A67B6699
GROUP NAME : WKLD2_GROUP1
IPADDR..PORT: 10.10.1.1..8020
MATCHES : 001 PROTOCOL : TCP
FLAGS : ANY DISTDVIPA
TCPNAME : TCPIP
JOBNAME : JOB1 ASID : 0034 RESOURCE : 0000096B
GROUP NAME : WKLD2_GROUP2
IPADDR..PORT: 10.10.1.1..8021
MATCHES : 000 PROTOCOL : TCP
FLAGS : DISTDVIPA
TCPNAME : TCPIP
JOBNAME : N/A ASID : N/A RESOURCE : N/A
LB INDEX : 01 UUID : 9A78BE9E
GROUP NAME : TIER2_GROUP1
IPADDR..PORT: 192.10.110.1..5001
MATCHES : 001 PROTOCOL : TCP
FLAGS :
TCPNAME : TCPIP
JOBNAME : JOB3 ASID : 0036 RESOURCE : 0000096D
GROUP NAME : TIER2_GROUP2
IPADDR..PORT: 192.10.110.1..6001
MATCHES : 001 PROTOCOL : TCP
FLAGS :
TCPNAME : TCPIP
JOBNAME : JOB4 ASID : 0037 RESOURCE : 0000096E

Key Agent State Change commands

- **MODIFY ageproc,ENABLE,…**
  - Signals the Agent to enable server applications (make them available to be load balanced to)
    - Server applications bound to a distributable dynamic VIPA must be enabled using the VARY TCPIP,,SYSPLEX,RESUME command
- **MODIFY ageproc,QUEISCE,…**
  - Signals the Agent to quiesce server applications (make them unavailable to be load balanced to)
    - Server applications bound to a distributable dynamic VIPA must be quiesced using the VARY TCPIP,,SYSPLEX,QUIESCE command

For more information…

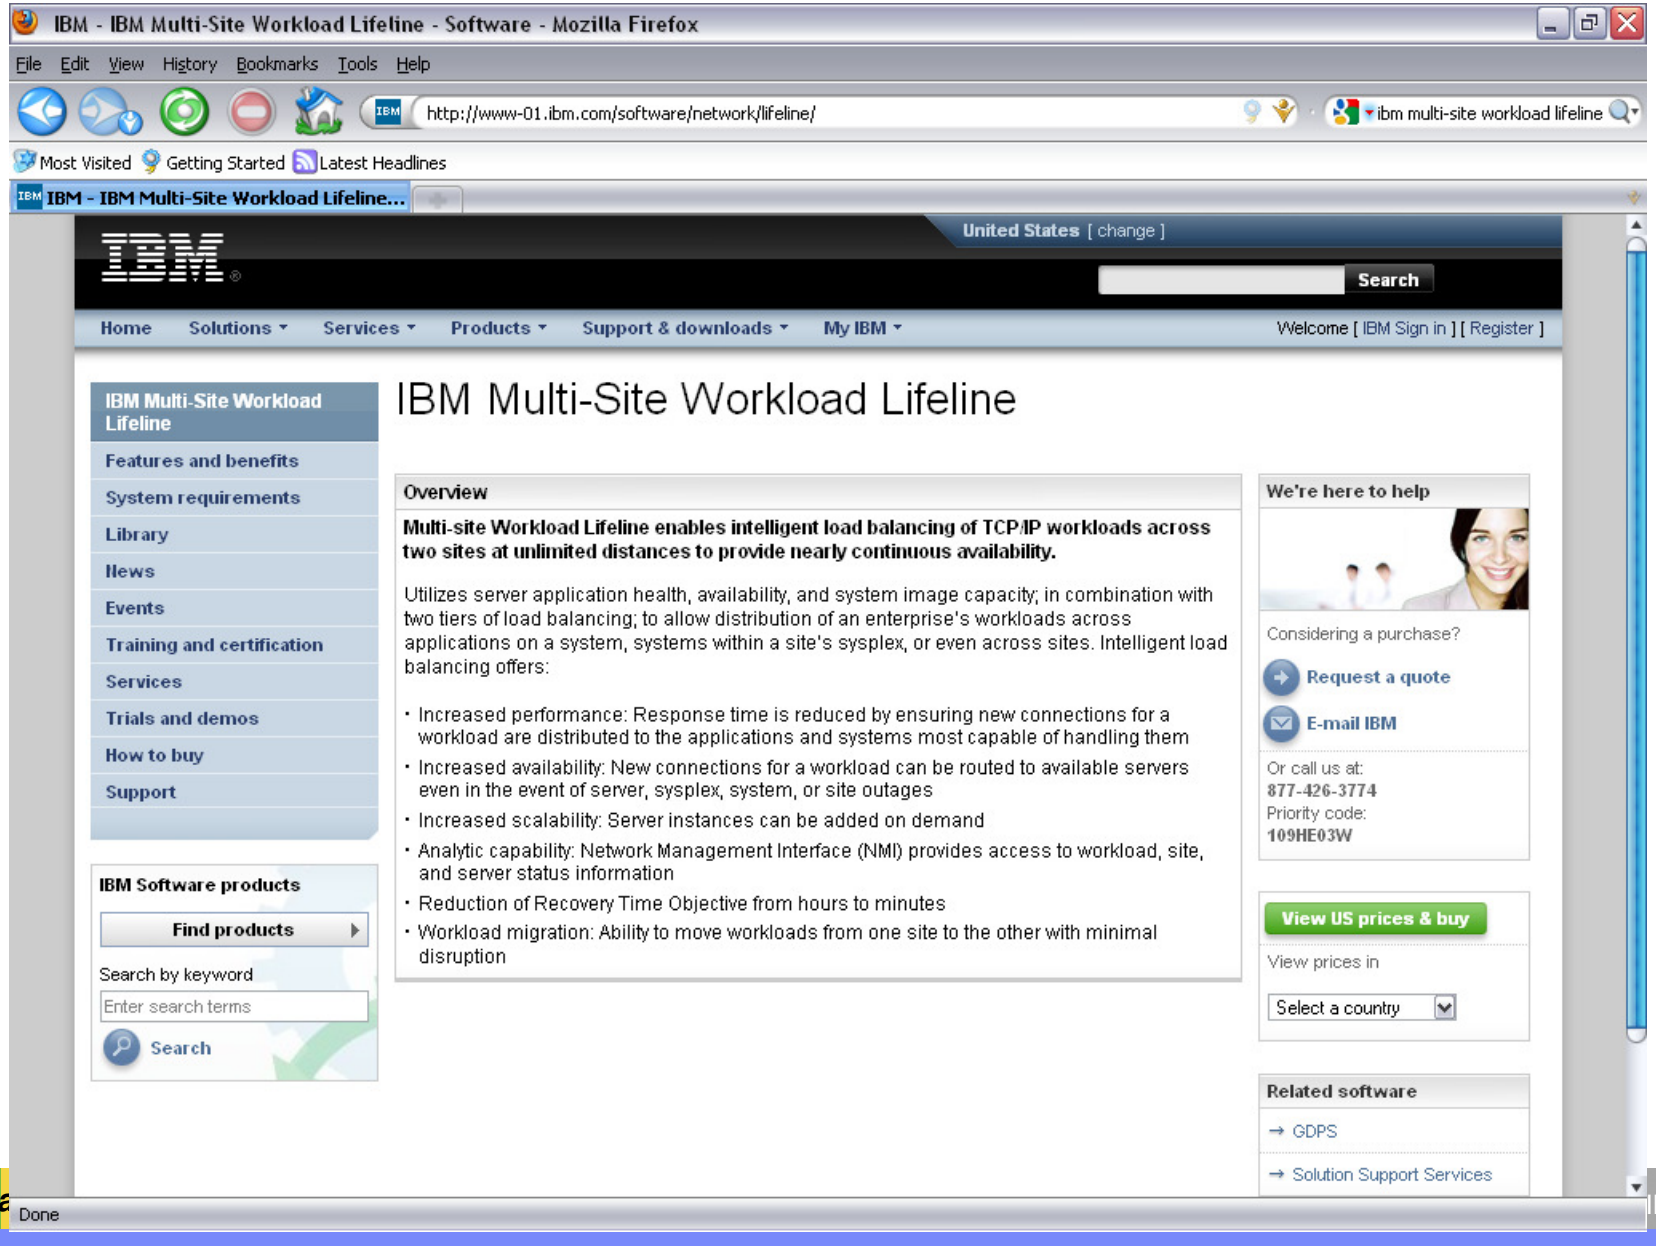

Please fill out your session evaluation

- Enabling Continuous Availability and Reducing Downtime with IBM Multi-site Workload Lifeline
- Session # 15197
- QR Code: