

::ISUG

Replication Server 12.6

techcast
series

Performance, Ease of Use and Reducing Administration

This live, interactive seminar will begin shortly

For audio, please dial:

Toll Free within the U.S. & Canada: 1-888-459-9364

International: 1-210-234-6662

Passcode: ISUG

techcast
series

Performance, Ease of Use and
Reducing Administration



Thomas Lamb
Moderator
Past President
Sybase International User Group

techcast
series

Performance, Ease of Use and
Reducing Administration



Irfan Khan
Presenter
*Director of Engineering
Sybase Evangelist Team*

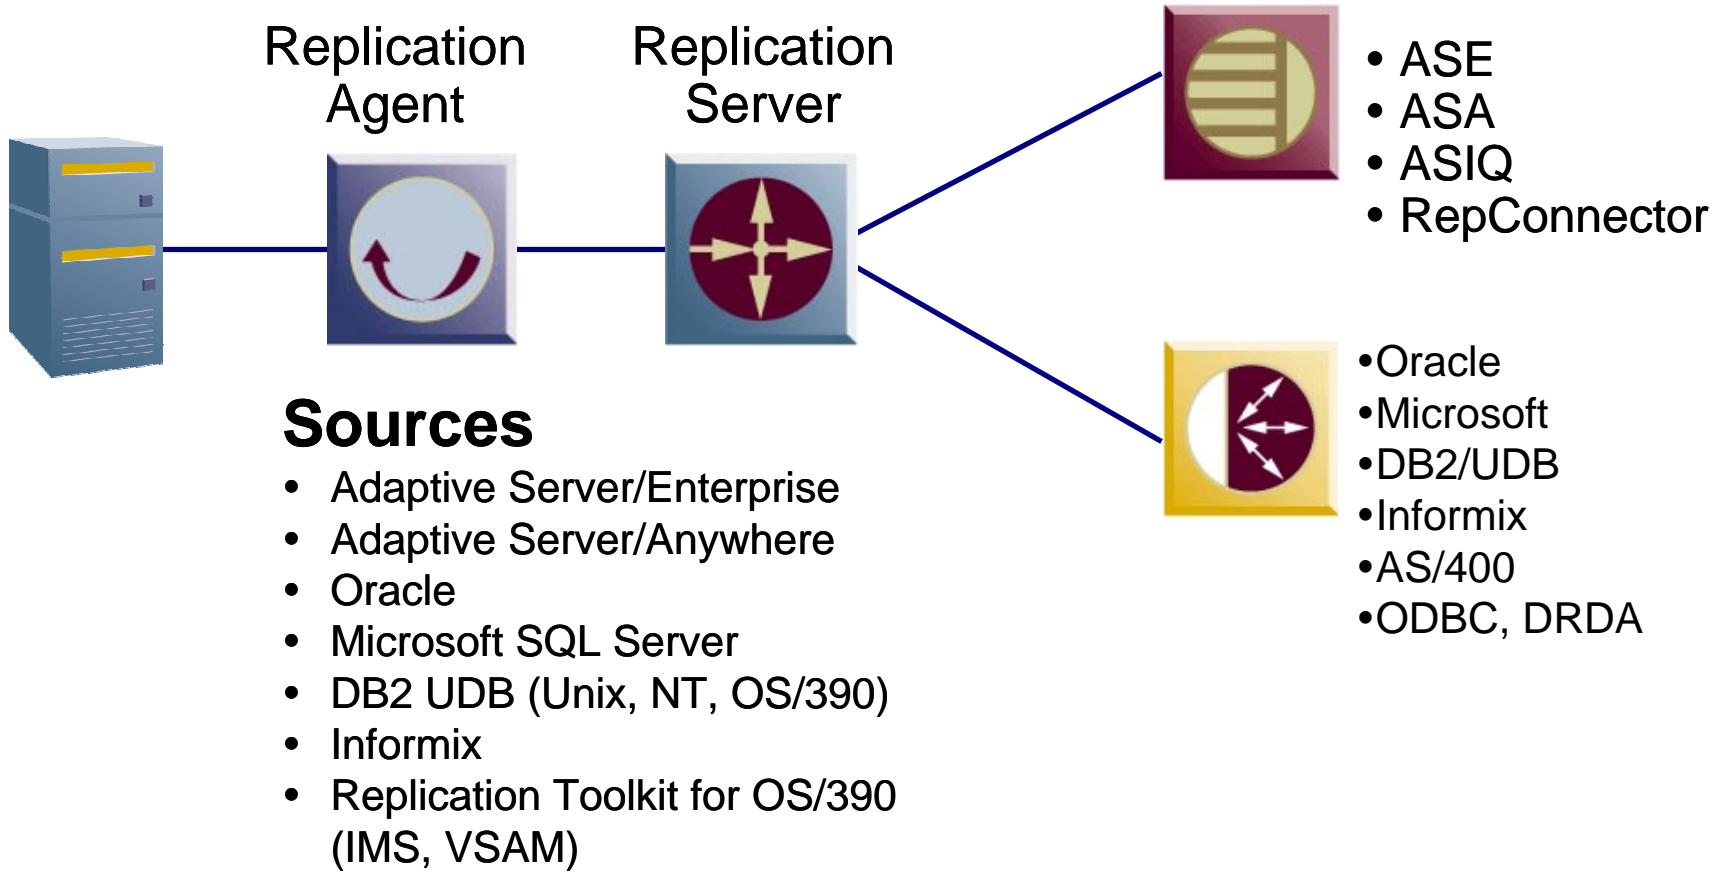
- **Introduction**
- **Replication Architecture**
- **RS 12.6 features**
- **Demo**
- **RS Futures Roadmap**
- **Q & A**

- **High Availability of data**
Planned, Unplanned downtime and Disaster Recovery
- **Data Distribution**
Decentralizations
Consolidations
Load balancing
- **Live decision support**
Separation of OLTP and DSS, reporting server

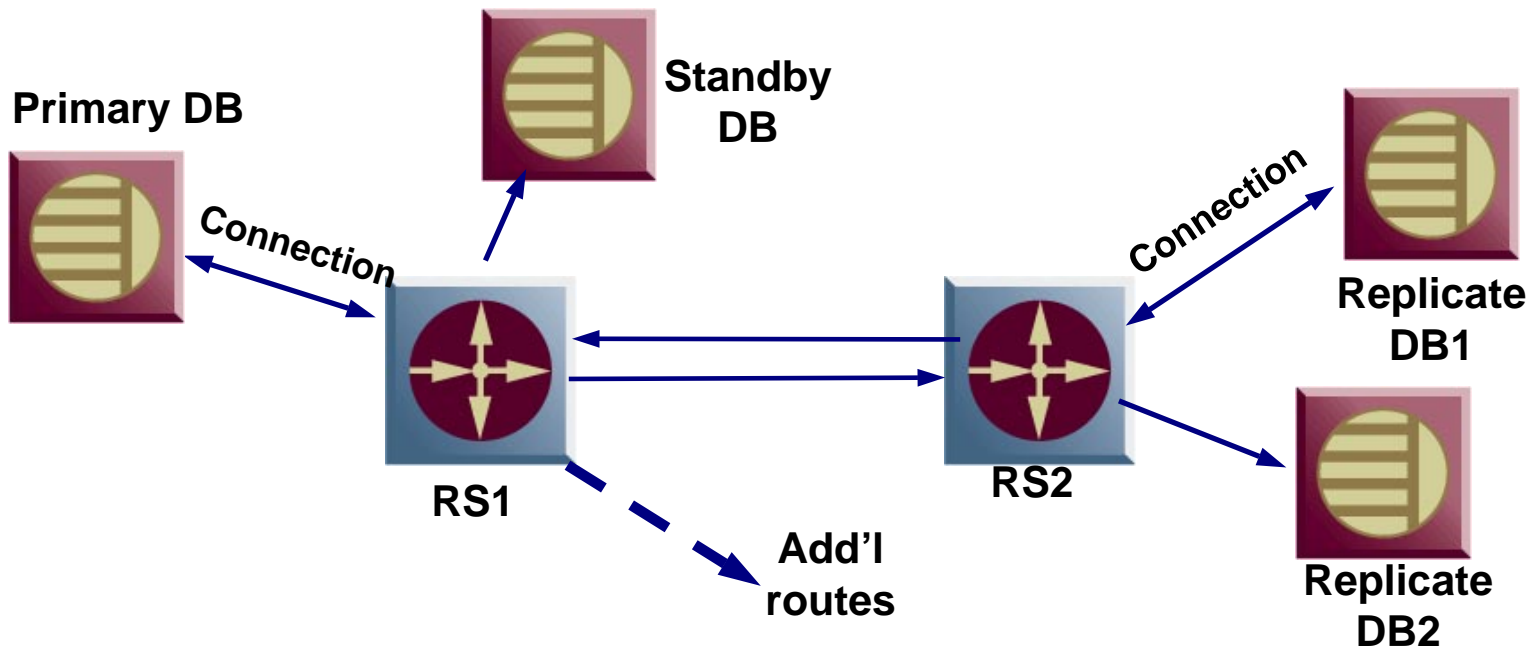
- Introduction
- **Replication Architecture**
- RS 12.6 features
- Demo
- RS Futures Roadmap
- Q & A

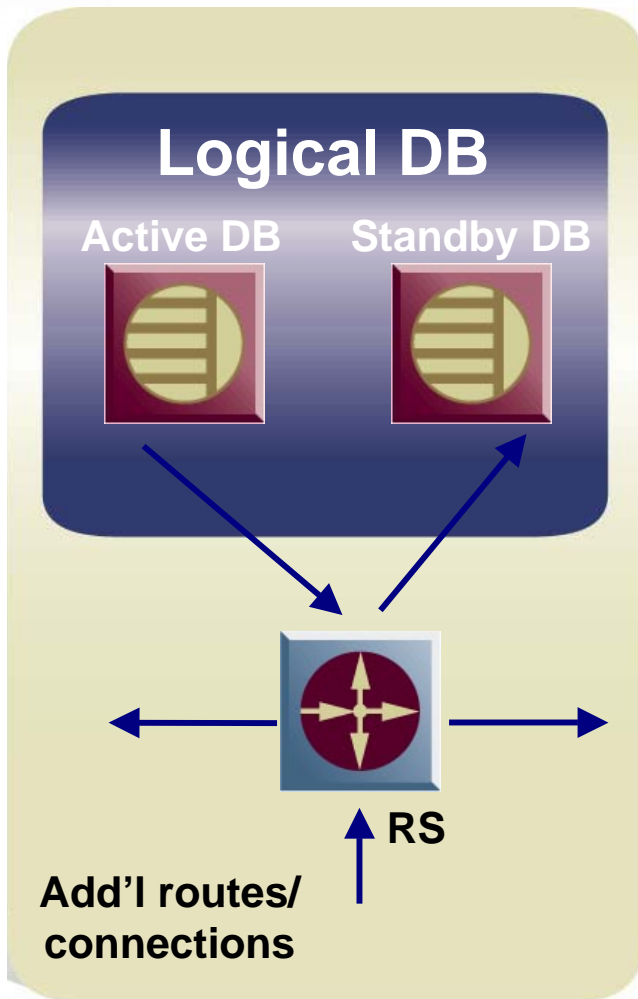
Primary Sites

Replicate Targets



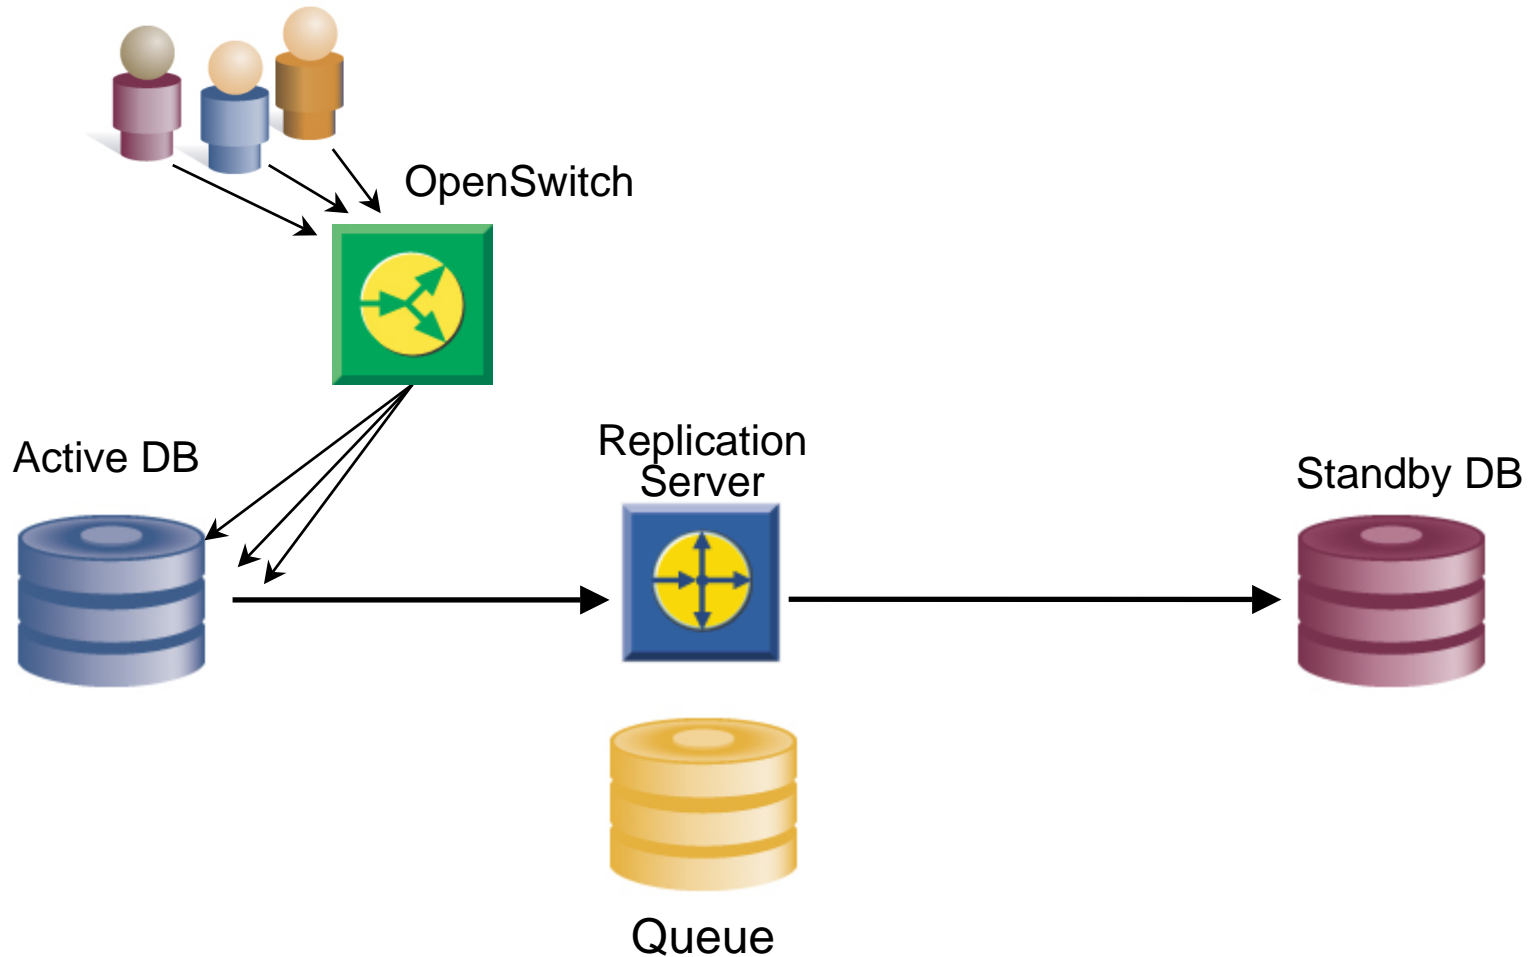
- System Configuration
 - Repservers manage connections to databases
 - Repservers exchange data via routes





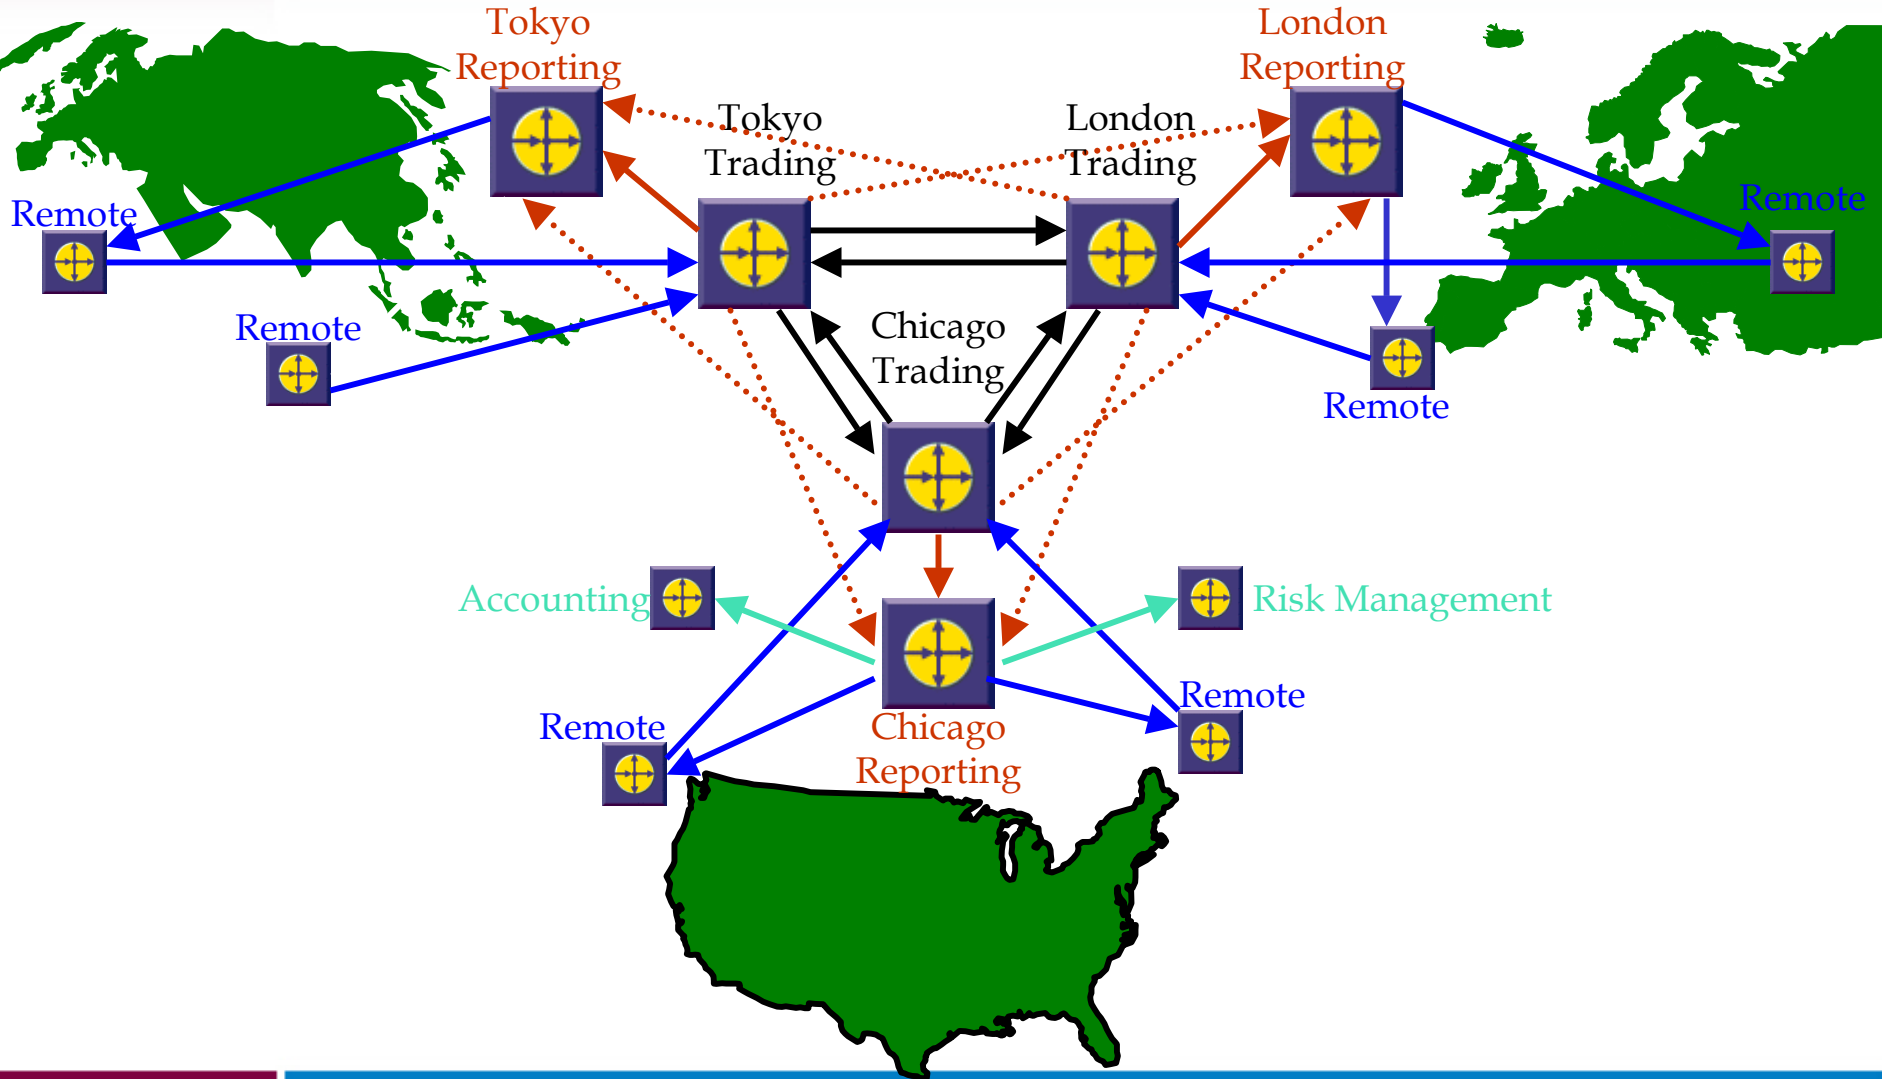
- Special Configuration: One Active, One Standby Database managed by One Rep Server
- May be configured to replicate the entire database – Complete Mirror! (publish-subscribe mechanism not required)
- DDL/Schema Change replication!
- Direction of Replication can be switched
 - Performs failover and failback
- Ease of Administration, Low latency
- Active-Standby pair viewed as a single logical DB at other replicate sites

Warm Standby with Open Switch



- Routes allow Replication Server environments to scale up to 100s of sites
 - Routing eliminates the need for a “star” configuration (each site connected to every other site) between numerous sites – sites could be managed in a hierarchical fashion
- A route is a connection between two Replication Servers
 - Allows them to exchange data over a physical link
 - Indirect Route is a logical connection between Replication Servers that have no physical link between them
 - Each data item passed on a route is marked with the “destination” site(s) – the receiver passes it on appropriately
 - Only one copy of the data item is passed over the network intended for multiple target sites – conserve bandwidth and processing

Global, Enterprise-wide distribution for a multi-national bank's trading system



- Function Strings allow users to configure the application of transactions at the target site. Function strings allow:

Data transformations, Operational transformations, conflict resolution, replication to proprietary targets (custom language), and more

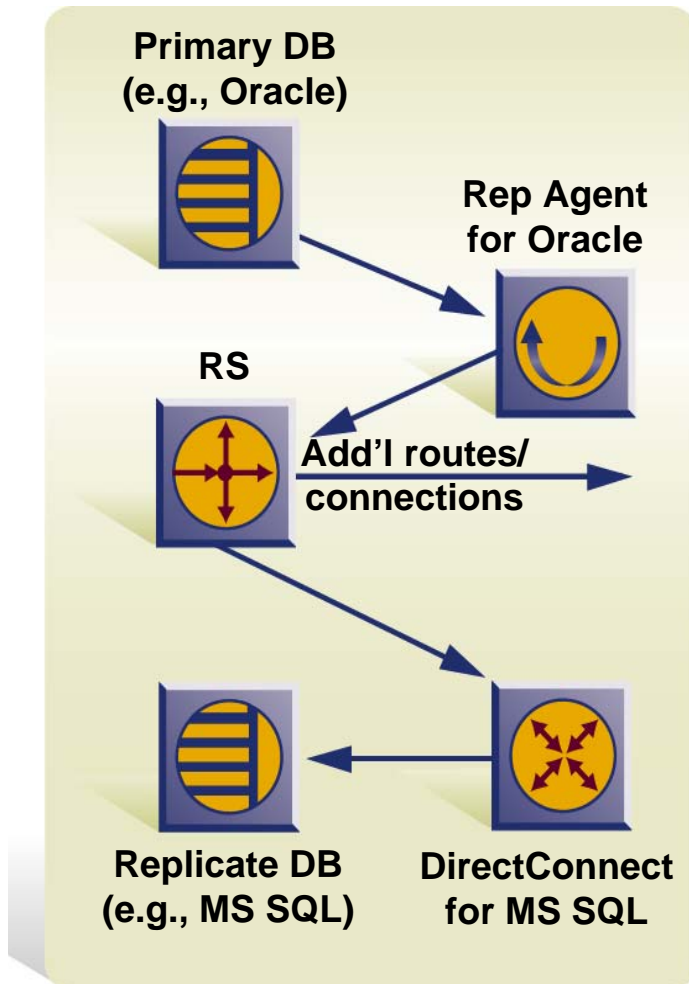
- A Function String is a template for each operation (insert/update/delete) that can be customized by the user

The function string, if changed by the user, is used to apply operations (transactions) at the target site

Data transformations like aggregation of columns, data type conversions, conversion between different DB formats etc. can all be accomplished

Deletes can be turned into inserts, inserts can be turned into stored procs, updates can be suppressed, each operation can be turned into a proprietary language to be sent to the target, etc.

Multi-Vendor Replication



- Replication Agents pull data out of non-Sybase sources
- DirectConnects are used as gateways to replicate into non-Sybase targets
- Support for vendor-specific datatypes, translations etc.
- Some Replication Agents read the transaction log (DB2), some build a virtual log with triggers (Oracle, MS SQL)

Replication
Server
Base
Package

Replication Server Base Server License



Replication Server Manager
Sybase Central Plug-in for
Replication Server

Replication Server - Advanced Security Option

Replication
Server
Options
Package

Option for
Oracle



RA




DC



ASE

Option for
Informix



RA




DC



ASE

Option for
Microsoft



RA



DC



ASE

Option for
DB2/UDB



RA



DC



ASE

- Introduction
- Replication Architecture
- **RS 12.6 features**
- Demo
- RS Futures Roadmap
- Q & A

- SMP Enabled Replication Server
- Other Performance Features
- Multi-Site Availability
- Embedded RSSD
- Sybase Central Java Plugin
- New Date and Time datatypes
- Clustering support for Solaris, HP, NT

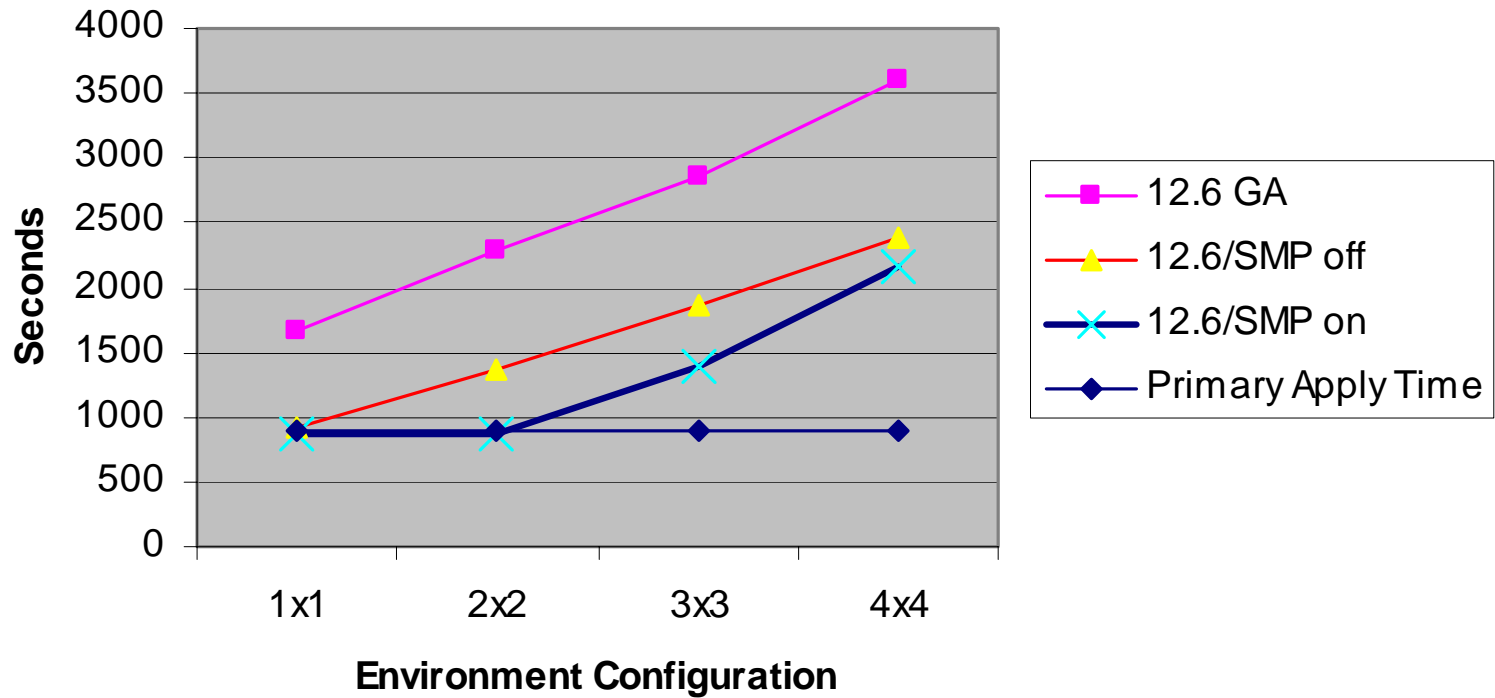
- Replication Server threads are managed by Sybase Open Server
- Sybase Open Server threading model puts thread management responsibility on the Operating System when SMP is enabled
- Using this “preemptive” capability, to the extent that the Operating System supports parallel execution of multiple user threads in a given process, RepServer is multi-processor enabled.
- Continued concentrated focus on improving performance in each release

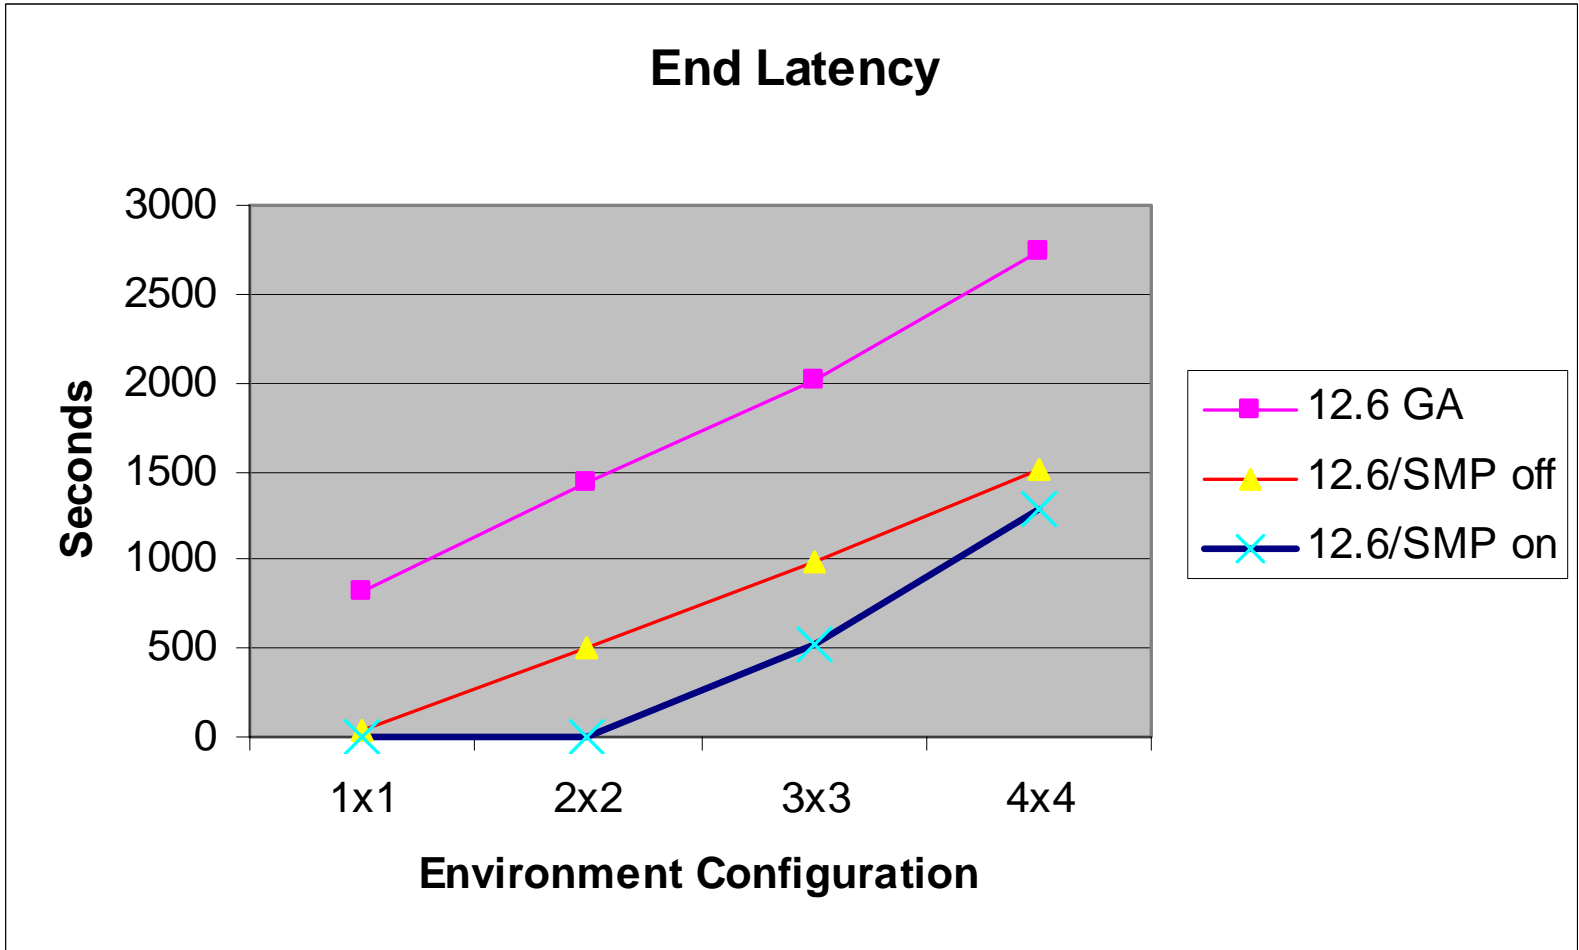
- SMP Performance white paper available titled “**Replication Server® 12.6 Performance**” at Sybase.com
- Performance benchmark “reference implementation” & Monitors & Counters analysis package available on SDN(Sybase Developers Network) as free download

	Primary Apply Time	Replicate Apply Time	Tran App Diff	Latency (start)	Latency (end)	Total Time	Throughput
1x1							
12.6 GA	902	1666	85%	51	815	1718	1.36
12.6/SMP off	902	912	1%	31	41	943	2.42
12.6/SMP on	902	861	-5%	43	2	904	2.55
2x2							
12.6 GA	903	2283	153%	53	1433	2336	1.91
12.6/SMP off	903	1373	52%	26	497	1400	2.78
12.6/SMP on	903	866	-4%	40	2	905	4.73

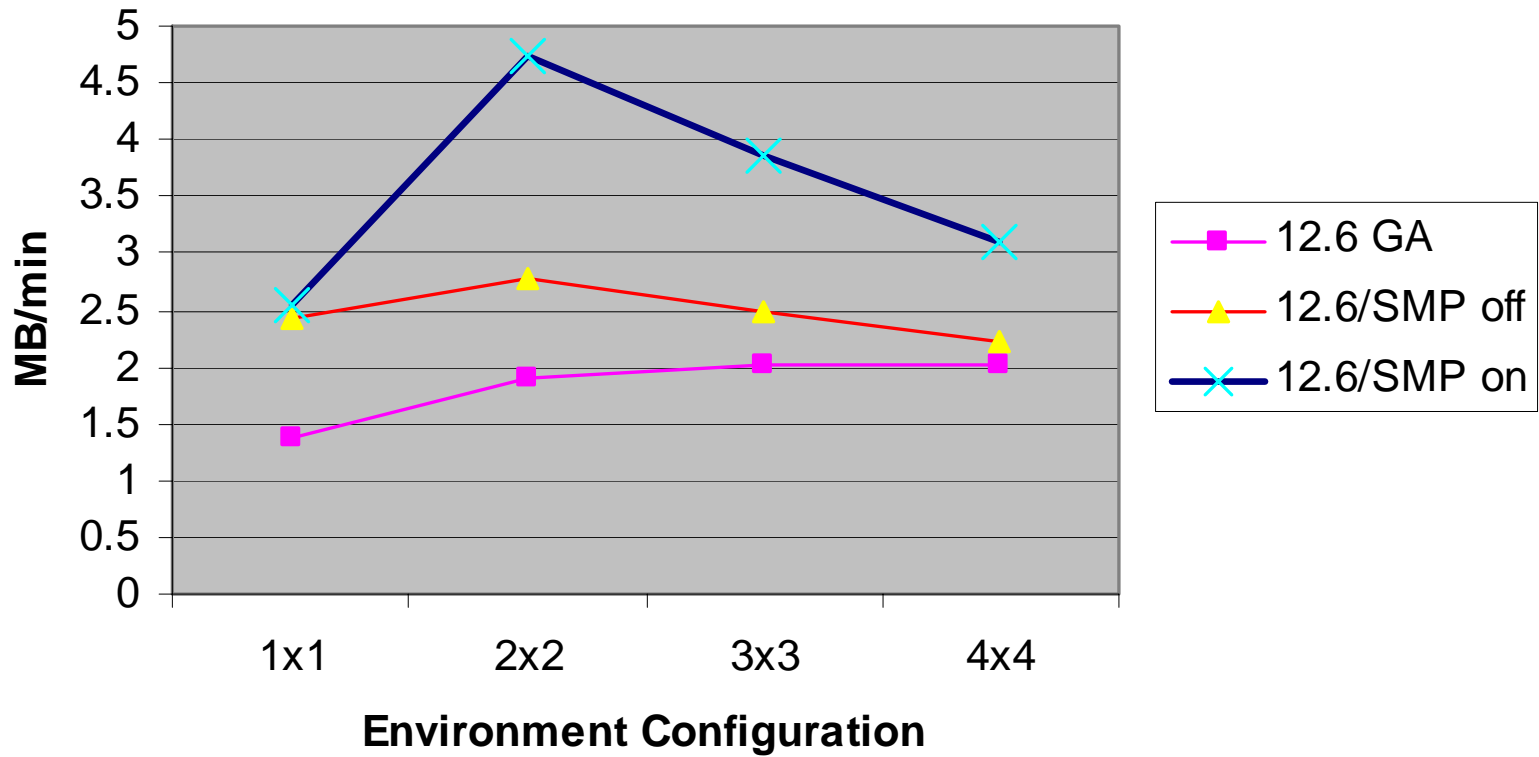
12.6 SMP vs. 12.6 GA on	Primary Apply Time	Replicate Apply Time	Tran App Diff	Latency (start)	Latency (end)	Total Time	Throughput
1x1	0%	-48%	-105%	-16%	-100%	-47%	88%
2x2	0%	-62%	-103%	-25%	-100%	-61%	148%

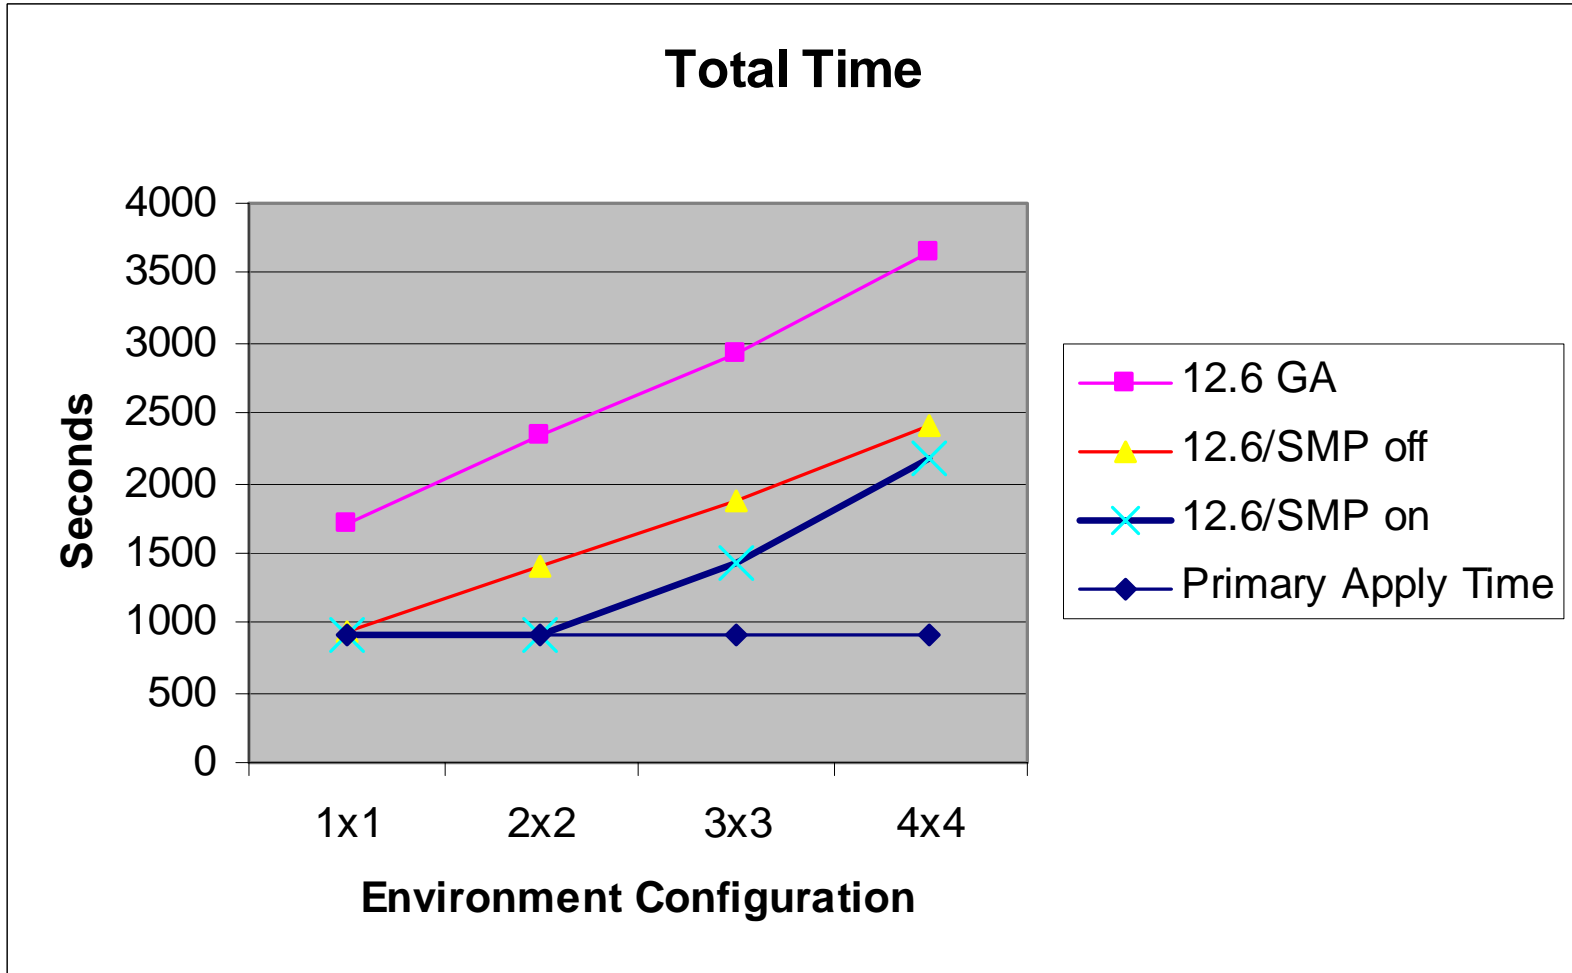
Replicate Apply Time





Throughput

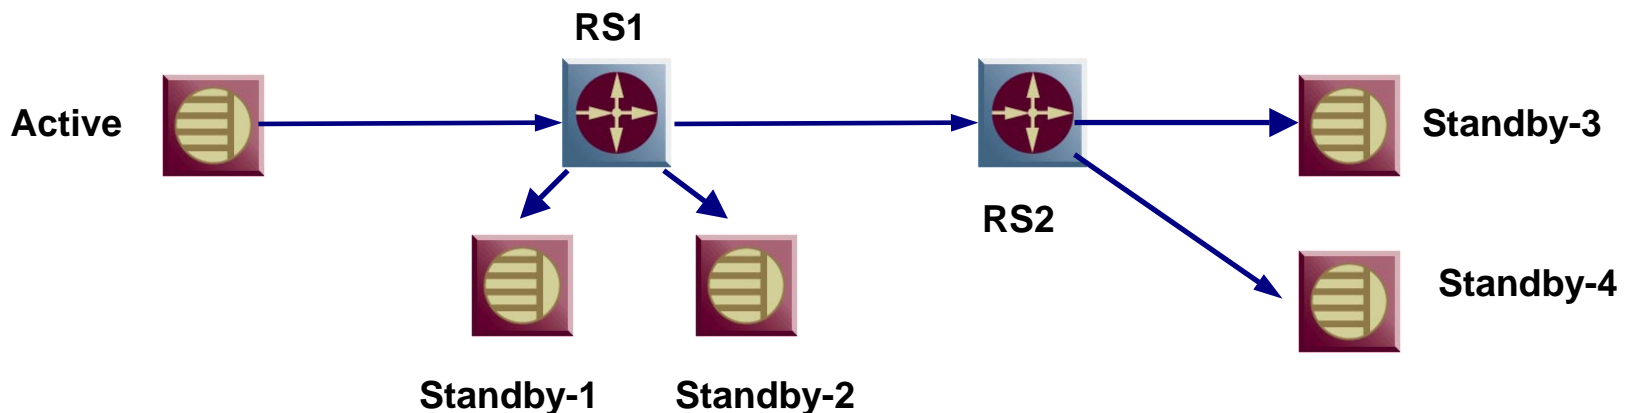




- Parallel DSI threads in RS attempt to apply transactions in the queue to the replicate DB in parallel
 - Commit order is still maintained – only the “body” of the transactions are applied in parallel
- Replication Server relies on a table in the replicate DB (rs_threads) to detect deadlocks and help sequence commits
 - Additional overhead on the replicate DB
 - Additional network roundtrips (between RS and replicate DB) for each transaction (group)
 - Limits the usability of Parallel DSI to ASE environments
- Deadlock detection for transactions distributed across Parallel DSI threads is now done internally in Replication Server
 - Commit sequencing is now completely internal to RS
 - Eliminates a network roundtrip for each transaction – reduced I/O
 - Less overhead for each transaction (group) – faster commit processing
- Removes the reliance on replicate DB
 - Eliminates a bottleneck (rs_threads) in the replicate DB
 - Allows non-ASE environments to take advantage of the Parallel DSI feature

- Warm Standby empty transaction processing
 - ASE generates “empty transactions” that contain only a “begin’ and “commit’
 - The overhead these empty transactions cause can be reduced with their removal, Warm Standby DSI will not be burdened by them anymore
 - RepAgent/ASE will also work to identify and filter some empty transactions before they are sent to Replication Server
- Thread contention improvements
 - Identify and reduce impact of serialized access to shared resources
 - SMP tuning in RS threads
- SPID Partitioning
 - Work with RepAgent teams to include SPID in LTL and enable Replication Server to partition transactions based on SPID
 - Use PDB SPID associated with the transaction to determine optimal DSI connections relative to incoming SPIDs

- Allows easy setup of multiple standby copies of a database
- Expands the table replication definition/subscription model to databases
- DDL can be replicated to all replicate sites, simply mark primary database with “sp_reptostandby”
- Configure environment without using rs_init, requires only Sybase Central Java Plugin for Replication Server




- Multiple database replicates can be used as read-only copies for:
 - Reporting
 - load distribution
 - decision support, etc.
- Any number of sites can be primary and replicate sites
 - Any site can be configure as a primary and replicate site but without applications connected for writes.
 - Once an event occurs(downtime, failure) such that the applications are redirected to the former replicated for use as the primary, transactions will flow from this site to the other replicates.
 - As a replicate site transactions are applied as the maintenance user
- No “switch” as provided by RS Warm Standby
 - Application must re-direct connections to new Primary Database, RepConnector is ideal solution
 - If all sites configured as primary sites no RS administration is required
 - Former Primary RS queues will drain and application can apply transactions to new primary

- With only two commands a full database with hundreds of tables can be replicated
 - “create replication definition ...”
 - “create subscription ...”
- All materialization methods are supported as with table replication
 - No materialization
 - Bulk with dump/load coordinated(database dump/load)
 - Bulk with manual coordinated(define/activate/validate)
- Supports replication of system procedures
- Supports Heterogeneous sources and targets
 - Non-ASE rep Agents cannot extract DDL from Heterogeneous sites

- MSA does NOT require table replication definitions
- Table replication definitions are defined to have greater precedence than database replication definitions
 - No table repdef – MSA sends all columns of all tables
 - One table repdef – MSA sends all columns and replicates the table using the table repdef
 - Multiple table repdefs on single table – MSA sends all columns and replicates the table using internal repdef which is a union of all repdefs
- Flexibility provided by MSA with table repdefs
 - Allow changes to column names and datatypes
 - Replicate tables of a particular owner
 - Minimal columns
 - Performance hints, ie. Primary key
 - Where clause available, not available on database subscriptions

- RSSD on an ASA(8.0.2) internally managed by RS(optionally), removes ASE requirement for hosting RS metadata
- Initialized/configured/started by rs_init in the background at Repserver installation time
- Simplify RS environment, reduces the number of RS dependencies and external components, reduces RS TCO




- Greatly benefits heterogeneous replication environments where ASE is not needed except for the RSSD
- Replication Server will manage the ASA without user involvement, RS will perform the following tasks:
 - Start
 - Stop
 - Configure
 - Perform backups
 - Recover from ASA failures including disk errors
 - ERSSD recovery systems are built into RS to work through most ASA failures
 - Requires no routine maintenance
- Existing environments can continue to use ASE as the RSSD
 - Minimize upgrade impact
 - New installations can choose between ASE and ASA

Open Switch Replication Coordination Module

- Coordinates activity between Open Switch and Replication Server at fail-over time
 - Automates client connection fail-over as well as Replication Direction Switch
 - Ensures all pending transactions in the queue are flushed to the standby before client connections are failed over
 - Coordinates 2 Open Switch servers to provide redundancy (if required)
- Open Switch configuration allows for multiple servers to be pooled for fail-over purposes
 - Read-only applications can be load balanced across a pool of servers
- Distributed with Open Switch 12.5

Warm Standby with Open Switch

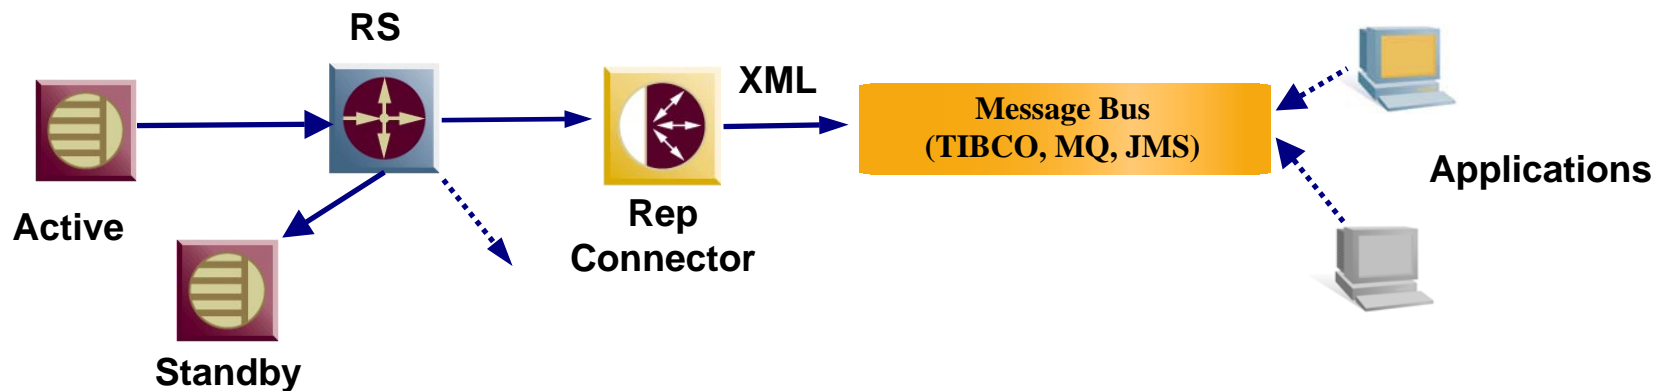


- Java based plug-in to Sybase Central to manage HA/Standby environments, support for MSA feature
- Integrated install/configurator for Standby environments – includes configuration of ASE, RS and Open Switch
- See Demo

- ASE datatype support for Date and Time
- Fully supported ANSI standard types
- Heterogeneous environment are simplified as datatype conversions are not required
- Updated HDS translations for new datatypes
- Mixed version support with automatic conversions between datatypes

- Native support for chinese character set gb18030-2000
- Chinese encoding standard published March 2000
- Required character set for business in China
- Multi-byte character set with 1-byte, 2-byte, and 4-byte forms.

- Integration with Message Bus, Workflow, Process Automation architectures
- Deliver database events to message bus or workflow environments for BPA/BPM
- Further interoperability with Integration/Workflow architectures in the future



- Introduction
- Replication Architecture
- RS 12.6 features
- **Demo**
- RS Futures Roadmap
- Q & A

- Introduction
- Replication Architecture
- RS 12.6 features
- Demo
- **RS Futures Roadmap**
- Q & A

Manageability and Ease of Use

- Reduce administration complexity
 - Automate or simplify recovery procedures
 - Enhance fault tolerance characteristics
 - Automated fault management
 - Simplify monitoring and performance management
 - Capacity Planning Tools
- Reduce Configuration Complexity
 - Simplify replication concepts
 - Modeling Tool support for replication topology and architecture design
 - Automated configuration deployment through modeling tool
 - Minimize deployment foot print – replication services well integrated into ASE

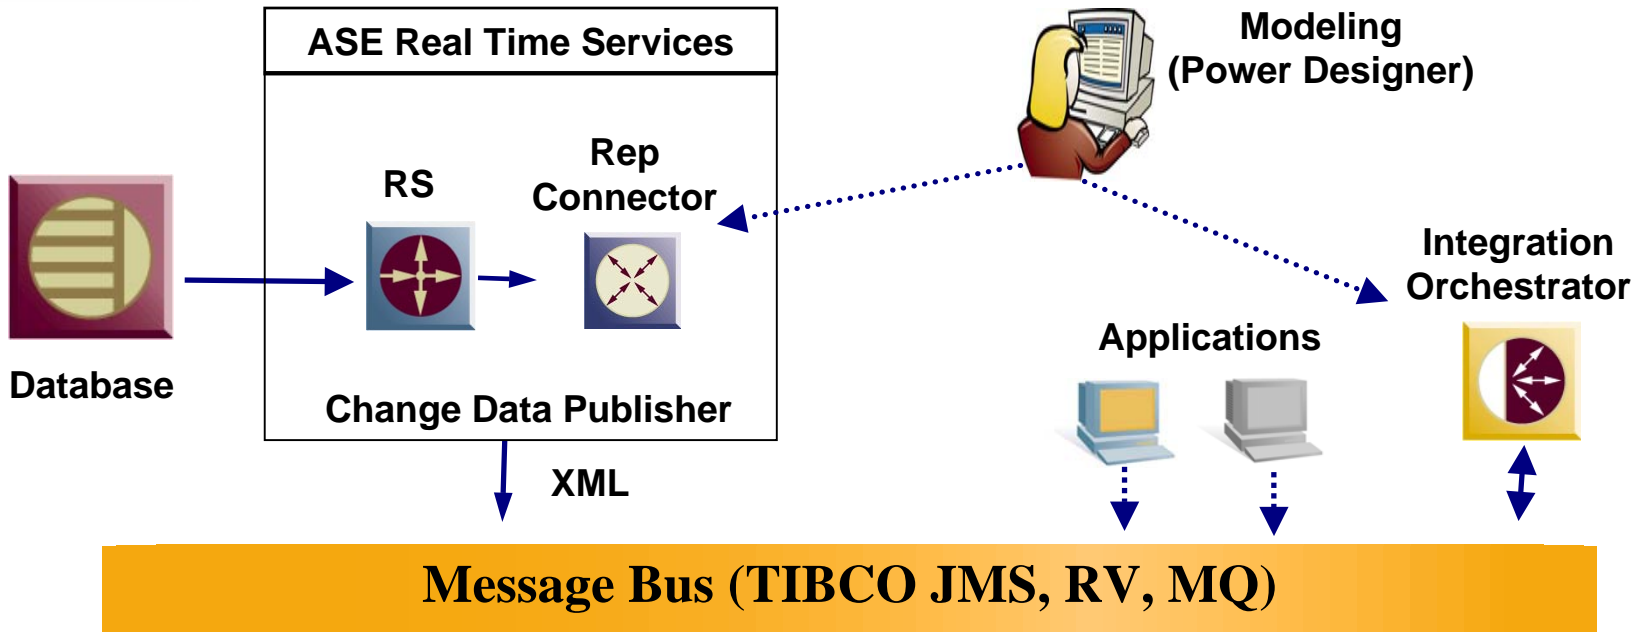
Integration, Heterogeneous Capabilities

- Log Based Replication Agent for Oracle
 - Minimizes application impact
 - Increased performance for replication from Oracle
 - To be followed by log based agent for Microsoft SQL Server
- Warm Standby for non-Sybase Databases
 - Allows easy DR setup for non-Sybase databases
- “Channel”-based Replication for Multi-master replication
 - Updates can come into any source and be distributed to other sites quite easily
 - Replaces the traditional publisher/subscriber concept with a replication Channel concept – any site can broadcast on a channel with any other receiving from the channel!

ASE Support

- Partition & Cluster Support
 - Replication from Partitions
 - Cluster availability for Replication Server
- XML Distribution
 - Distribute XML Content stored in ASE
 - Content-aware distribution
- External File System Support
 - Distribute content on the file system

Integration



- Rep Connector is already available
 - Integration with Message Bus, Workflow, Process Automation architectures
 - Deliver database events to message bus or workflow environments for BPA/BPM
- Modeling to simplify the configuration and setup in the future

A New Benefit for Sybase Administrators & Developers

- Forum for exchanging samples, tools, scripts, etc.
- New features enable community collaboration
 - Download samples created by Sybase or external users
 - Leverage contributions of others to extend Sybase products
 - Contribute code or start your own collaborative / open source project with input from other product experts
- **Any SDN member can participate**
 - Log in using your MySybase account via SDN
- www.codexchange.sybase.com
 - Or via SDN at www.sybase.com/developer

- Introduction
- Replication Architecture
- RS 12.6 features
- Demo
- RS Futures Roadmap
- Q & A