solidDB, extending IBM Data Server Portfolio for extreme speed

Paola Lubet
Senior Vice president ,
Solid, an IBM Company
plubet@us.ibm.com

Agenda

- **Solid, Company Background and Target Market**
- **Solid, Technology Overview and Competition**
- **Solid and IBM, the Road Ahead**

solid.
an IBM Company

Solid Company Background and Target Market

About Solid...

- Founded in 1992 in Helsinki, Finland
- Headquartered in Cupertino, CA
- Worldwide operations

Recognized leader for in-memory relational database technology

- Hundreds of customers worldwide including market leaders Airbus, Alcatel, Cisco, Motorola, Nokia-Siemens, NEC, Nortel
- Over 3 Million deployments of Solid software
- Strong presence in Communications
- Proven technology

What does it offer?

Extends IBM Data Server Portfolio with in-memory database technology for Extreme Speed

- **Solid is a leading provider of relational in-memory database solutions**
- **IBM solidDB™, extends IBM's DB2 and IDS data server products with:**
  - micro-second response time and
  - high transaction throughput

  …using the familiar SQL language
- **IBM solidDB is available today as a standalone product**
- **In 2008 also available as SQL Database Cache able to work seamlessly with DB2 and IDS**

IBM solidDB: The Value

Extreme Speed

- Latency: Measured in microseconds
- Volume: Tens of thousands of concurrent transactions

Extreme Availability

- Uninterrupted data access - 99.999% availability
- Failover to hot standby in 50-500 ms

Low Cost

- Easy to deploy and administer within applications
- Small footprint – 10MB or less

Solid provides proven technology, providing over 3,000,000 deployments worldwide with real-time data access

Solid: 100s of Customers in a Variety of Industries

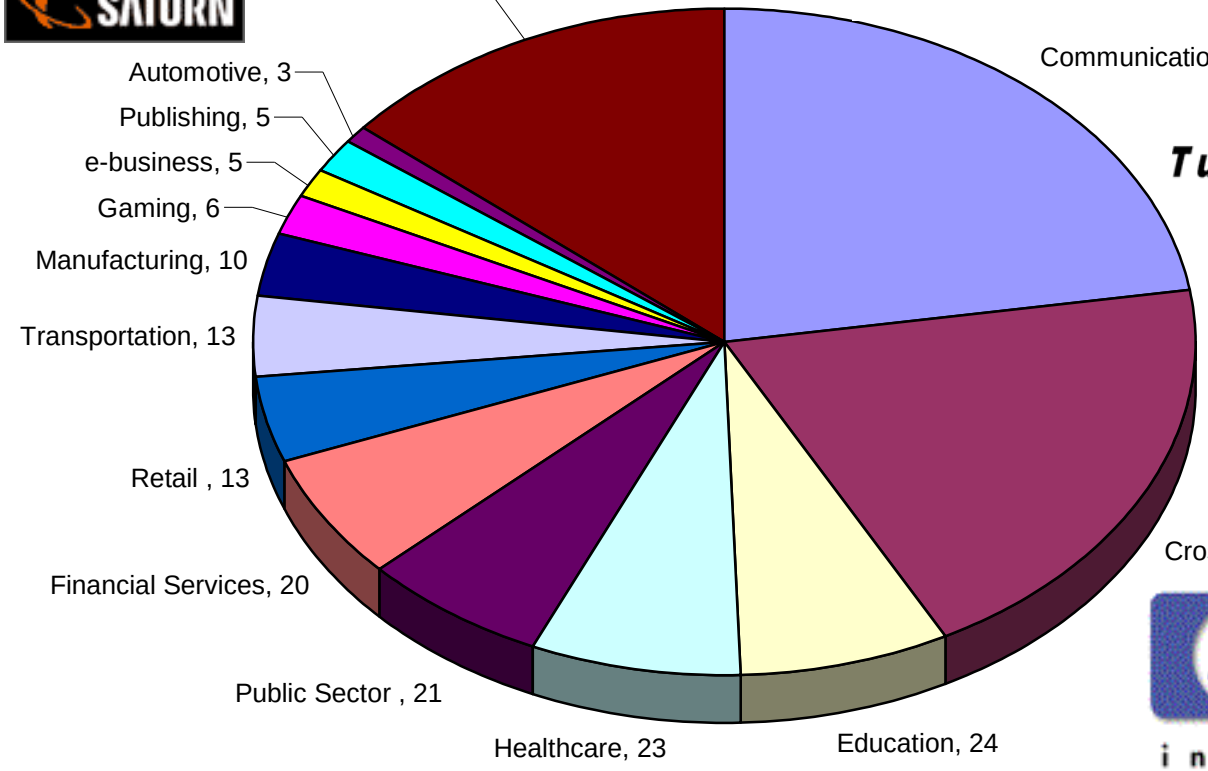

Solid Delivers Business Value To ISVs and NEPs

- **BridgePort Networks** allows Voice-over-IP subscribers to use their regular cell phone numbers for internet-based calls. It leverages solidDB to instantly register and authenticate internet users within their converged network.

BridgePort Networks
Where Mobile meets VoIP

- **Cisco's Wireless Location Appliance** is a location solution that relies on solidDB to simultaneously tracks thousands of devices for location trending, rapid problem resolution, and radio frequency capacity management.

CISCO

- **One of top 10 NEPs** uses solidDB to route mobile phone users requests to the appropriate service provider within 20 microseconds.

- **TruePosition** uses solidDB to implement 911 location services so that mobile phone users can be immediately located and promptly assisted.

TruePosition
Global Wireless Location Solutions.

Solid Technology Addresses Growing Market Trends

Increasing Demand For Immediate Response Time

- There is a growing market for applications which require extreme speed; **at least 10x that of a traditional relational database**

Increased Demand For Extreme Transaction Throughput

- “By 2011, more than 20% of new TP application projects will require support for more than 10,000 concurrent accesses or more than 500 transactions per second, up from less than 10% in 2006 (*)”

… Around-the-Clock and at a Low Cost

- Businesses need to keep data available during planned and unplanned outages, 24 x 7 x 365
- Businesses seek to accelerate deployments and reduce administration costs for new and existing applications.

(*) Source: Gartner 2007

Together IBM solidDB and DB2/IDS

Telecom Service Connect

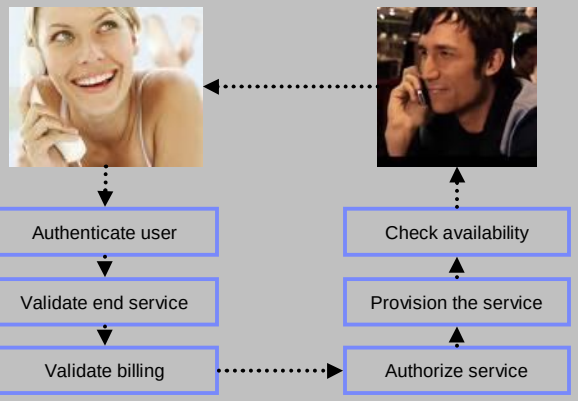

End-to-end cycle must be <1 second

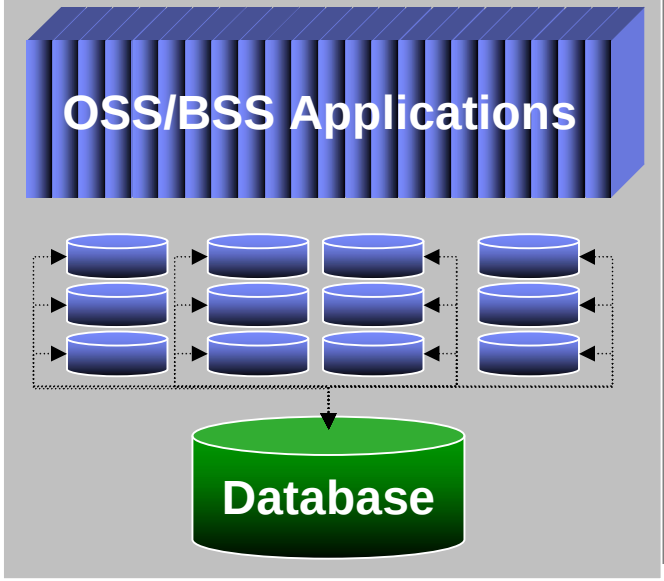

Common Requirements

End Users Applications

Both scenario require:

- immediate response time
- 99.999% data availability

Applications

- Dozens of customized SQL applications
- Speed/response time IS the competitive advantage

In-Memory Database Layer

- micro-second response time
- Scalability
- Highly Availability

Master Database*

- Transactional data
- Customer data

Financial Services Trading

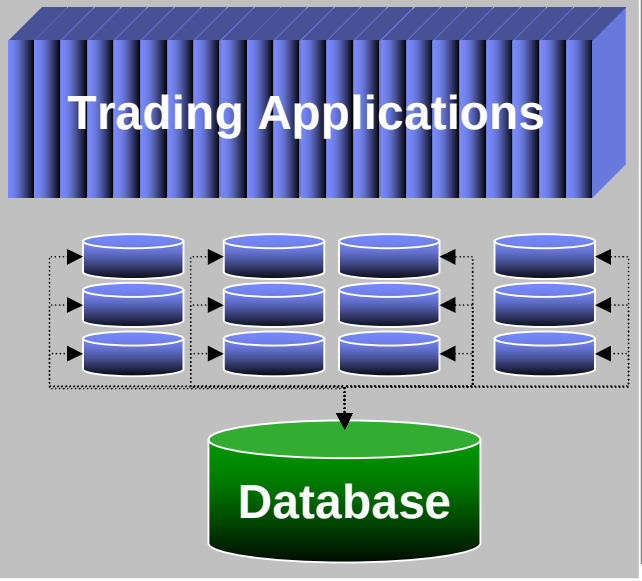

Solid and IBM at Work: A Proof of Concept

- **Application validates and updates credentials of millions of mobile phone users**
  - solidDB 6: Real-time user validation
  - DB2: User and service provisioning
  - Changes propagate from DB2 to solidDB
- **Delivering Extreme transaction Throughput ***
  - IBM DB2 + (4 instances) solidDB 6
    - **61,390** transactions per second
- **Achieving response time average of less than 1 millisecond for read operations**

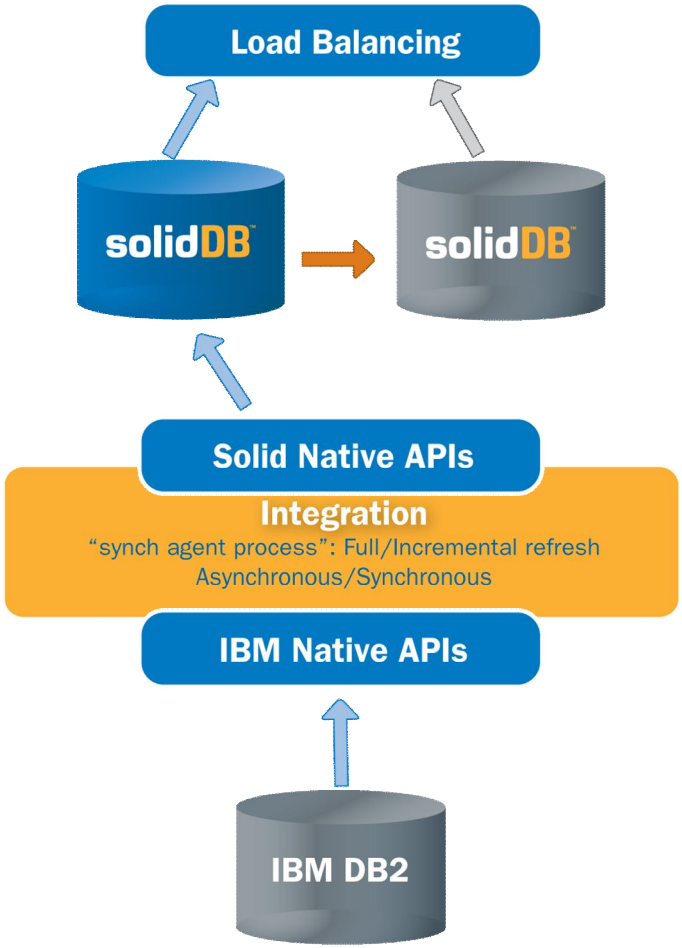

* TelecomOne Benchmark

solid.
an IBM Company

IBM solidDB 6 Built for Extreme Speed

Technical Overview

IBM solidDB: Key Capabilities

- **in-memory relational database with on-disk engine**
  - solidDB keeps data in main memory tables OR disk based tables in the same database
  - Applications can take advantage of its capability through standard ODBC, JDBC, SQL interfaces (SQL92)
  - C based runtime engine
  - Indices are optimized for in-memory searches and cache sensitivity
  - Query optimization takes into account cost functions for RAM operations, which are different than with disk access
- **Instant failover**
  - solidDB maintains two copies of the data synchronized at all times
  - In case of system failure, applications can recover access to solidDB well under 1 second without loss of data
- **Embeddable**
  - solidDB can be deployed in a client/server configuration, or as a linked-library embedded in the application
  - Linked library mode collocates application with engine thread giving lowest possible latency

solidDB In-memory Engine Supports Three Table Types To Optimize Durability And Performance

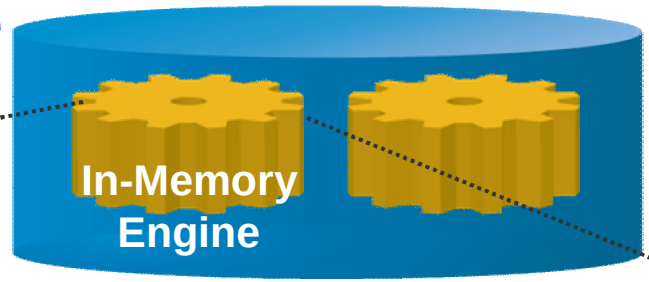

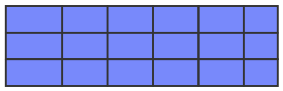

Persistent in-memory tables

- Logged, checkpointed
- Response times more predicable – checkpointing does not block access
- Secondary indices maintained only in-memory
- All indices equally fast

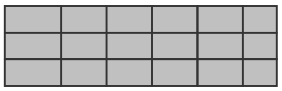

Transient in-memory tables

- Last until database server shuts down
- All users see them
- No logging – faster than persistent tables

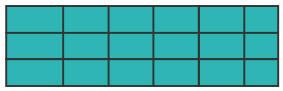

Temporary in-memory tables

- Last until connection closes
- Private to the connection
- No concurrency, no locking – faster than transient tables

solidDB In-memory Performance

- OS Version: RHEL 4.0 + SP 4
- TM1 Benchmark
- RHEL 4.0 + SP4
- 2 socket quad core 2.66 GHz Xeon
- 24 Gb RAM

- 38,000 txn/sec
- GET_SUBSCRIPER <20 us (single row read)
- UPDATE_SUBSCRIBER updates row in two tables

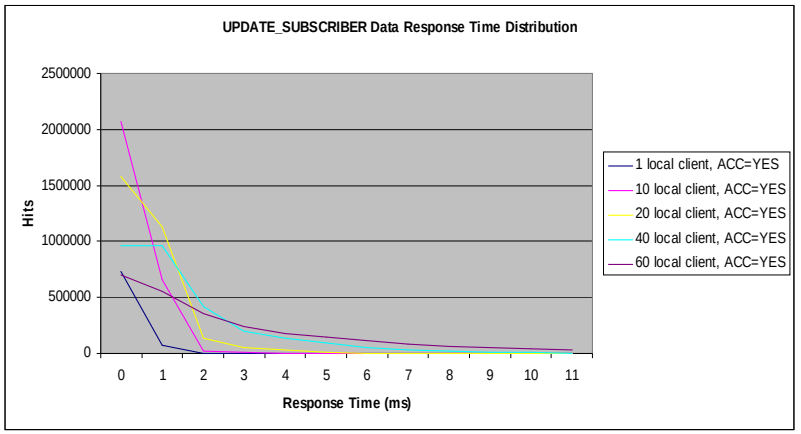

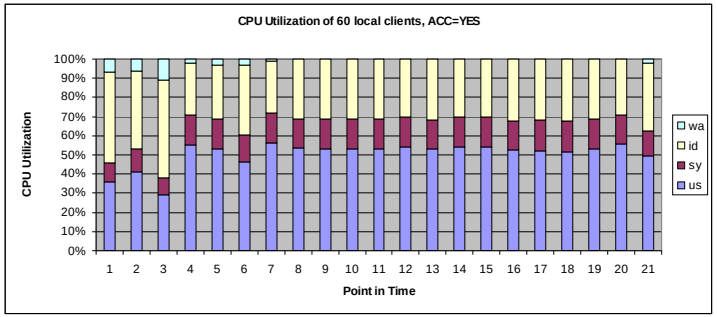

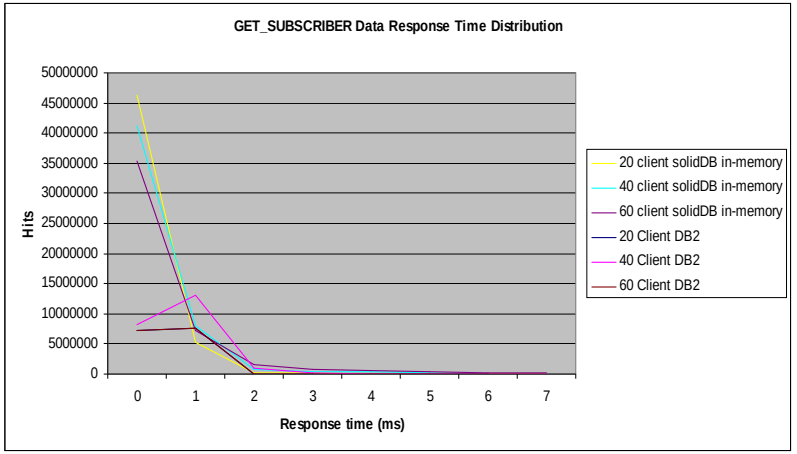

solidDB High Availability provides up to 99.999% database up-time delivering sub-second failover time

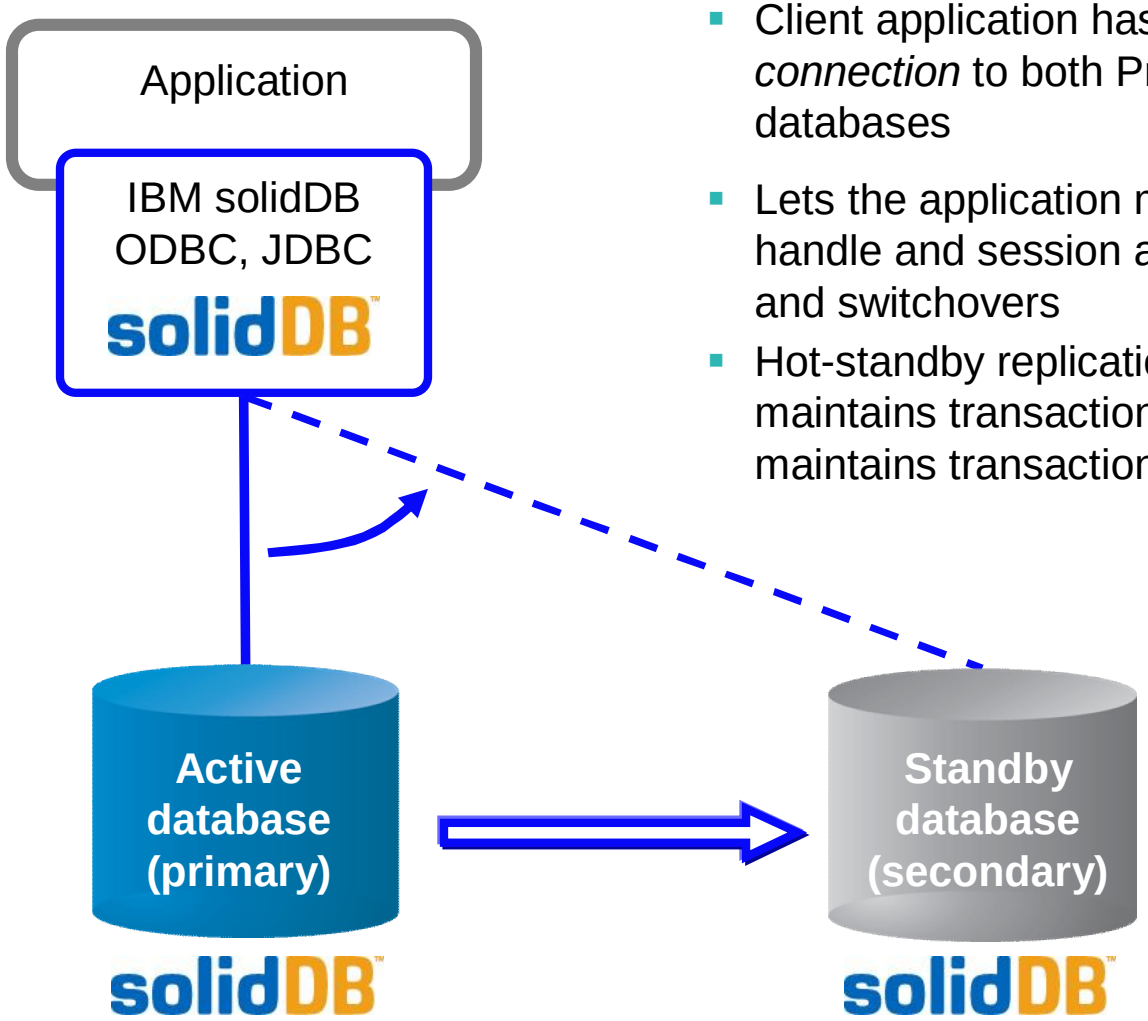

- Client application has *only one logical connection* to both Primary and Secondary databases
- Lets the application maintain connection handle and session attributes over failovers and switchovers
- Hot-standby replication checkpoints the state, maintains transactional consistency, maintains transaction durability

solidDB High Availability Configuration Achieves Right Trade-off Between Performance, MTTR, and Durability

Chose the Appropriate Durability

- Flexible synchronous and asynchronous logging at primary and secondary allow appropriate durability tuning

Choose the Appropriate Transaction Replication Protocol

- Transaction can be committed at Primary, then transmitted to Secondary, or …
- … transaction can be acknowledged/executed at Secondary before it gets committed at Primary
- Configurable globally (server), per session, or per transaction

Leveraging solidDB HA Architecture for Extreme Speed: Flexible Logging Architecture

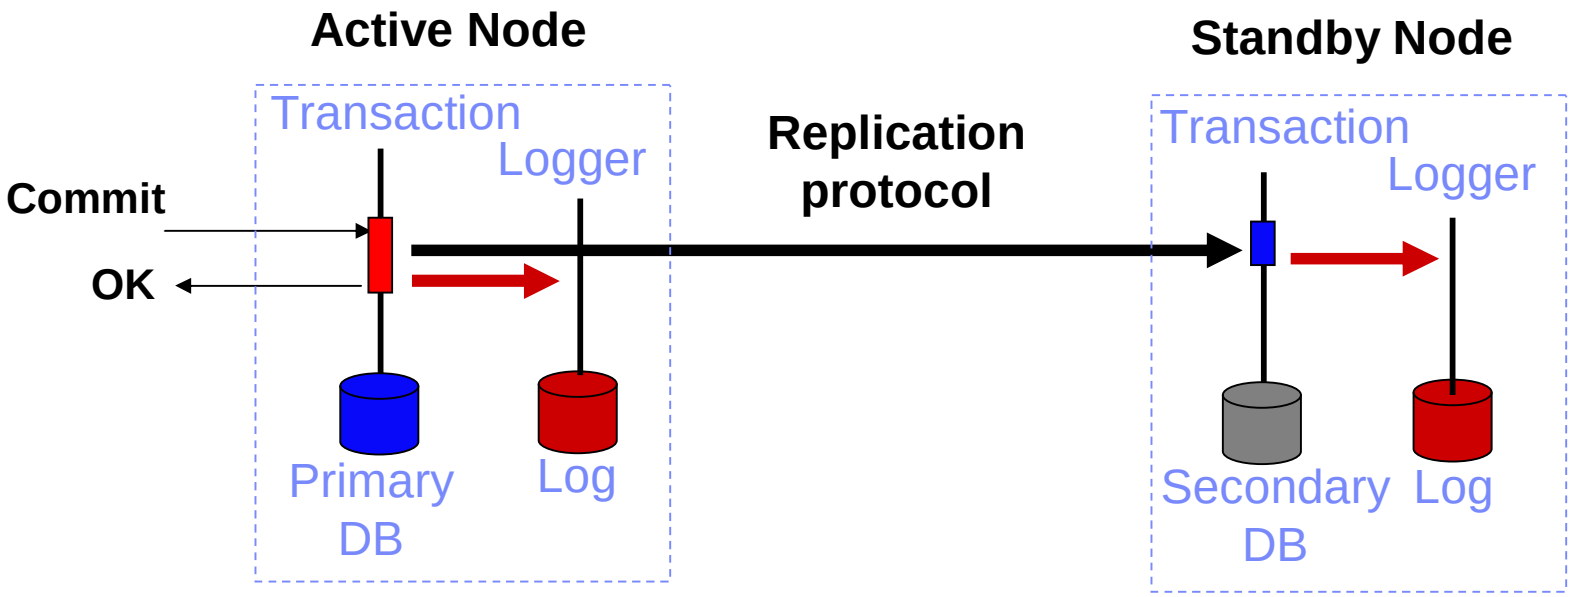

There are two ways to maintain transaction durability:
- **disk-based durability (transaction logging)**
- **network-based durability (transactional replication)**

Impact of Asynchronous Logging on Performance

75-150% throughput increase using asynchronous disk-based logging in the active node

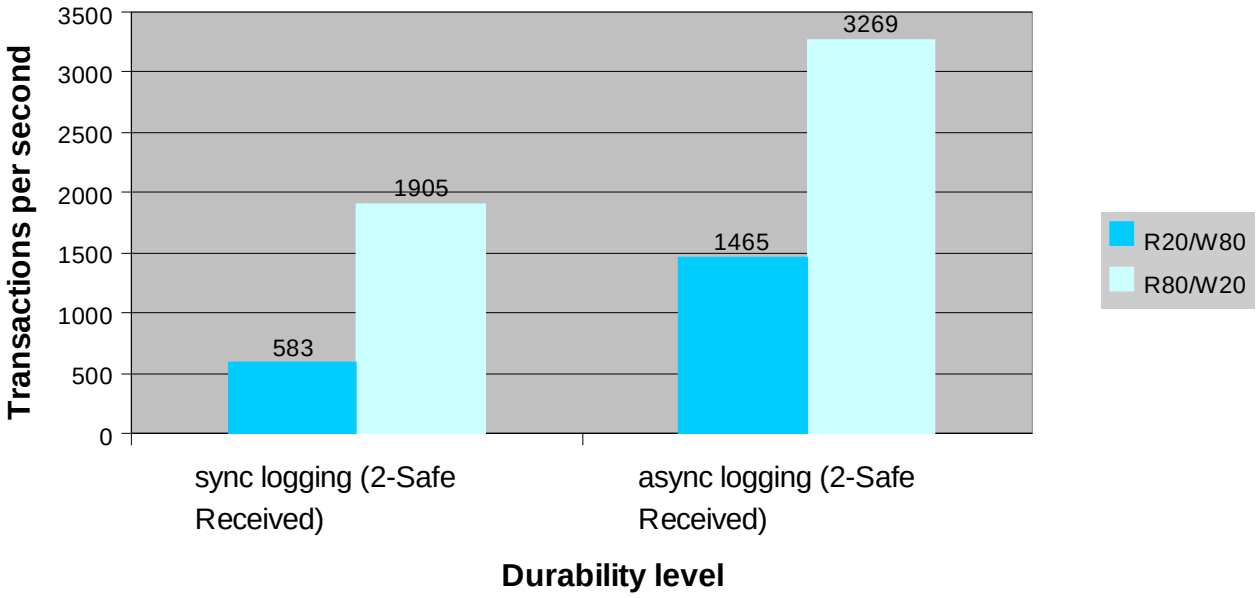

Transaction committed at Primary and then transmitted to Secondary: *1-Safe Replication*:

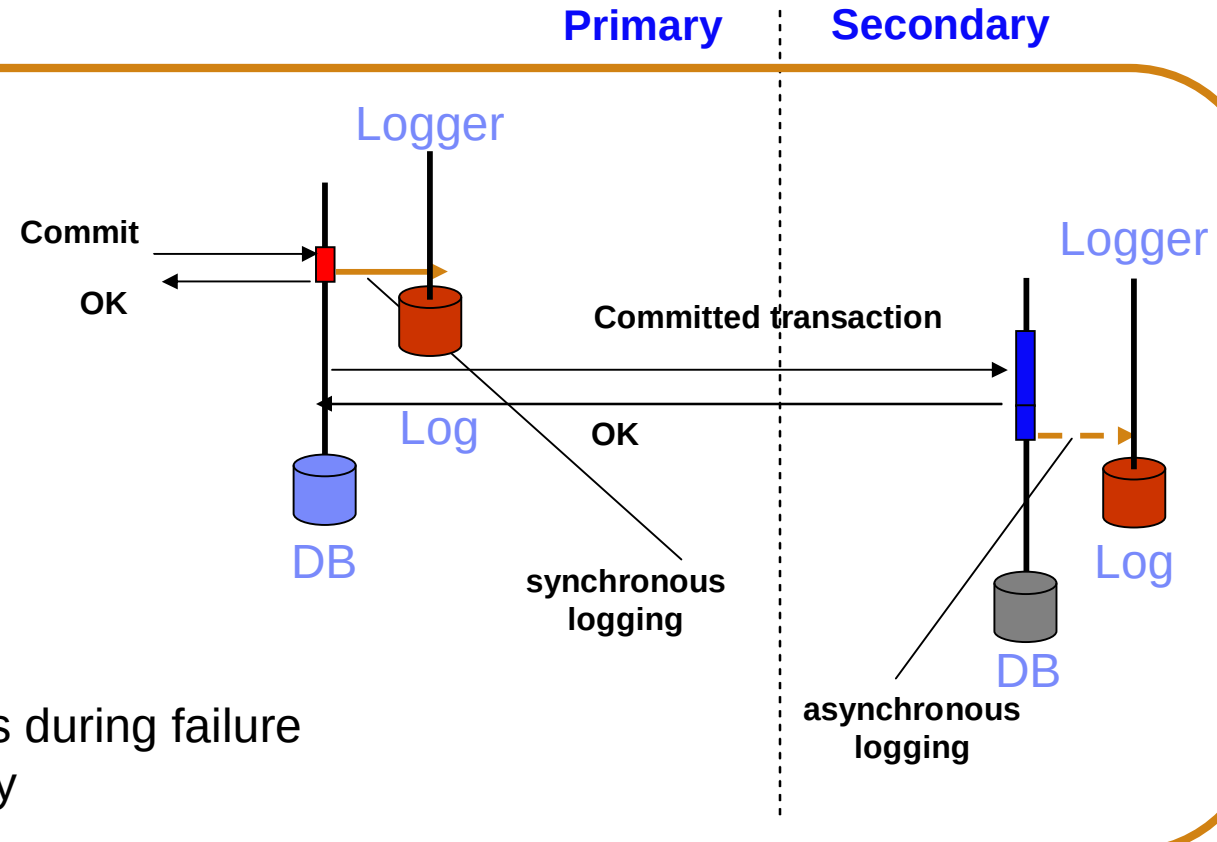

- Very Low latency
- Can lose transactions during failure
- Used for site recovery

Transaction committed after it has been acknowledged by Secondary: *2-Safe Replication with different levels of durability*

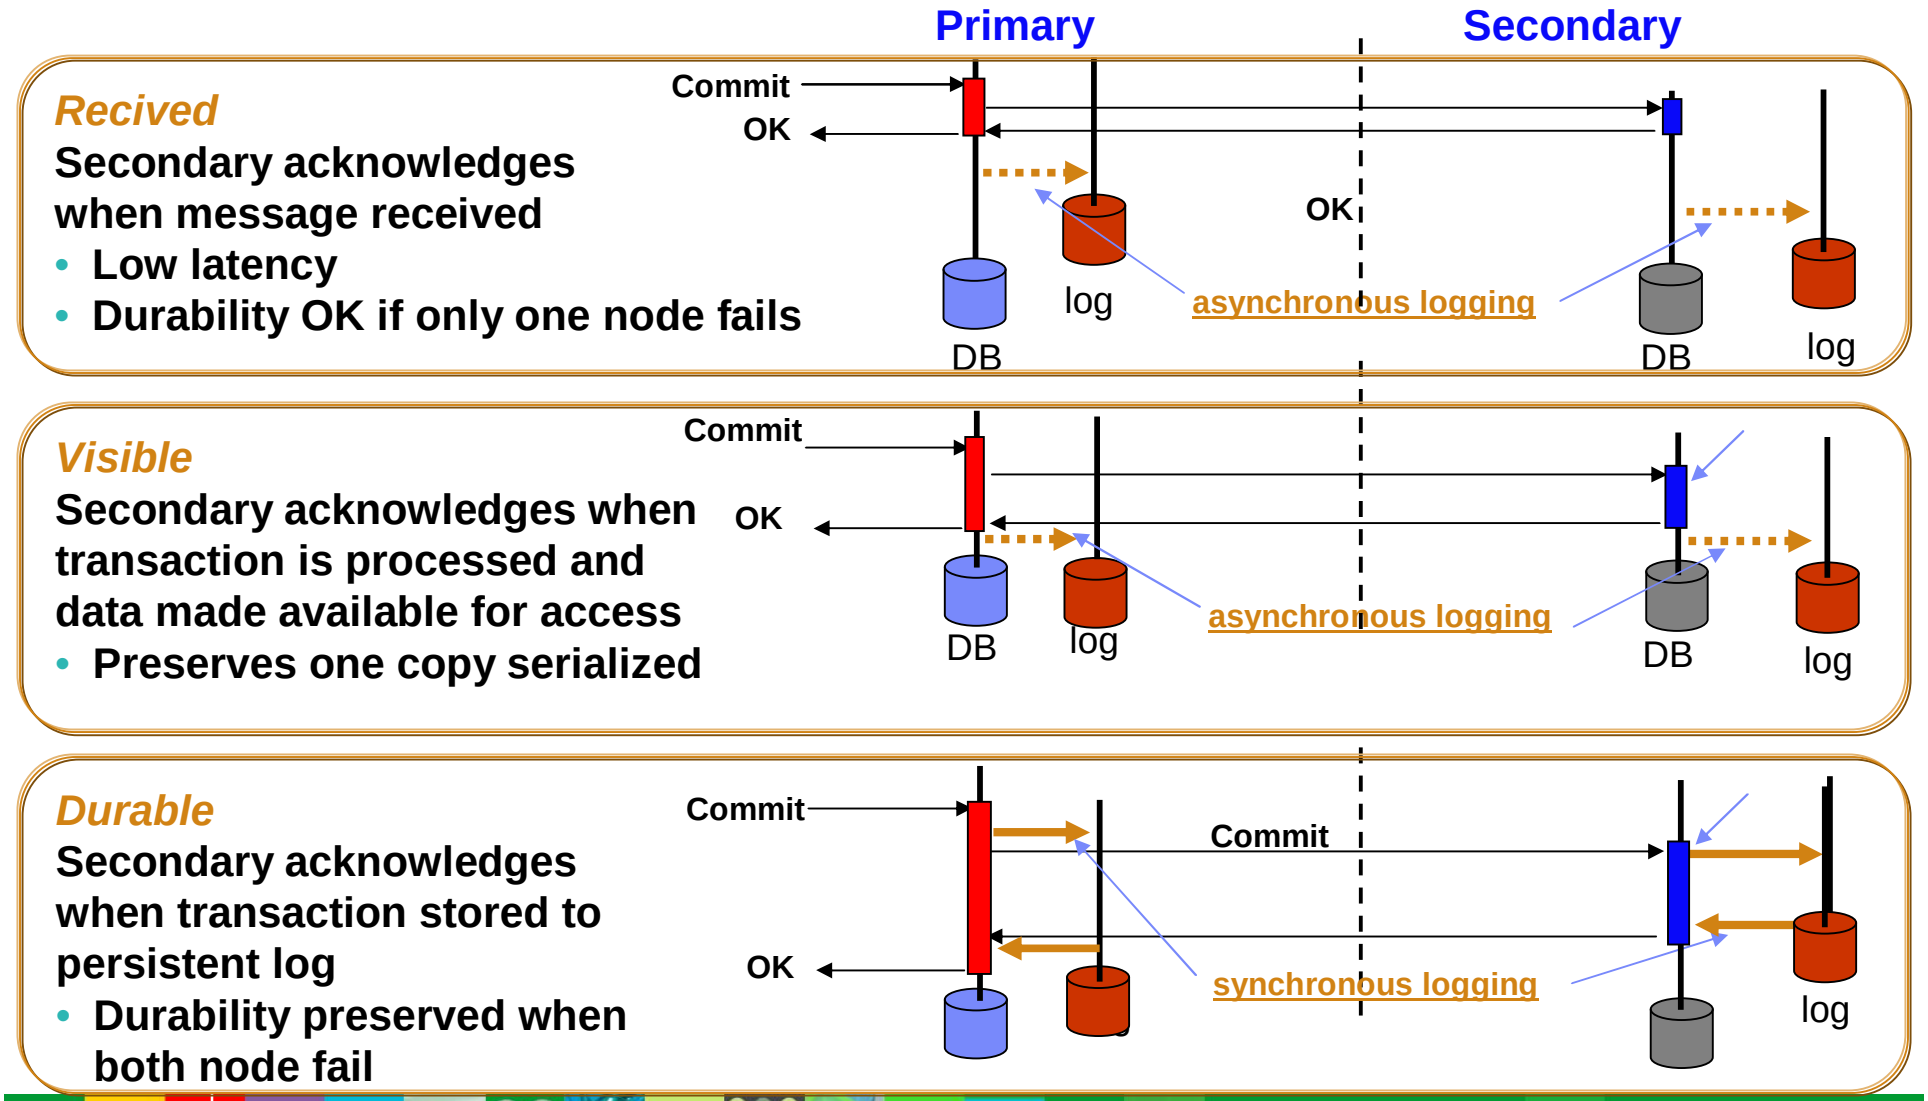

Comparing different replication protocols and durability levels in hot-standby configuration

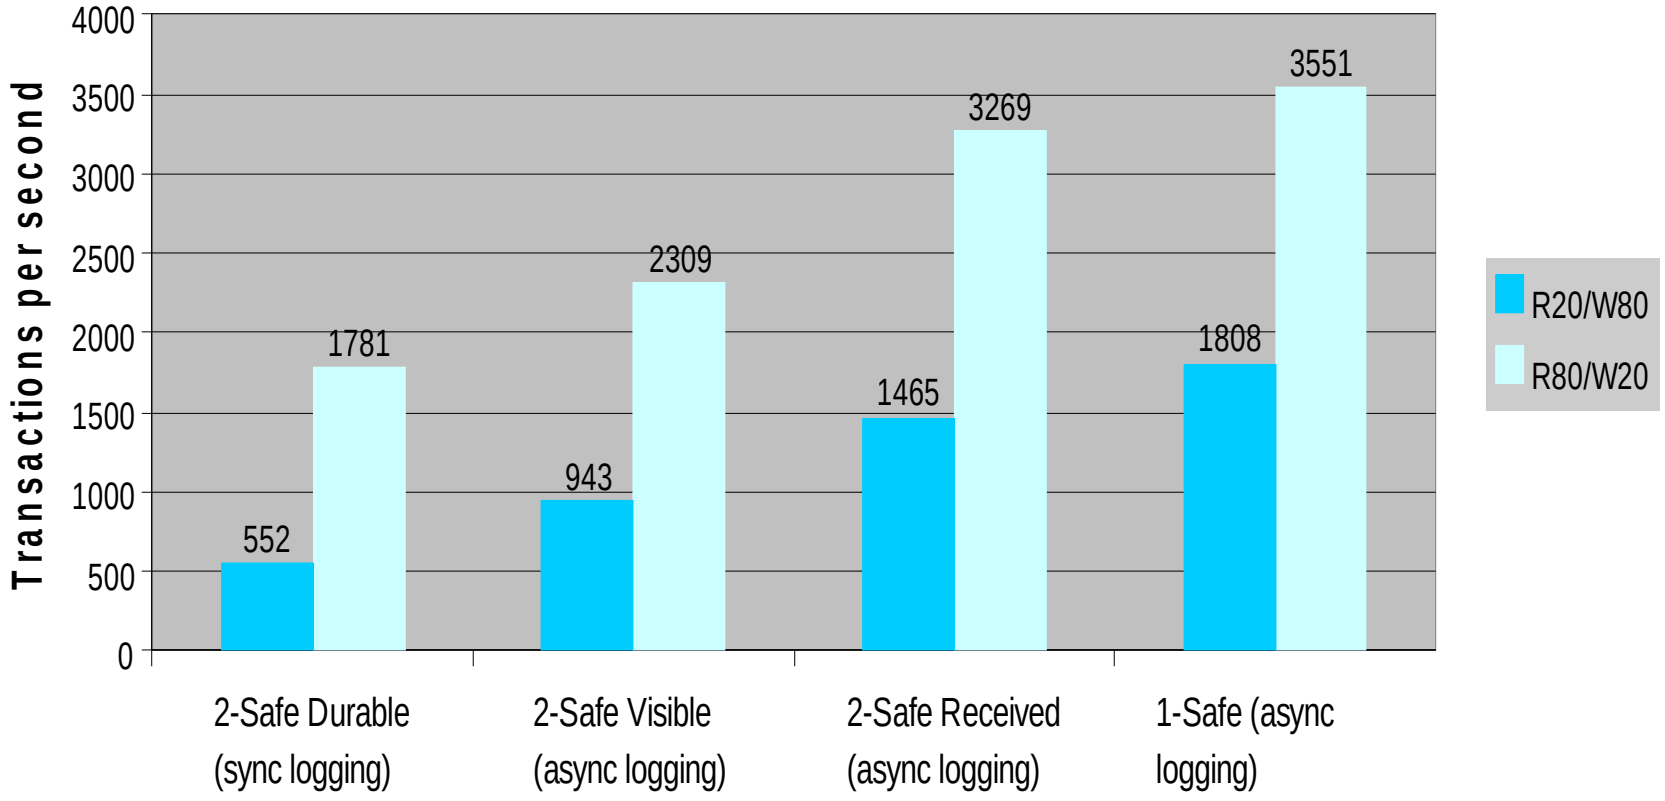

solidDB Load Balancing in HA Configurations: Improves Throughput Scaling Read Operations

- **Read-only transactions are distributed between primary and secondary hot-standby nodes**
- **All write transactions are directed to the primary node.**

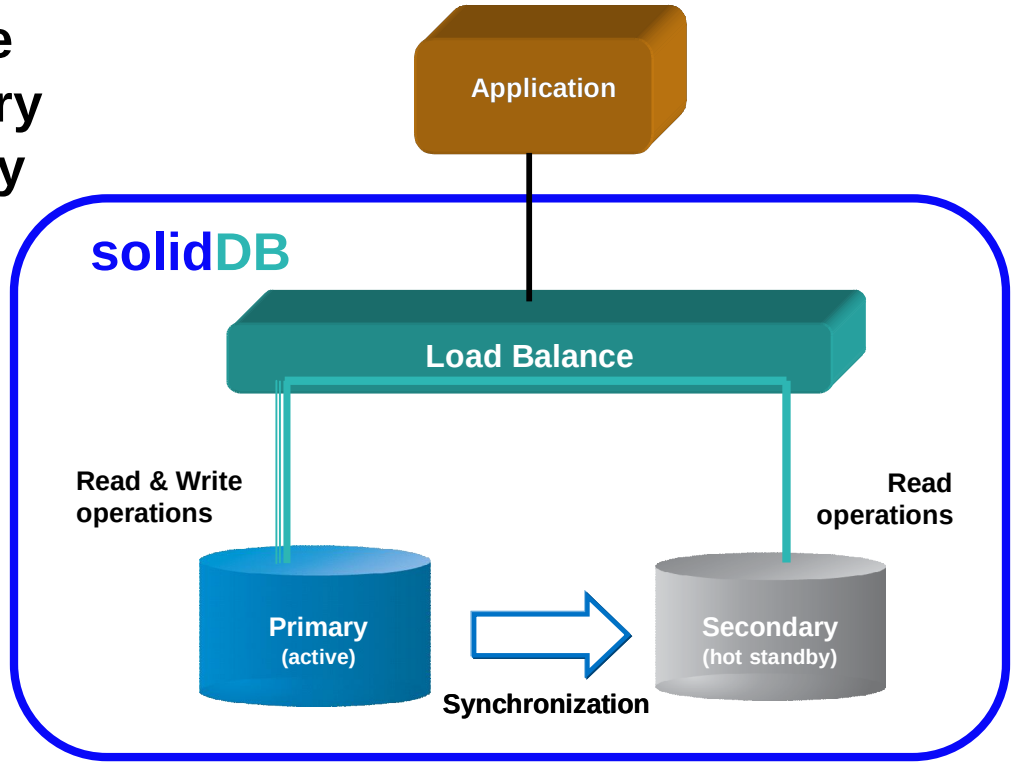

Load Balancing Impact on Transaction Throughput

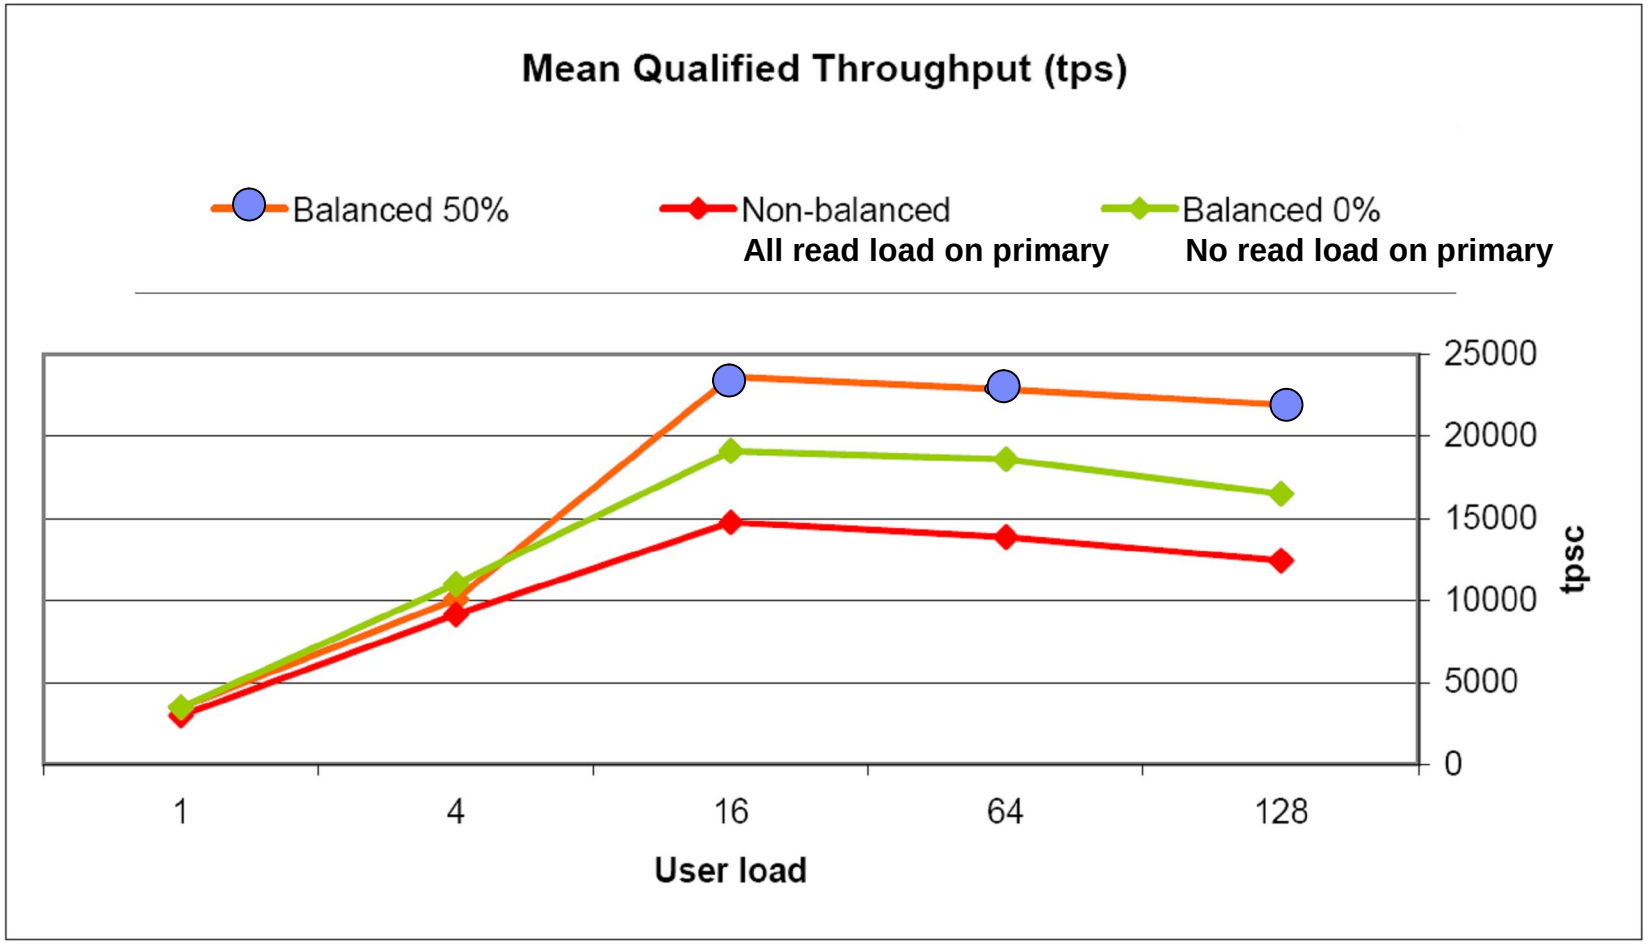

Telecom One Benchmark: 80/20 read/write workload

Additional Deployment Architectures for Extreme Speed

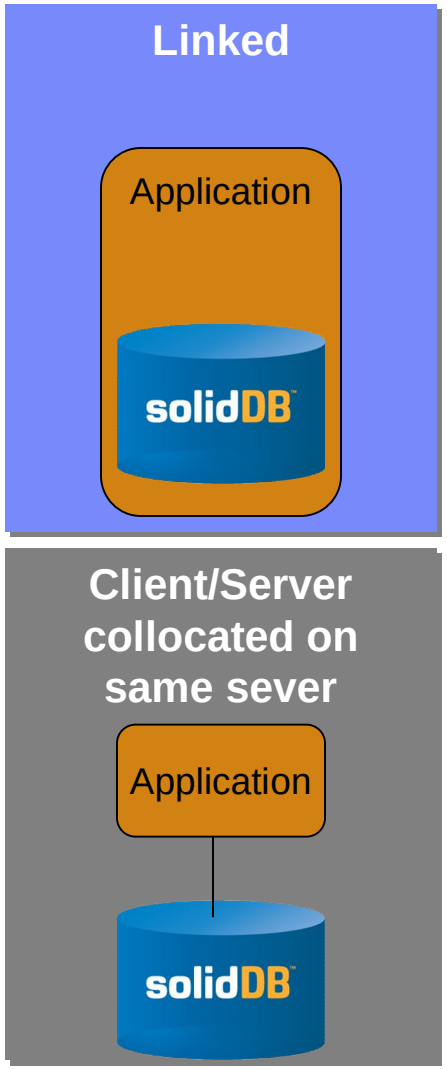

~40% throughput increase using solidDB as a linked library

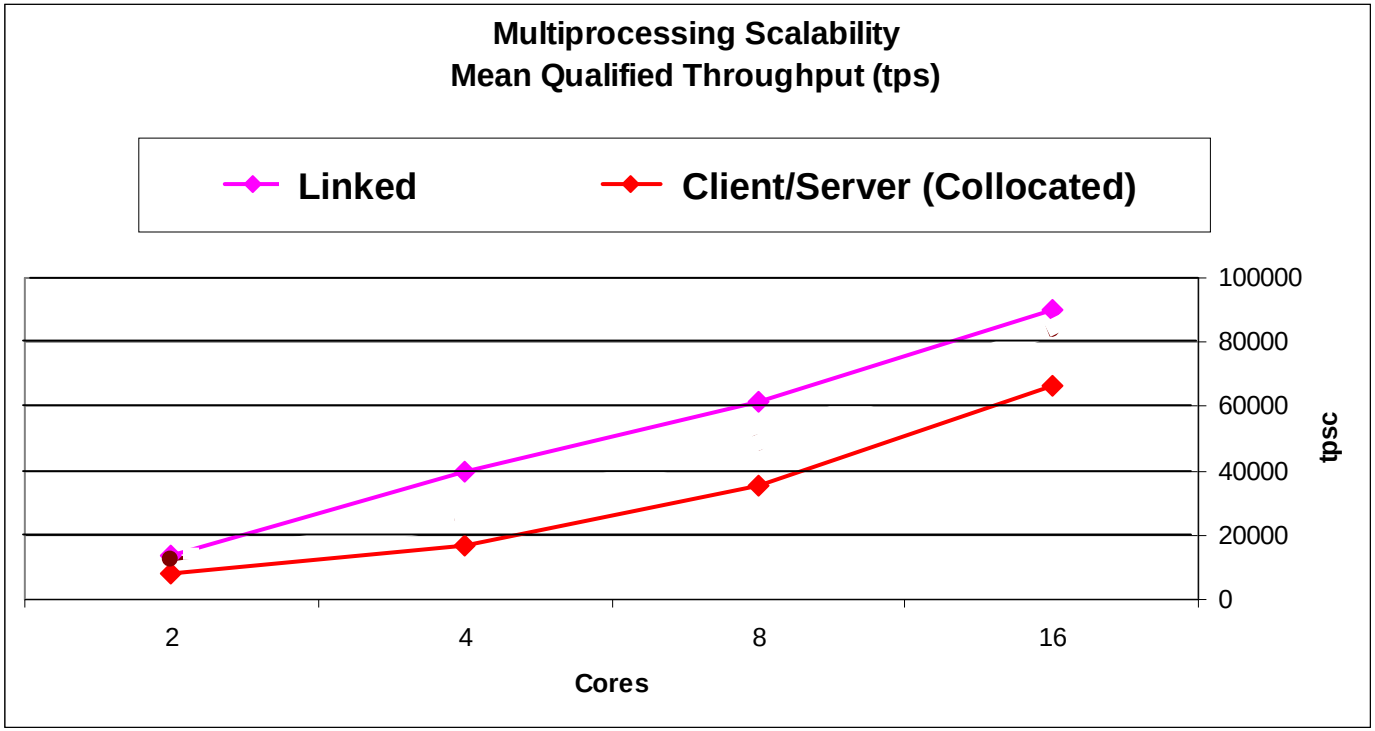

solidDB Scales on Multi-core/Multi-processor architectures for Extreme Speed

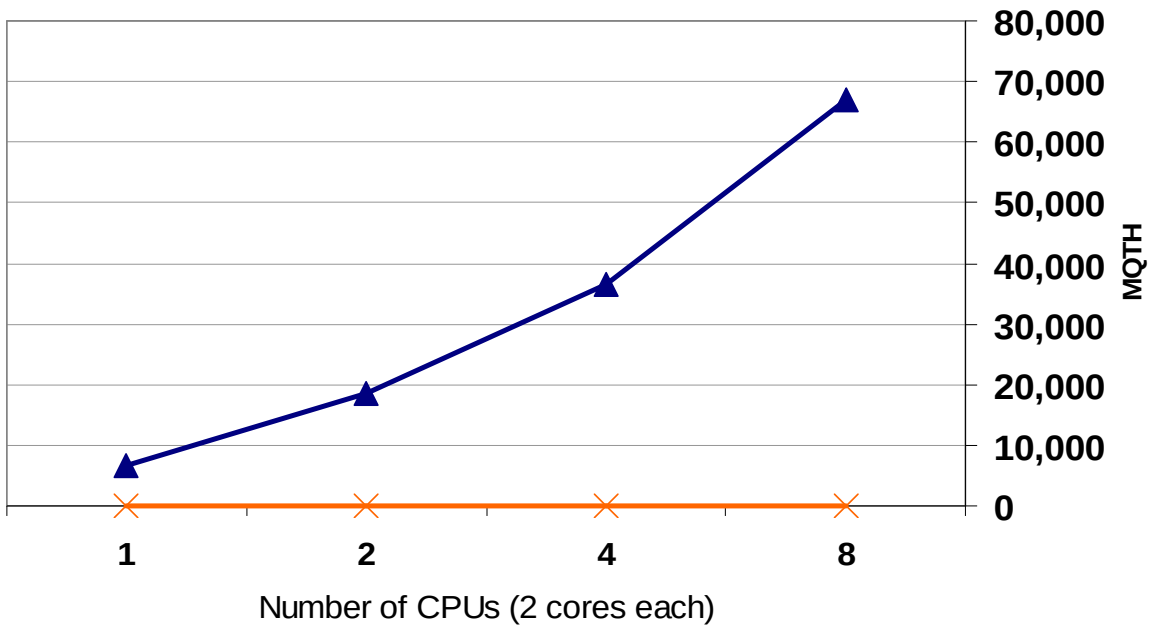

- **Telecom One Benchmark**
- **Simulates 80 / 20 reads/writes workload**
- **100.000 subscribers**
- **32 concurrent clients on the same server as solidDB 6**
- **8CPU dual-core AMD on SuSe 10**

Additional Functionalities: solidDB provides built-in advanced replication

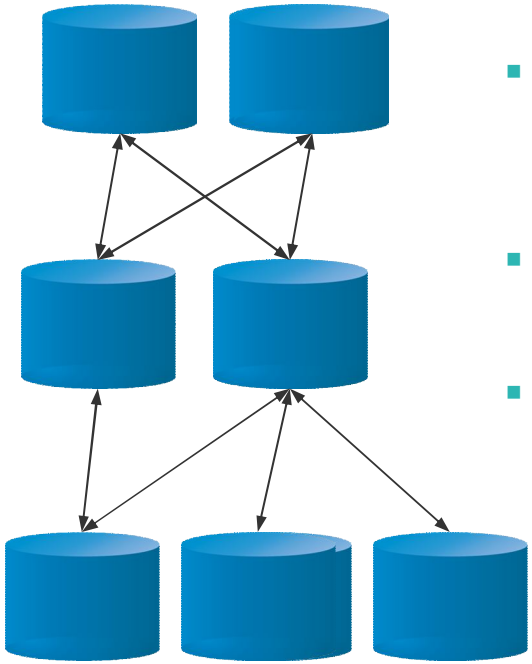

- **Bi-directional, publish and subscribe replication capabilities between solidDB databases with data consistency**
- **Originally designed for occasionally connected databases, improved for ATCA/blade-centric computing with fast, reliable server connections**
- **Built into the solidDB process, no extra component/process needed to configure or to manage**
- **Concepts:**
  - Master and replica databases
  - Publications and subscriptions for transferring new and changed data from the master database to the replica database
  - Can propagate changes from a replica database to the master database
  - Asynchronous store and forward messaging for implementing safe and reliable communication between the master and a replica

Who is the competition?

Oracle TimesTen

IBM solidDB vs. Oracle TimesTen

solidDB wins on:

- more transactions per second in networked client/server configurations, as demonstrated by the Telecom One Benchmark.
- never loses transactions solidDB can be configured in such a way that no transactions can be lost in case a database node becomes unavailable unlike Oracle TimesTen HA configuration, where transactions can be lost.
- superior database programmability with its support for stored procedures and triggers, which is a capability Oracle TimesTen lacks.
- for applications where database logging can be disabled, solidDB provides superior performance, concurrency, high availability and replication

Contrary to TimesTen, solidDB doesn't not yet:

- allow multiple processes on the same server as the database must be linked with the database for shorter response times
- integrate with back-end database servers. This will be available starting in June 2008

solid.
an IBM® Company

Coming up: solidDB DB2/IDS Integration

solid.

What is Available Today and Product and Roadmap

Today – Embedded In-Memory Database

Embedded into Communications Devices

Today – In-Memory Database for Low Latency Data Access

Financial Services and Telecom

Coming up - SQL Database Cache

Feature of IBM data servers

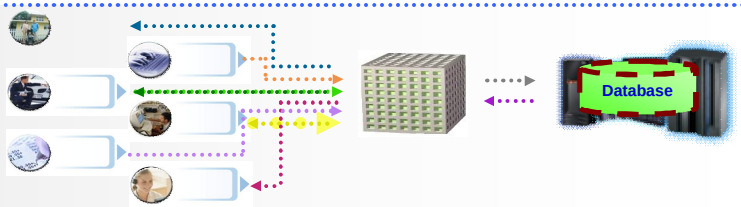

Low Latency Data Access

- **Query response time in sub-millisecond (<20 us) range**
- **Continuous Availability characteristics**
- **Latency sensitive schema will be in memory**
  - Database could include on disk schema for queries that are not latency sensitive
- **Data will be persisted to backend database (DB2 or IDS)**
  - Can persist direct transactional load or aggregation or transformation of that load
- **Improve replication setup and configuration through Data Studio configuration**

Latency Schema: Cache Owned

- **All Insert, Update, Delete operations go directly against front end cache**
- **Low latency reads also go against front end cache**
- **Operations asynchronously written to backend database**
- **Persistence of data cannot impact latency or availability**
  - Asynchronously push data to outboard process
  - Persist data through JDBC or ODBC connection
  - Plug-in point for transformation/aggregation logic
- **Reporting and other non-latency sensitive reads are executed against backend**
- **Backend database provides disaster recovery**

Built-in, dynamically queriable, and programmatically configurable high availability state machine

PRIMARY ACTIVE	**Secondary server in SECONDARY ACTIVE state**
PRIMARY ALONE	**Servers not connected**
PRIMARY UNCERTAIN	**Abnormal state, transactions pending (not committed until state changes)**
SECONDARY ACTIVE	**Primary server in PRIMARY ACTIVE state**
SECONDARY ALONE	**Servers not connected**
STANDALONE	**Standalone operation**
OFFLINE	**Waiting for a netcopy (spare)**

solidDB Carrier-Grade High Availability state diagram