Information Management

IBM

# DB2 Connect Application Development Update ODBC / CLI and .NET

**Brent Gross - gross@ca.ibm.com**

November 13, 2011

# Agenda

- ADO.NET Providers
- DB2 .NET Provider
- .NET Features
- ODBC / CLI Update
- pureQuery Client Optimization
- Performance Hints and Tips
- Function Rollout Overview

# ADO .NET Providers

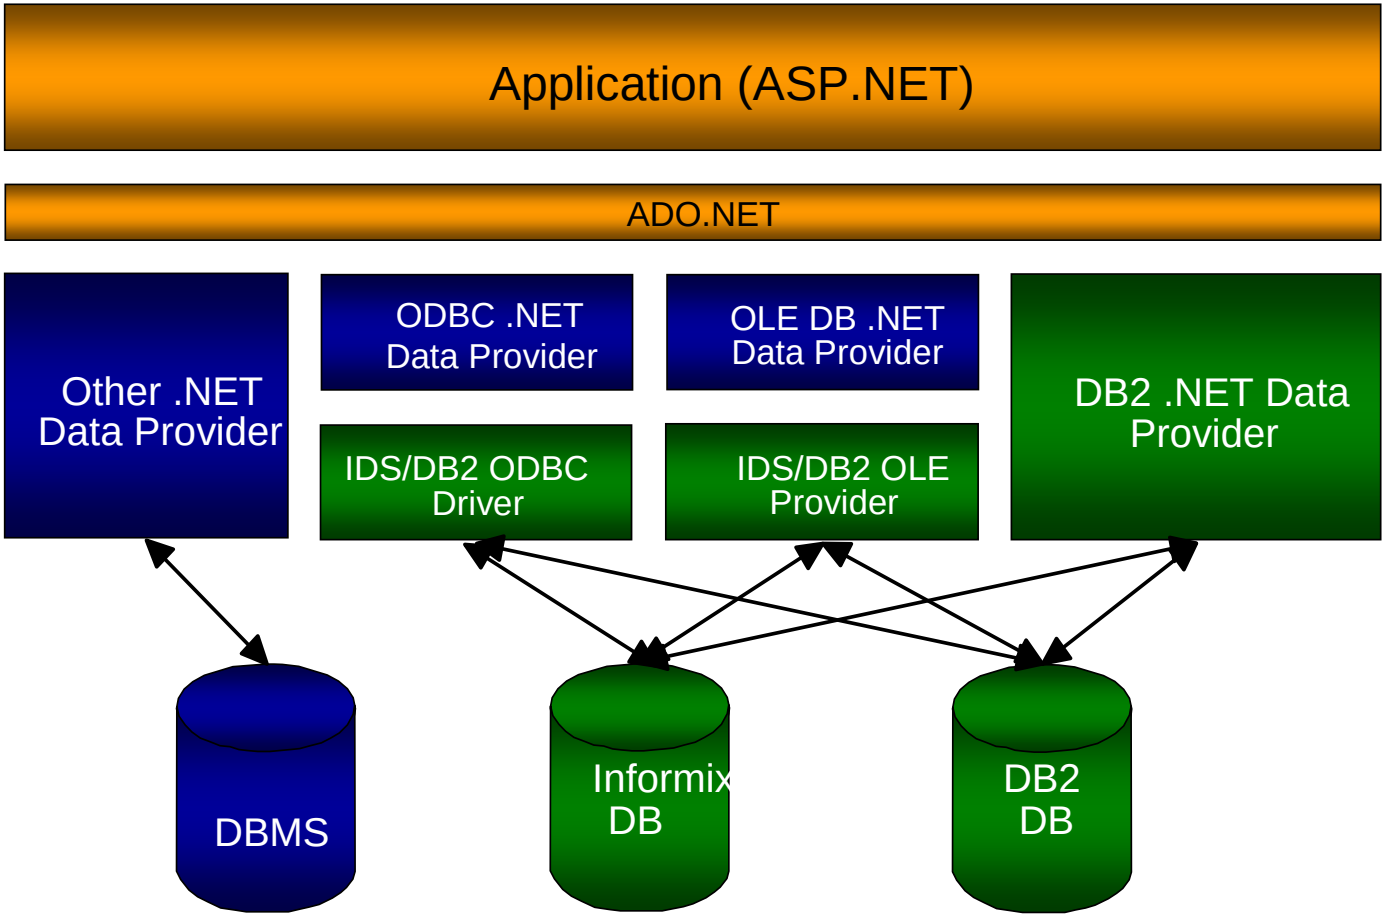

# ADO .NET Data Providers

- **DB2 .NET Data Provider**
- **OLE DB .NET Data Provider**
- **ODBC .NET Data Provider**
- **Customer desire for bridge providers is to enable easy application development against multiple vendor data sources**
  - This need is largely gone with the Framework (FW) 2.0 (circa 2005) DbXxx base classes and DbFactory
  - Have seen much less emphasis on bridge providers
  - MS is following this trend – there is no Entity Framework support for ODBC .NET or OLE DB .NET

# DB2 V8.1.4 .NET Data Provider Comparison

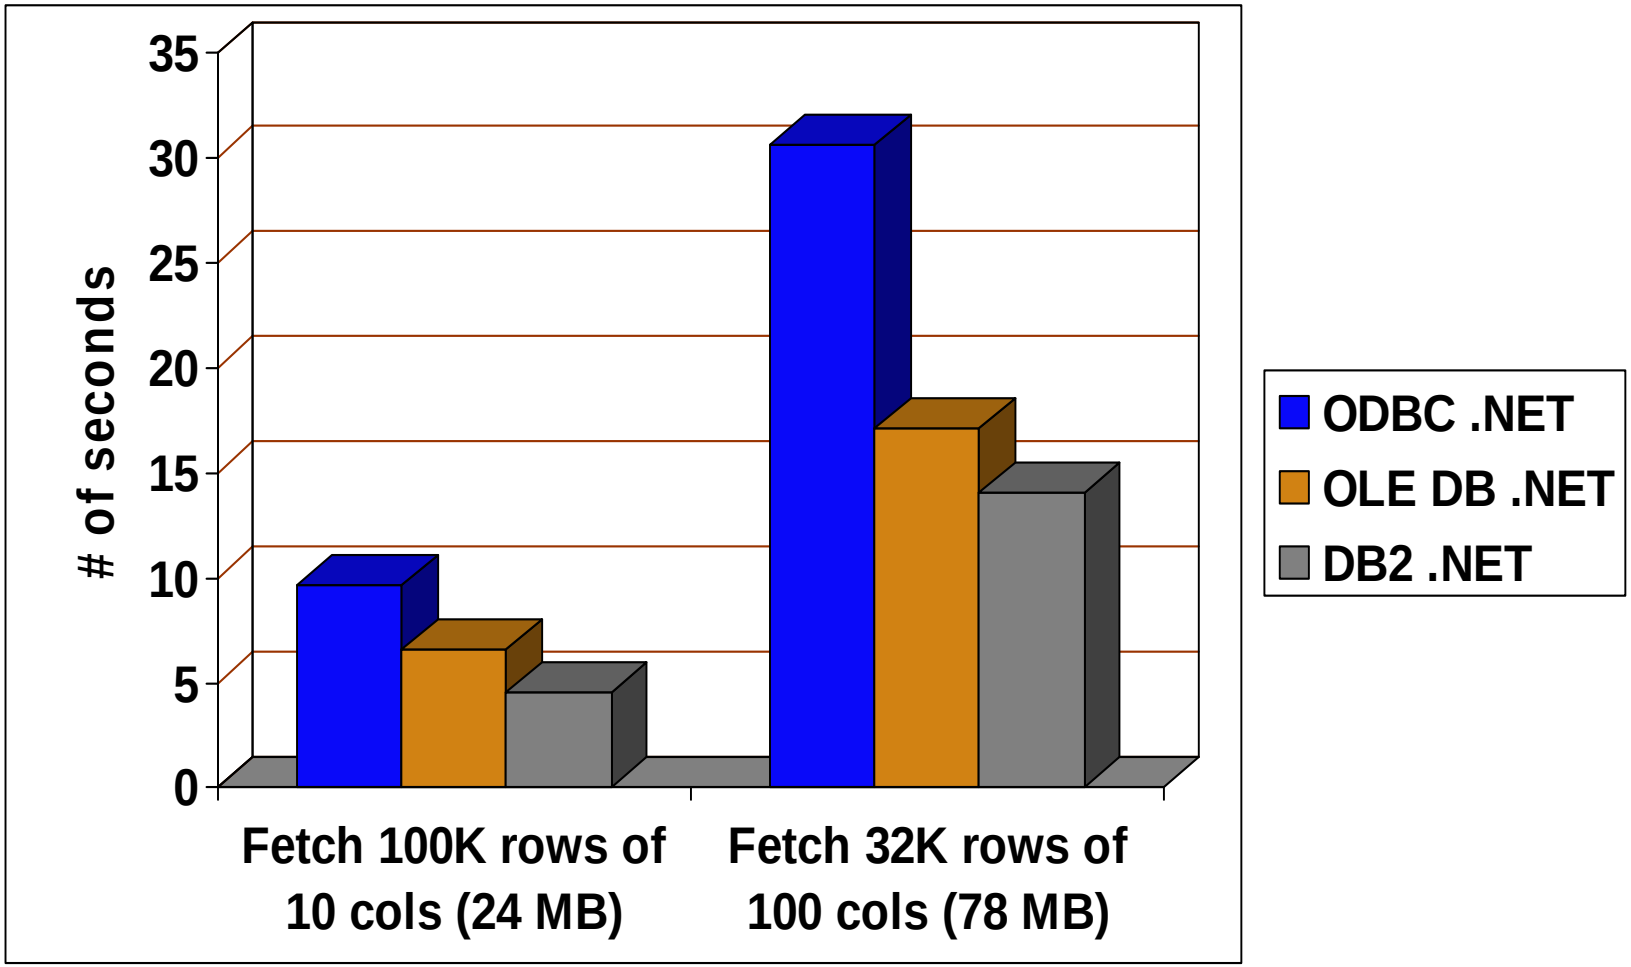

# .NET Data Provider Comparison

- **Bridge providers do not have newer data types and features**
  - DB2BulkCopy, DB2ResultSet
  - XML - integration with XmlReader
  - Decimal Float
- **Difficulty in diagnosing problems**
  - Limitations in MS Bridge providers or ODBC / OLE DB interfaces
  - Connection pooling often an issue
- **IBM VS Addins for DB2 .NET provider only**
- **As usage of bridge providers drops off, so does testing**

# DB2 .NET Data Provider

- **IBM's .NET support is designed to operate in the .NET environment as .NET programmers expect**
- **VS Addins fully leverage rapid application development**
  - VS Addins also exploit IBM unique features
- **Current .NET and Visual Studio skills transportable**
- **DB2 and Informix servers supported**
  - Informix as of V9.5

# DB2 .NET Data Provider

- **DB2 9, 9.5, 9.7**
  - Backwards compatible to any V9.1 LUW server or DB2 Connect server
  - V8, V8.2 are out of service
- **DB2 Connect 9, 9.5, 9.7**
  - Support for z/OS V9 and iSeries V5R4 in V9.1 FP2
- **Informix Cheetah**
  - Supported with V9.5 and later

# VS Addins Function Summary – VS 2008 / 2010

- **Server Explorer**
  - Enumeration - filtering
  - View Data, Script
  - Drag n drop – win app, web app
  - Schema evolution - developers
- **Stored Procedures**
  - SQL Procedure designer
  - SQL keyword highlighting and Intellisense
  - End to end debugging – merged call stack, step into, run to breakpoint
  - CLR Procedures
- **XML**
  - XML Schema, generate sample
  - XML index designer
- **Web Services**
  - IIS, DB2
  - UDF to consume (not CLR)

# Early features in .NET

- **Accounting information (SQLESETI)**
  - Properties on DB2Connection object
- **DB2Connection.Enlist**
  - Allows delayed enlistment in distributed transaction
- **Specify isolation level for distributed transaction**
  - Connection string keyword “isolation level”

# Early features in .NET

- **Named parameters**
  - E.g.. SELECT * from table where department = @dept
  - Implemented in DB2 .NET
- **Chaining**
  - DB2Connection.BeginChain, EndChain
  - Allows non-atomic batching of non-query statements
  - Array of DB2Error objects indicate individual results
- **Multiple Concurrent Data Readers**
  - Enables natural processing of parent – child relationship scenarios

# Provider features in DB2 9

- **XML**
  - Full exploitation of XML Datatype in .NET
    - DB2Xml data type
      - Ability to retrieve XmlReader, string from DB2Xml
    - DB2DataReader.GetDB2Xml, GetXmlReader
    - Bind in string or byte[] to XML column
  - Natural integration with .NET built in XML related classes
    - Bind in any XmlReader as a parameter
    - DB2XmlCommand
      - Direct xquery execution against XML
    - DB2XmlAdapter
      - Read only (Fill) of XPathDocument

# Provider features in DB2 9

- **XML Schema**
  - DB2DataReader.GetXmlSchemaCollection (Set)
    - Retrieve schema(s) used to validate instance of XML document
  - RegisterXmlSchema
    - Registers one or more XML schemas in the DB2 repository

# Provider features in DB2 9

- **DB2ResultSet**
  - First .NET scrollable cursor
  - Scrollable, Sensitive, Updateable and SkipDeleted options
    - Specific combinations determined by DB2 server
    - LOB columns will force forward only
  - Server side locking semantics
  - Retrieve values just like DB2DataReader
  - Set methods to update values
  - DB2UpdateableRecord to insert a row

# ODBC – Typical Applications

- **ODBC / CLI user interchangeably**
  - Think of CLI as ODBC ++
- **ISV applications**
  - Every database has an ODBC driver
  - Has been a mainstay for nearly two decades
- **Open source drivers**
  - Current Perl, Python, Ruby drivers written to vendor ODBC libraries
- **Popular end user apps**
  - Excel, Access

# ODBC – Latest Additions

- New server types – XML, DecFloat, enhanced timestamp
- Sysplex and pureScale support
- Streaming API improvements
- Allow codepage to be set individually per connection handle
- SQL Compatibility support
- Result Sets from Anonymous blocks
- CLI trace using db2trc
- z/OS V10 support

# pureQuery Client Optimization

- **pureQuery client optimization enables dynamic SQL execution and static execution for *existing* applications without any code changes**

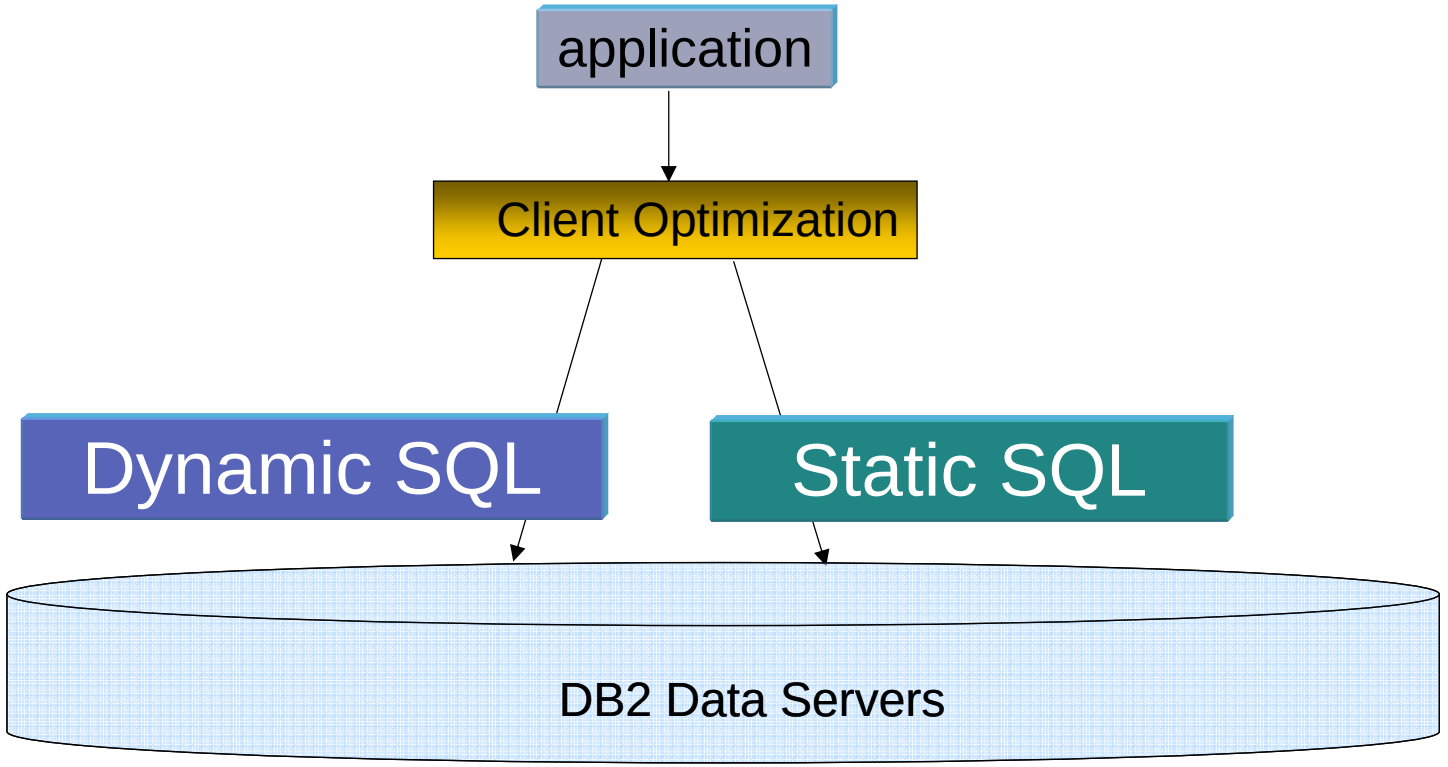

# How does pureQuery client optimization work?

Capture/Configure → Bind → Execute

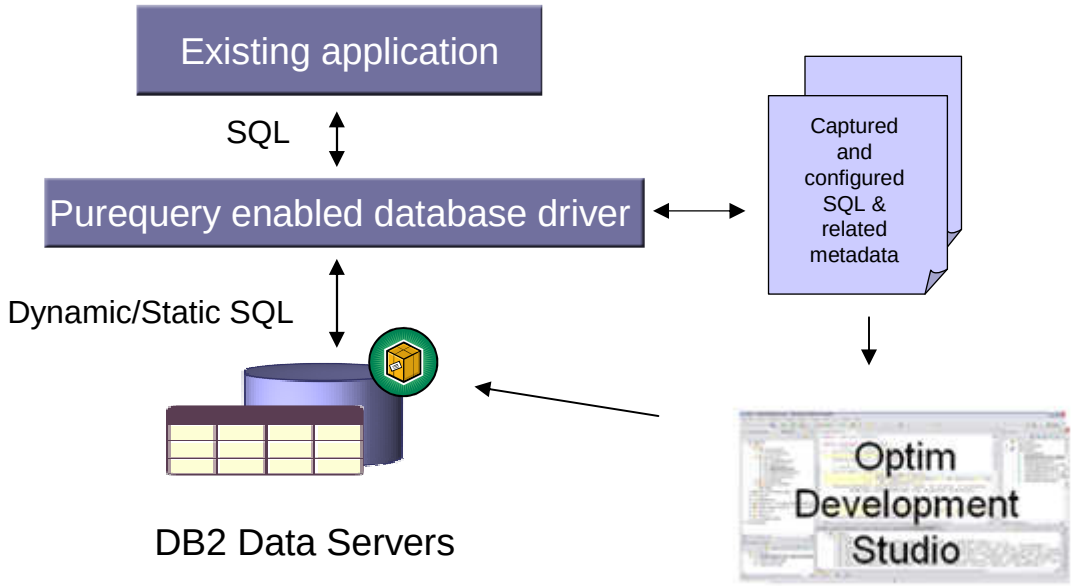

# Database Access and pureQuery
## Many On-ramps for New and Existing Applications

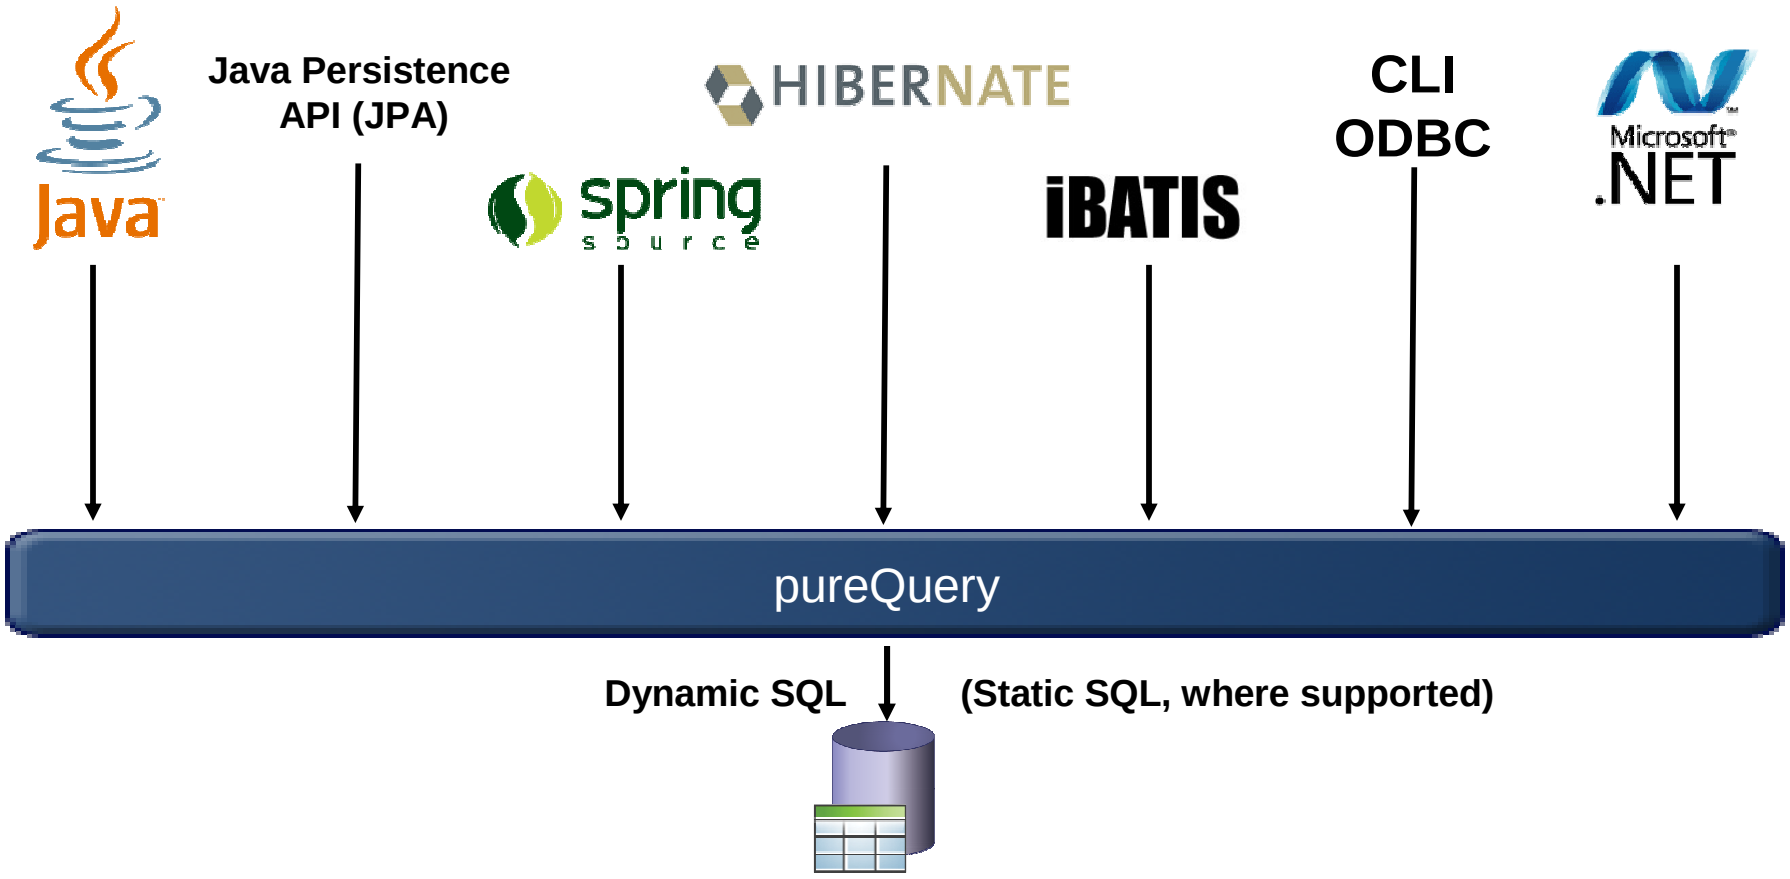

# SQL Execution - Dynamic vs. Static

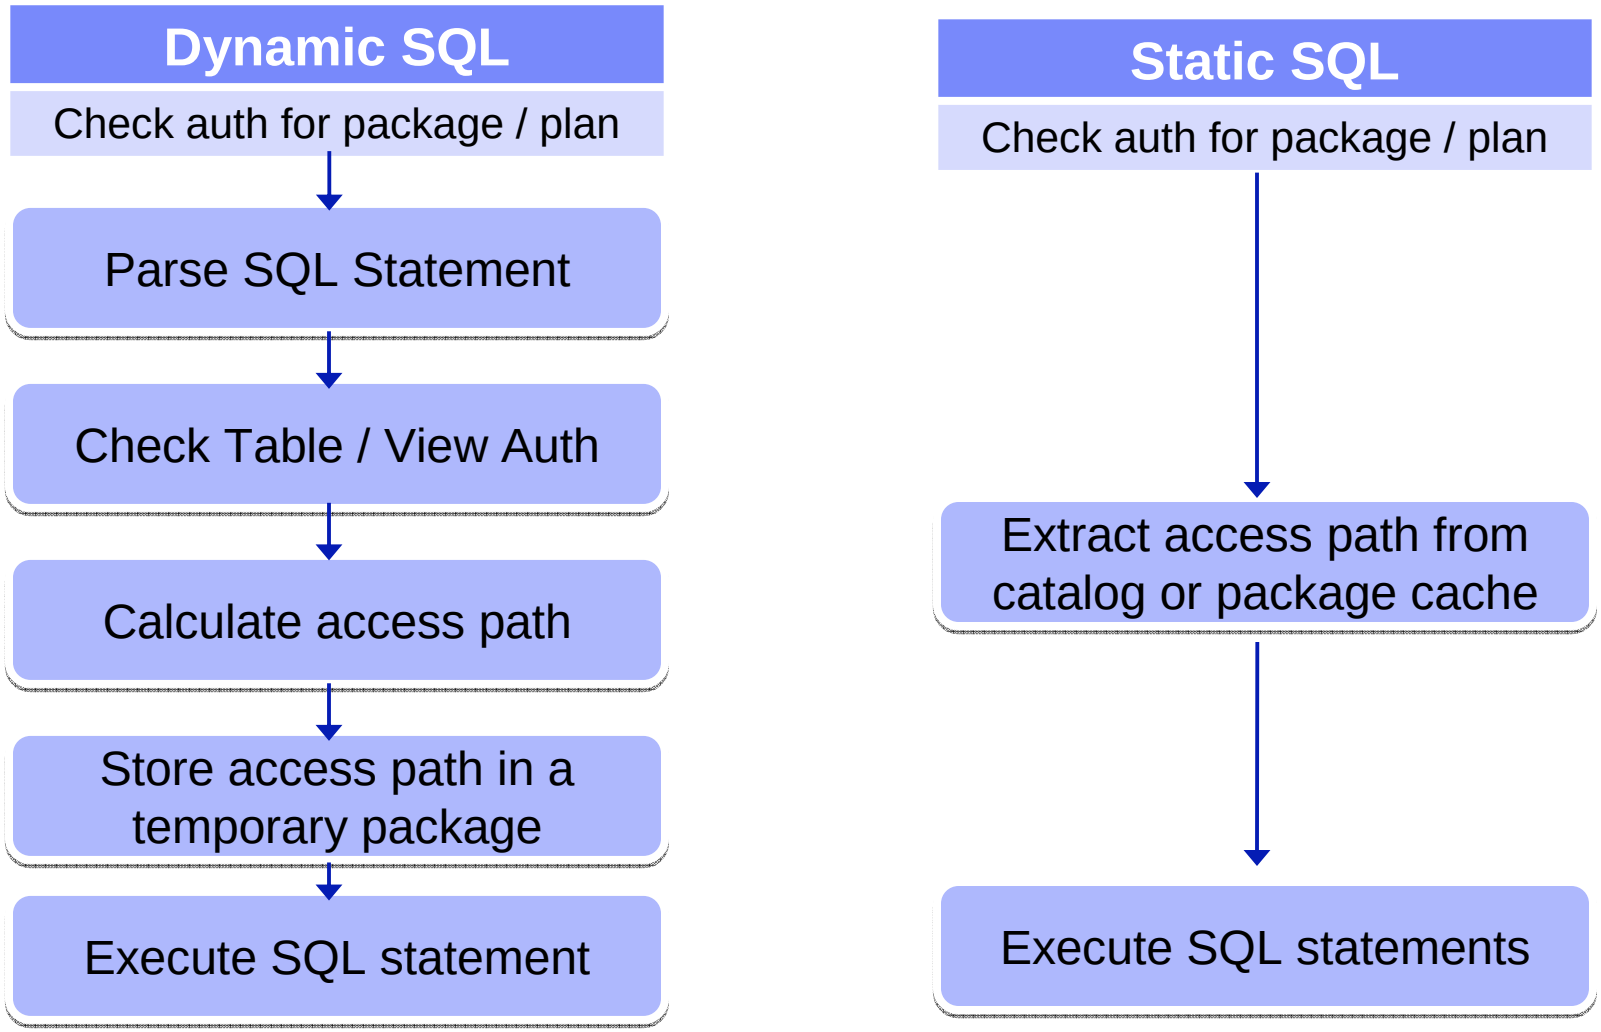




20

# How well does it work?

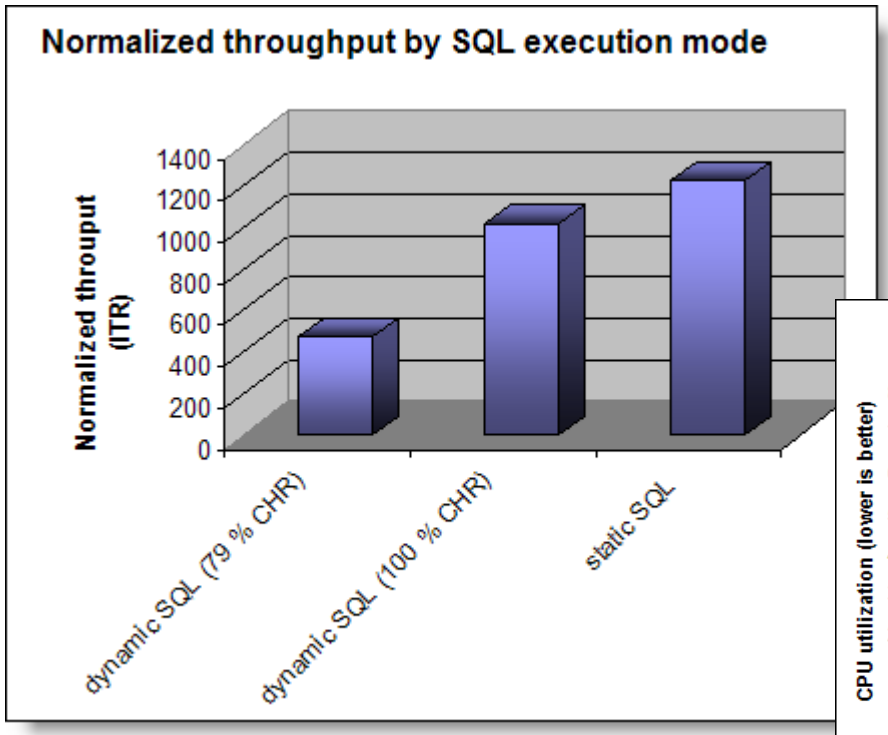


- IRWW – OLTP application
- Application runs in IIS and accesses DB2 for z/OS

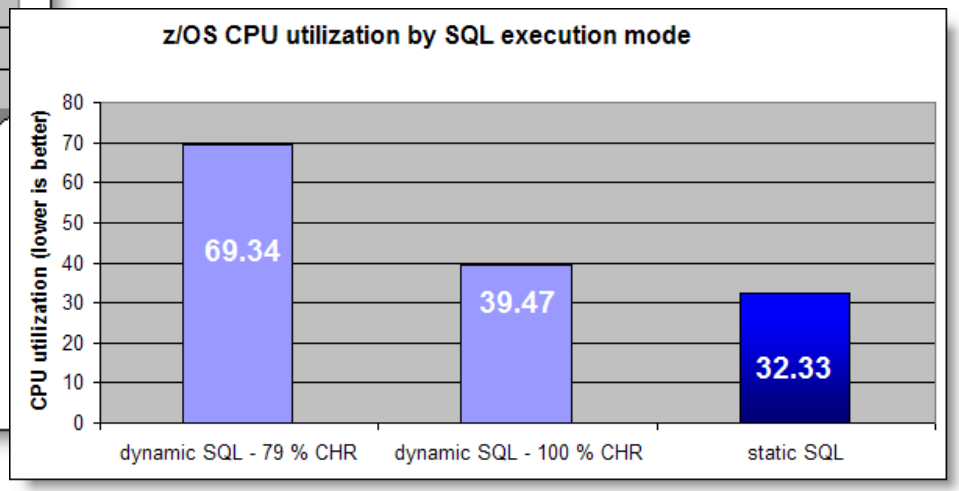


- Throughput during static execution increased by 159% over dynamic SQL execution assuming a 79% statement cache hit ratio

***Any performance data contained in this document were determined in various controlled laboratory environments and are for reference purposes only. Customers should not adapt these performance numbers to their own environments as system performance standards. The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.***

# Improve Performance Without Changing Code- e.g. literal replacement

- **Review SQL coming from applications**

```
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 35213 AND ORDER_METHOD_CODE = 35213
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 31842
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 27368 AND CUST_CC_ID = 27368 ORDER BY CUST_ORDER_DATE ASC
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 20022 AND CUST_ORDER_STATUS_CODE = 2 ORDER BY CUST_ORDER...
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 48679
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 22302 AND CUST_ORDER_STATUS_CODE = 2 ORDER BY CUST_ORDER...
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 11161 AND CUST_ORDER_STATUS_CODE = 2 ORDER BY CUST_ORDER...
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 18494 AND ORDER_METHOD_CODE = 18494
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 40932 AND CUST_CC_ID = 40932 ORDER BY CUST_ORDER_DATE ASC
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 31968 AND CUST_ORDER_STATUS_CODE = 3 ORDER BY CUST_ORDER...
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 7652 AND CUST_ORDER_STATUS_CODE = 2 ORDER BY CUST_ORDER
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 4427
SELECT * FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_CODE = 7149
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_COD
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_COD
SELECT COH.CUST_ORDER_NUMBER, COH.CUST_ORDER_DATE, COH.CUST_ORD
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_COD
SELECT COUNT(*) FROM GOSALESCT.CUST_ORDER_HEADER WHERE CUST_COD
```

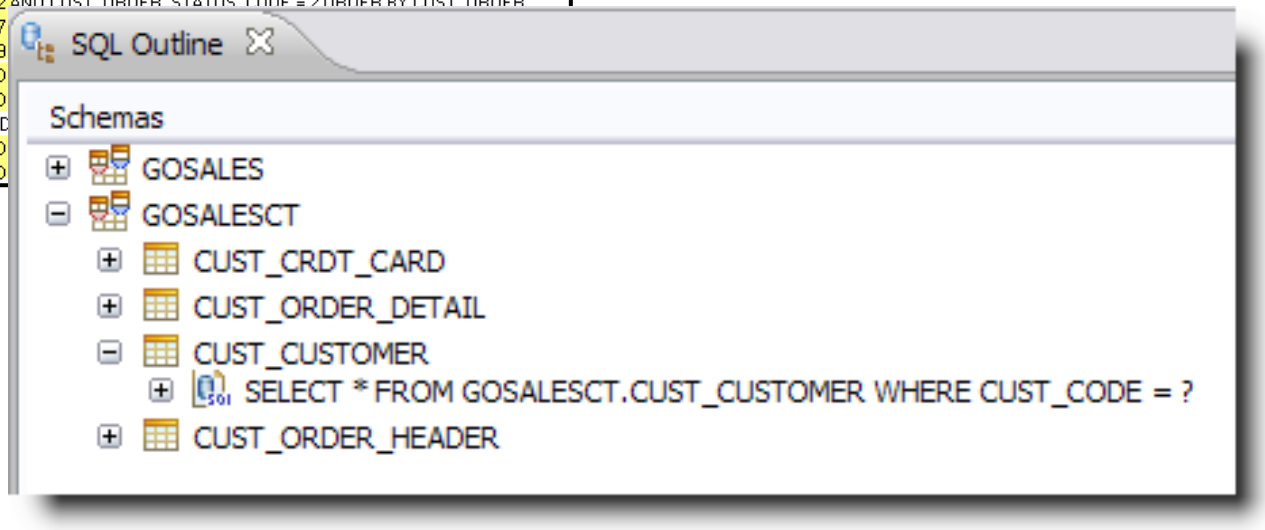


- **Consolidate SQL by replacing literals with parameter markers at runtime**
- **Optimize database resources (e.g. dynamic statement cache) usage**

# Performance Hints – Reduce .NET Memory Usage

- **None of these are bugs, but they will help prevent out of resource errors on the server**
- **Managed environment – why do I care about memory**
  - Try and make the system work less to cleanup after you
- **Select only what you need**
  - Applies to rows as well as columns
    - DB2DataReader.Read assumes all columns will be Get'd
  - BIG cost in ASP.NET
- **Free early to avoid generation 2 allocations**
  - The longer an object lives, the harder it is to delete it
  - The more long term objects the system has, the hard it has to work to free memory

# Performance Hints – Reduce .NET Memory Usage

- **Connections – open late, close early**
  - Maximize the use of connection pooling
- **Access columns in order selected**
  - This will reduce internal buffering and make the most of streaming
- **Other objects – close / dispose, set reference to null**
  - DB2DataReader.Close, DB2Command.Dispose
  - Allows the server to free locks and clean up locators earlier
  - It will eventually get done in the Finalizer, which is run even less often than the garbage collector

# Productivity Tips

- **Named parameters (.NET)**
  - Readability
    - Select * from EMPLOYEE where FNAME = @firstname and LNAME = @lastname
    - Parameters.Add( “lastname”, … );
    - Parameters.Add( “firstname”, … );
  - Allows binding parameters in any order

- **Use Parameter markers vs. building statement from strings**
  - Higher hit rate on server side shared statement cache
  - Avoids SQL injection attacks when parameter values come from user input

# DB2 .NET Notes - DB2Command

- **Command Types:**
  - CommandType.Text (default)
    - SQL statement
  - CommandType.TableDirect
    - Table name equivalent to “SELECT * FROM <tableName>”
  - CommandType.StoredProcedure
    - Stored procedure name equivalent to “CALL <spName> (?,…)”, where the # of parameter markers depends on the # of DB2Parameters specified.
    - Does not allow use of named parameters, parameters must be added in order required

# DB2 .NET Notes - DB2DataReader

- **In order to provide the best data retrieval performance, the Get methods do not perform data type conversions, e.g. an integer column must be retrieved using GetInt32(), and not GetInt16() nor GetInt64().**
  - Strong typing is a characteristic of .NET
  - DB2 .NET provider does allow all types to be retrieved as a string.
- **DataReader.GetValue() can be used to retrieve the best type (as an object) for the column**
  - Performance cost to encapsulate type in an object instance (boxing)

# Function Rollout Overview

- **New function rollout limited to Fixpack stream of current release only**
  - V9.5 FP3 was last set of enhancements for V9.5
  - Current enhancements are in V9.7 stream
- **All active releases fully supported**
  - APARs backfit based on customer request and severity
- **Unless stated, function applies to**
  - DB2 for z/OS V8 and later
  - DB2 for IBM i V5R4 and later

# Leading up to V9.1

- **Cross API**
  - Single client instance install
  - DB2 Runtime Client (> 100 MB)
  - DB2 Client (almost 500 MB)
    - DB2 Connect Personal Edition
- **ODBC / CLI**
  - Terms used interchangeably
- **.NET**
  - Framework (FW) 1.0, 1.1
  - Visual Studio (VS) VS.NET 2002, VS.NET 2003
  - 32-bit only

# V9.1

- **Cross API**
  - Multiple copy support (concept of default copy)
  - XML datatype (DB2 for z/OS V9 only)
- **.NET**
  - FW 2.0, VS 2005
  - Dropped FW 1.0
  - DB2ResultSet (scrollable, server locking semantics)

# V9.1 FP2

- **.NET**
  - 64-bit Provider
  - End to End SQL Procedure debugging for z/OS V9 servers
    - Single step from application into SQL stored procedure
    - Merged call stacks

# V9.5

- **ODBC / CLI**
  - cli driver package available
    - Approx 15 MB, designed for ISVs to embed
- **.NET**
  - VS 2008 supported
  - FW 3.0, 3.5 supported

# V9.5 FP1

- **Cross API**
  - 18 character location names
    - Previously only 8 characters supported

# V9.5 FP2

- **Cross API**
  - SSL
- **ODBC / CLI**
  - Interleaved SQLPutData for dealing with stream data
- **.NET**
  - Enterprise Library (via CodePlex)
  - Entity Framework public beta
  - VS 2008 WPF and WWF integration
  - End to end SQL procedure debugging for z/OS V8 servers

# V9.5 FP3

- **Cross API**
  - IBM Data Server Driver Package (ds driver)
    - db2dsdriver.cfg configuration file introduced
- **ODBC / CLI**
  - SQLCreatePackage API (similar to bind command)
- **.NET**
  - pureQuery static profiling (DB2 Connect Advanced Edition)
  - Entity Framework at FW 3.5 SP1 level
  - Schema node in VS Server Explorer

# V9.5 FP3

- **Cross API**
  - z/OS Sysplex exploitation
    - Client reroute and workload balancing
      - Functions previously only available with DB2 Connect Server
    - Seamless failover
      - Failures in first SQL can be seamlessly rerouted to alternate server without application awareness
    - Configurable in all client packages with dsdriver.cfg file
    - Primary scenario is application servers
      - Performance benefit by removing gateway

# V9.7

- **ODBC / CLI**
  - Interleaved fetch and rollback during streaming
- **.NET**
  - Remove any accidental db2cli.ini support
  - Dropped FW 1.1
  - Program Name and Program ID client information properties
  - VS tooling for pureQuery static profiling

# V9.7 FP1

- **Cross API**
  - Environment variable for config file location
  - Ability to reload config file for ACR section
  - Failback to primary support for client affinity
- **ODBC / CLI**
  - pureQuery static profiling
  - Interleaved insert with streaming
- **.NET**
  - Full alter support for procedures in VS
  - VS 2008 LINQ over XML reference (z/OS only)

# V9.7 FP2

- **ODBC / CLI**
  - Command line to register ODBC data source
- **.NET**
  - 32-bit and 64-bit coexistence
  - Entity Framework Filtering
  - Override for with hold cursors

# V9.7 FP3a

- **Cross API**
  - Unlimited Edition server based license key (z/OS only)
  - DB2 for z/OS V10 exploitation
    - Binary XML
    - Timestamp precision plus timezone
    - Currently committed semantics
    - Extended indicators
    - Explain modes

# V9.7 FP3a (continued)

- **Cross API**
  - Performance Expert Extended Insight integration
- **ODBC / CLI**
  - Network statistics API
  - Retrieve last member used on connection
  - Instance based client support with dsdriver.cfg
- **.NET**
  - Entity Framework canonical functions
  - .NET 4.0, VS 2010, Entity Framework 4.0 beta

# V9.7 FP4

- **Cross API**
  - Sysplex migration support for z/OS V10
- **ODBC / CLI**
  - db2cli32 in 64-bit builds
  - db2cli to report on unrecognized settings
  - CLI load enabled for async operation
- **.NET**
  - .NET 4.0, VS 2010, Entity Framework 4.0 GA
  - High precision timestamp and decfloat mapping into dataset
  - Instance based client support with dsdriver.cfg

# V9.7 FP5

- **Cross API**
  - Schema filter in connection string
  - Passphrase up to 100 characters
  - Command line tool to add entry to dsdriver.cfg
  - Alternate group failover
  - Password in dsdriver.cfg file
- **ODBC / CLI**
  - ODBC 3.8
- **.NET**
  - Block for n rows override capability
  - Ability to disable auto-rebind
  - CommandBuilder CompareRowVersion support
  - Array Input (batching extension)
  - MS Trace Integration

# High Priority Requirements

- **Cross API**
  - z/OS SSL authentication
  - Seamless improvements for Sysplex, pureScale
  - Client Info default values
- **ODBC / CLI**
  - Validate improvements
- **.NET**
  - Additional filtering criteria
  - Integration with Optim tooling

# Summary

- Multiple provider options → DB2 .NET
- DB2 .NET Data Provider
- Special support for zOS – sysplex, pQ
- Ongoing exploitation of .NET features
- Continue to enhance ODBC interfaces
- Continue to exploit .NET and server strengths

# Online references

- **Developer Works Visual Studio Zone**
  - http://www.ibm.com/developerworks/db2/zones/vstudio/
- **DB2 .NET FAQ**
  - http://www.ibm.com/developerworks/wikis/display/DB2/DB2+and+.NET+FAQ
- **IBM .NET Forum**
  - http://www.ibm.com/developerworks/forums/dw_forum.jsp?forum=467&cat=19

# More samples and tutorials

- **VS Addins Tutorials**
  - VS Addins with IDS
    - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0703jayakumar/?S_TACT=105AGX11&S_CMP=LP
  - Good overview of .NET with VS Addins
    - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0710jayakumar/?S_TACT=105AGX11&S_CMP=LP
    - http://www.ibm.com/developerworks/db2/library/techarticle/dm-0711kumar/?S_TACT=105AGX11&S_CMP=LP