

A Peek Into The Future Direction: **Data Server Technology**

*Gene Fuh, Ph.D.
Distinguished Engineer and Senior Manager
Query Technology, DB2 for z/OS
IBM Silicon Valley Laboratory*



DB2 for z/OS Technical Forum
Taipei, October 5-6, 2009

Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

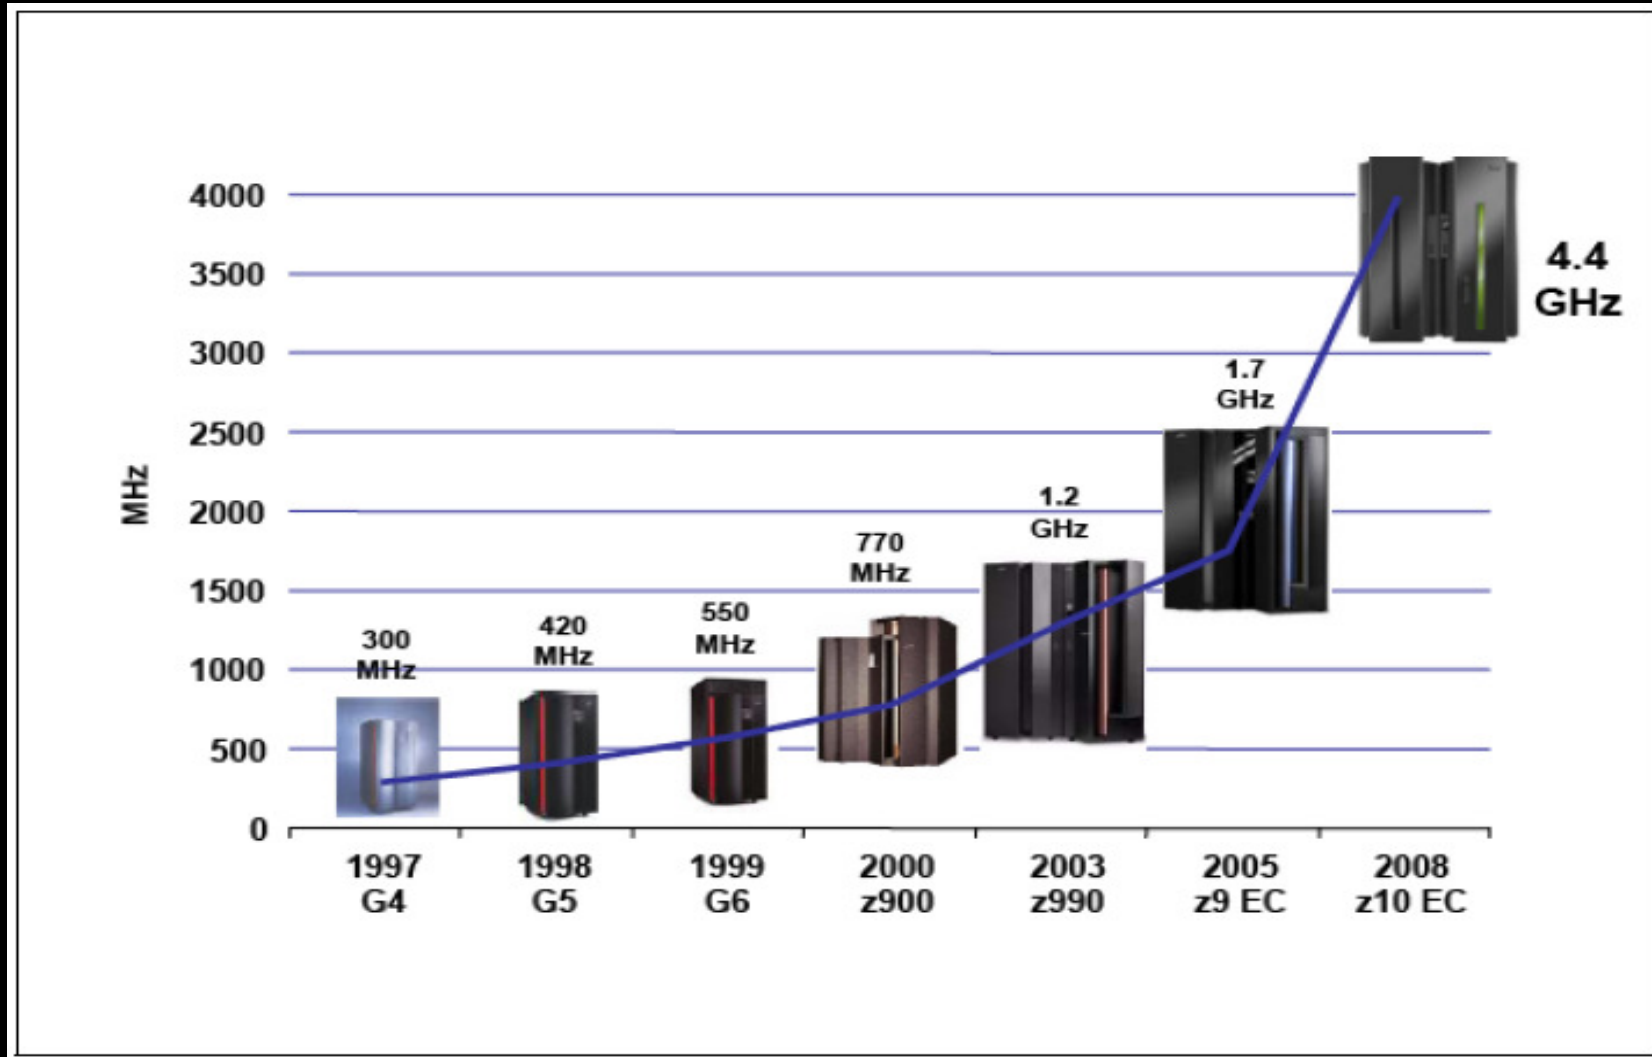
IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- **CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR**
- **ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.**

Mainframe Evolution



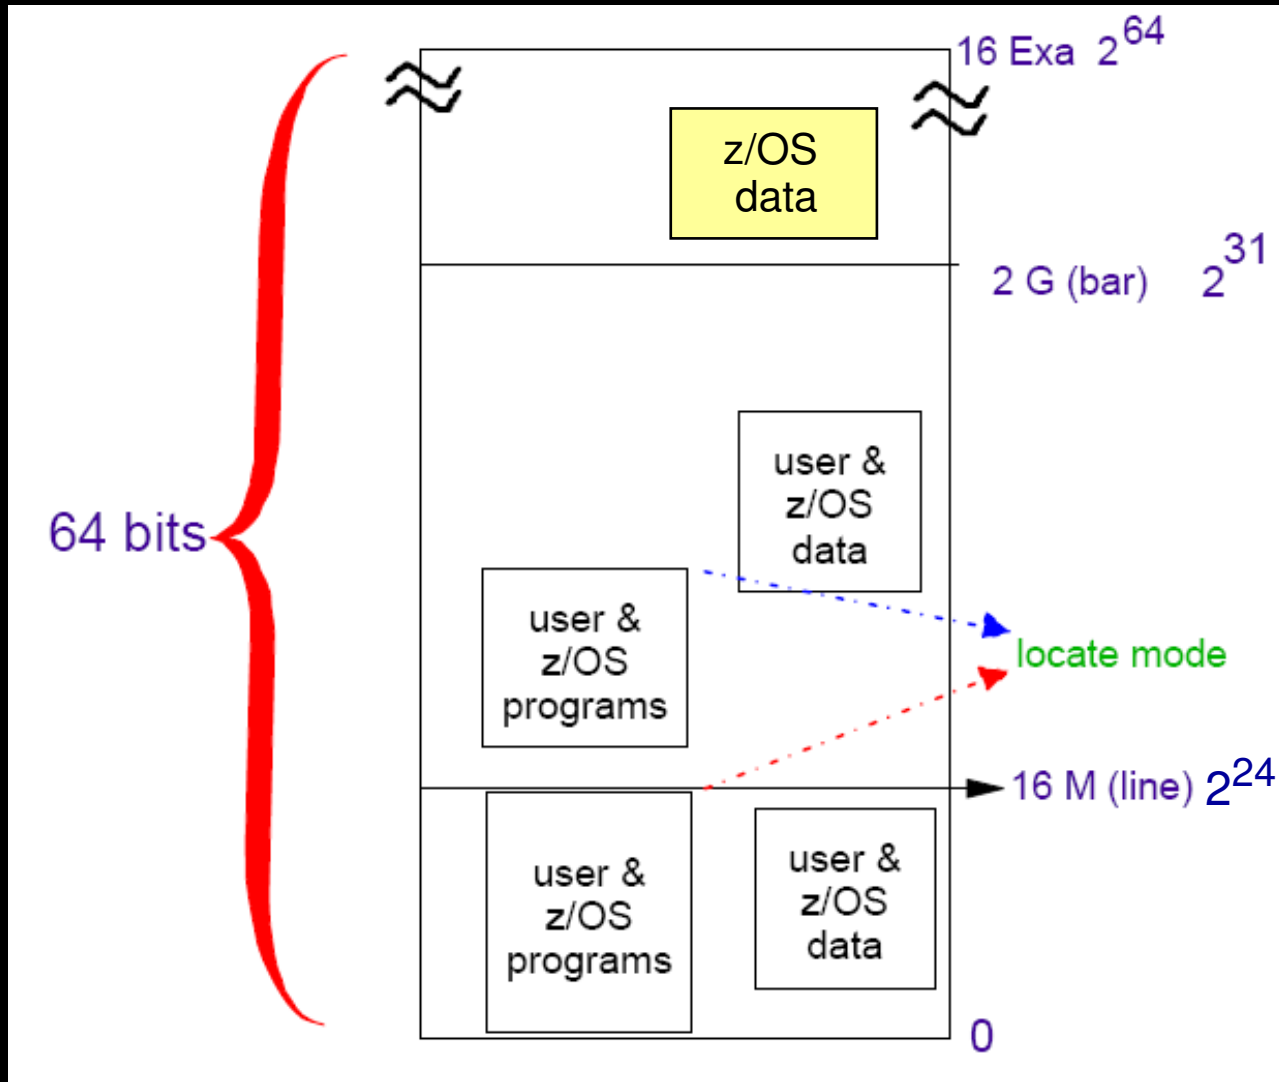
z/Hardware

- Also known as **Mainframe**
- **Most recent incarnation**
 - ▶ z10 Family
- **Z10 Major Characteristics**
 - ▶ 4.4 Ghz Microprocessors
 - ▶ Quad-core chips
 - ▶ Max. Number of Core = 77
 - User Core(*) = 64
 - SAP (System Assist Processor) = 11
 - Spare Core = 2
 - ▶ Max. memory Size = 1.5 Tb
 - ▶ Hardware Accelerators
 - Cryptography
 - Data Compression
 - Decimal Floating Point arithmetic



(*) Std., ICF, IFL, zIIP, zAAP

z/OS: Address Space Concept



DB2 for z/OS

- **“Born” in 1983**
- **RDBMS for z System platform**
 - ▶ Exploits the z platform
- **Hybrid engine**
 - ▶ Relational data
 - ▶ ORDBMS extensions
 - ▶ XML documents
- **Parallel query engine**
- **Cluster support on Parallel Sysplex**
 - ▶ Data Sharing architecture
 - ▶ High Availability & Parallelism
- **Distributed Data support**



Programming Interfaces

- **Developer communities**

- ▶ COBOL, PL/I, C, C++,
- ▶ REXX™, APL2®, Assembler, Fortran
- ▶ Java™ (JDBC / SQLJ)
- ▶ .NET (C#, VB .NET)
- ▶ PHP
- ▶ Perl
- ▶ Python
- ▶ Ruby on Rails
- ▶ Toad for DB2



* Interesting Facts about DB2 for z/OS

- ***Used by...***
 - ▶ The top 59 banks in the world
 - ▶ 23 of the top 25 US retailers
 - ▶ 9 of the top 10 global life/health insurance providers
- ***Performance, Performance, Performance***
 - ▶ Delivered peak online throughput of 15,353 transactions per second for a large Asian bank benchmark. 18 million transactions completed in 20 minutes.
 - ▶ Supports the world's largest known peak database workload - 1.1 Billion SQL statements per hour at UPS
 - ▶ The world's largest known transaction processing database – 23.1 TB at UK Land Registry



Strategic Areas At a Glance - *Addressing Corporate Data Goals*

<p>Application Enablement</p>	<ul style="list-style-type: none"> • pureXML enhancements • Temporal database • Last Committed reads • SQL improvements that simplify porting
<p>RAS, Performance, Scalability, Security</p>	<ul style="list-style-type: none"> • Wide range of performance improvements • More online schema changes • Catalog restructure for improved concurrency • Fine grained access control • Hash access to data • New DBA privileges with finer granularity
<p>Simplification, Reduced TCO</p>	<ul style="list-style-type: none"> • Full 64-bit SQL runtime • Auto stats • Data compression on the fly • Query stability enhancements • Reduced need for REORG • Utilities enhancements
<p>Dynamic Warehousing</p>	<ul style="list-style-type: none"> • Moving sum, moving average • Many query optimization improvements • Query parallelism improvements • Advanced query acceleration



DB2 X: Business Security & Compliance Needs

- Protect sensitive data from privileged users
 - ▶ SYSADM without data access
 - Separate authority to perform security related tasks
 - Allow EXPLAIN without execute privilege or ability to access data
 - Audit privileged users
-
- “As of” query, temporal or versioned data
 - Fine grained access control
 - Allow masking of value
 - Restrict user access to individual cells



Use disk encryption

Product						
Tkr	Name	ClosePrice	StartDate	EndDate	PrSysStart	PrSysEnd
IBM	International Business Machines Corp.	94	1/1/2009	1/2/2009	1/1/2009	12/31/9999
..
IBM	International Business Machines Corp.	91	1/30/2009	1/31/2009	1/30/2009	12/31/9999
..
IBM	International Business Machines Corp.	92	3/2/2009	3/3/2009	3/2/2009	12/31/9999
..
IBM	International Business Machines Corp.	95	3/31/2009	4/1/2009	3/31/2009	12/31/9999
..

Product has one row per trading day for each product transacted by firm.

Position							Trade						
AccountID	Tkr	Qty	PosStrt	PosEnd	PosSysStrt	PosSysEnd	AccountID	Tkr	Qty	TrdStDate	TrdEndDate	TrdSysStrt	TrdSysEnd
54667	IBM	1000	1/30/2009	3/2/2009	1/30/2009	12/31/9999	54667	IBM	1000	1/30/2009	1/31/2009	1/30/2009	12/31/9999
54667	IBM	2000	3/2/2009	3/31/2009	3/2/2009	12/31/9999	54667	IBM	1000	3/2/2009	3/3/2009	3/2/2009	12/31/9999
54667	IBM	3000	3/31/2009	12/31/9999	3/31/2009	12/31/9999	54667	IBM	1000	3/31/2009	4/1/2009	3/31/2009	12/31/9999

- Valid Time set by application*
- Valid Time set by application*
- DB2 transaction time when row image became current (set on insert & update)*
- DB2 transaction time when row image ceases to be current (set on insert, changed by update & delete)*

Business Scenario 1: On 4/15/2009 we are asked to account for 500 shares of IBM bought by client on 2/16/2009

Product						
Tkr	Name	ClosePrice	StartDate	EndDate	PrSysStart	PrSysEnd
IBM	International Business Machines Corp.	94	1/1/2009	1/2/2009	1/1/2009	12/31/9999
..
IBM	International Business Machines Corp.	91	1/30/2009	1/31/2009	1/30/2009	12/31/9999
..
IBM	International Business Machines Corp.	92	3/2/2009	3/3/2009	3/2/2009	12/31/9999
..
IBM	International Business Machines Corp.	95	3/31/2009	4/1/2009	3/31/2009	12/31/9999
..

Posiiton							Trade						
AccountID	Tkr	Qty	PosStrt	PosEnd	PosSysStrt	PosSysEnd	AccountID	Tkr	Qty	TrdStDate	TrdEndDate	TrdSysStrt	TrdSysEnd
54667	IBM	1000	1/30/2009	3/2/2009	1/30/2009	12/31/9999	54667	IBM	1000	1/30/2009	1/31/2009	1/30/2009	12/31/9999
54667	IBM	2000	3/2/2009	3/31/2009	3/2/2009	12/31/9999	54667	IBM	1000	3/2/2009	3/3/2009	3/2/2009	12/31/9999
54667	IBM	3000	3/31/2009	12/31/9999	3/31/2009	12/31/9999	54667	IBM	1000	3/31/2009	4/1/2009	3/31/2009	12/31/9999

Before
After

Posiiton							Trade						
AccountID	Tkr	Qty	PosStrt	PosEnd	PosSysStrt	PosSysEnd	AccountID	Tkr	Qty	TrdStDate	TrdEndDate	TrdSysStrt	TrdSysEnd
54667	IBM	1000	1/30/2009	3/2/2009	1/30/2009	4/15/2009	54667	IBM	1000	1/30/2009	1/31/2009	1/30/2009	12/31/9999
54667	IBM	2000	3/2/2009	3/31/2009	3/2/2009	4/15/2009	54667	IBM	1000	3/2/2009	3/3/2009	3/2/2009	12/31/9999
54667	IBM	3000	3/31/2009	12/31/9999	4/1/2009	4/15/2009	54667	IBM	1000	3/31/2009	4/1/2009	3/31/2009	12/31/9999
54667	IBM	1000	1/30/2009	2/16/2009	4/15/2009	12/31/9999	54667	IBM	500	2/16/2009	2/17/2009	4/15/2009	12/31/9999
54667	IBM	1500	2/16/2009	3/2/2009	4/15/2009	12/31/9999							
54667	IBM	2500	3/2/2009	3/31/2009	4/15/2009	12/31/9999							
54667	IBM	3500	3/31/2009	12/31/9999	4/15/2009	12/31/9999							

Shaded cells represent rows in history



Posiiton						
AccountID	Tkr	Qty	PosStrt	PosEnd	PosSysStrt	PosSysEnd
54667	IBM	1000	1/30/2009	3/2/2009	1/30/2009	4/15/2009
54667	IBM	2000	3/2/2009	3/31/2009	3/2/2009	4/15/2009
54667	IBM	3000	3/31/2009	12/31/9999	4/1/2009	4/15/2009
54667	IBM	1000	1/30/2009	2/16/2009	4/15/2009	12/31/9999
54667	IBM	1500	2/16/2009	3/2/2009	4/15/2009	12/31/9999
54667	IBM	2500	3/2/2009	3/31/2009	4/15/2009	12/31/9999
54667	IBM	3500	3/31/2009	12/31/9999	4/15/2009	12/31/9999

Current History

Valid Time

Transaction Time

Today is 4/16/2009, report the client's position on 2/20/2009?

Tkr	Qty
IBM	1500

Today is 4/16/2009, on client's 2/28 statement what position did we report for 2/20/2009?

Tkr	Qty
IBM	1000

Fine Grain Access Control - Requirements

- SQL row filtering
 - ▶ Who can see what rows
 - ▶ Applicable to SELECT, INSERT, UPDATE, DELETE, MERGE

```
CREATE ROW FILTER employeeData ON employees  
      WHERE VerifyGroup(SESSION_USER, 'Manager', 'HR') = 1  
FOR ALL ACCESS;
```

**Only authorization IDs of 'Manager'
group and 'HR' group can
access the records**

Fine Grain Access Control - Requirements

- SQL column masking
 - ▶ Who can see what value for a column
 - ▶ Applicable to the output of outermost subselect
 - No impact to comparison, grouping, ordering, etc.

```
CREATE COLUMN MASK Compensation ON employees  
FOR COLUMN salary  
RETURN case VerifyGroup(SESSION_USER, 'Manager')  
      when 1 then salary else NULL end case;
```

Only authorization IDs of 'Manager'
group can read the *salary* column

Fine Grain Access Control – Benefits


- Easy to maintain and implement security policies in DB2
 - ▶ More flexibility via SQL
 - ▶ Separation of security logic and application logic
 - ▶ No need to manage many views; no view's updatability issue
 - ▶ No need to overload applications
- Tighter security
 - ▶ Data-centric within database
 - ▶ No backdoor to bypass views or applications
 - ▶ More granularity via row permissions and column masks
- User friendly
 - ▶ No need to remember various view or application names
 - ▶ Allowing usage of ad-hoc query tools, report generation tools

Plan Stability – DB2 9


- At REBIND, save old copies of packages
 - Catalog tables
 - Directory
- Two flavors
 - BASIC and EXTENDED
 - Controlled by new ZPARM PLANMGMT
 - Default is OFF
 - Also supported as REBIND options
- REBIND PACKAGE ...
 - PLANMGMT(BASIC)
 - 2 copies: Current and Previous
 - PLANMGMT(EXTENDED)
 - 3 copies: Current, Previous, Original
- Most bind options can be changed at REBIND
 - But a few must be the same ...
- REBIND PACKAGE ...
 - SWITCH(PREVIOUS)
 - Switch between current & previous
 - SWITCH(ORIGINAL)
 - Switch between current & original
- FREE PACKAGE ...
 - PLANMGMTSCOPE(ALL) – Free package completely
 - PLANMGMTSCOPE(INACTIVE) – Free all old copies
- Catalog support
 - SYSPACKAGE reflects active copy
 - SYSPACKDEP reflects dependencies of all copies
 - Other catalogs (SYSPKSYSTEM, ...) reflect metadata for all copies
- Invalidation and Auto Bind
 - Each copy invalidated separately
 - Auto bind replaces only the current copy – previous and original are not affected

Plan Stability – DB2 9

REBIND ... PLANMGMT(BASIC)



REBIND ... SWITCH(PREVIOUS)

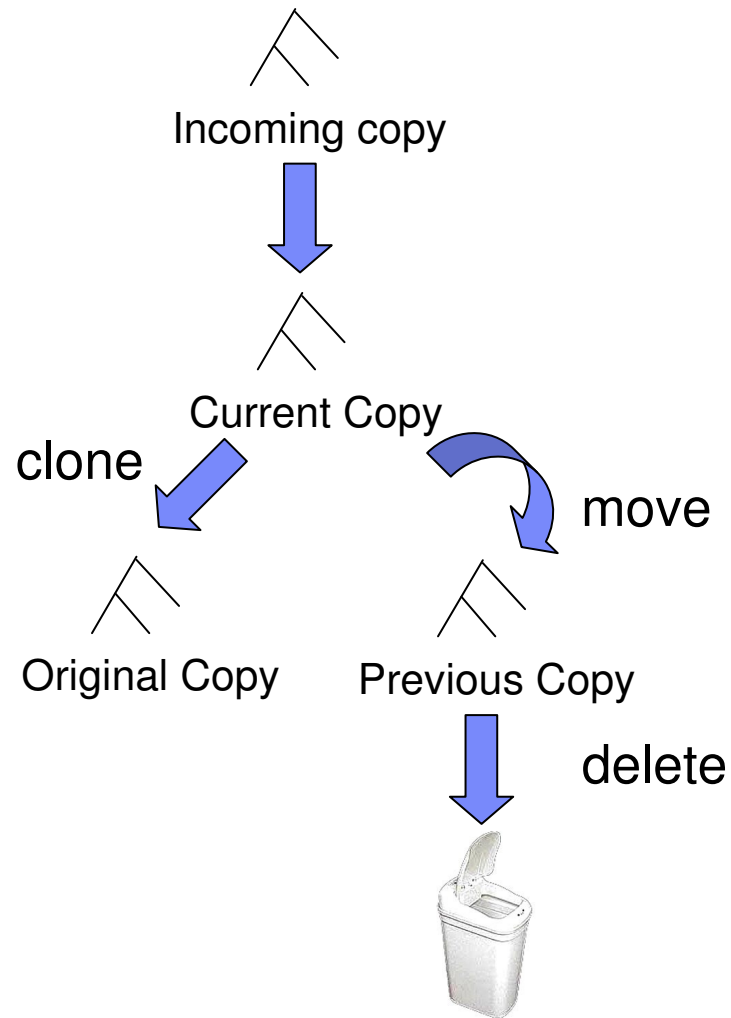


Similar processing for:


- EDM Sections
- SYSPACKDEP dependencies
- SYSPACKAGE record saved in EDM

Plan Stability – DB2 9

REBIND ... PLANMGMT(EXTENDED)

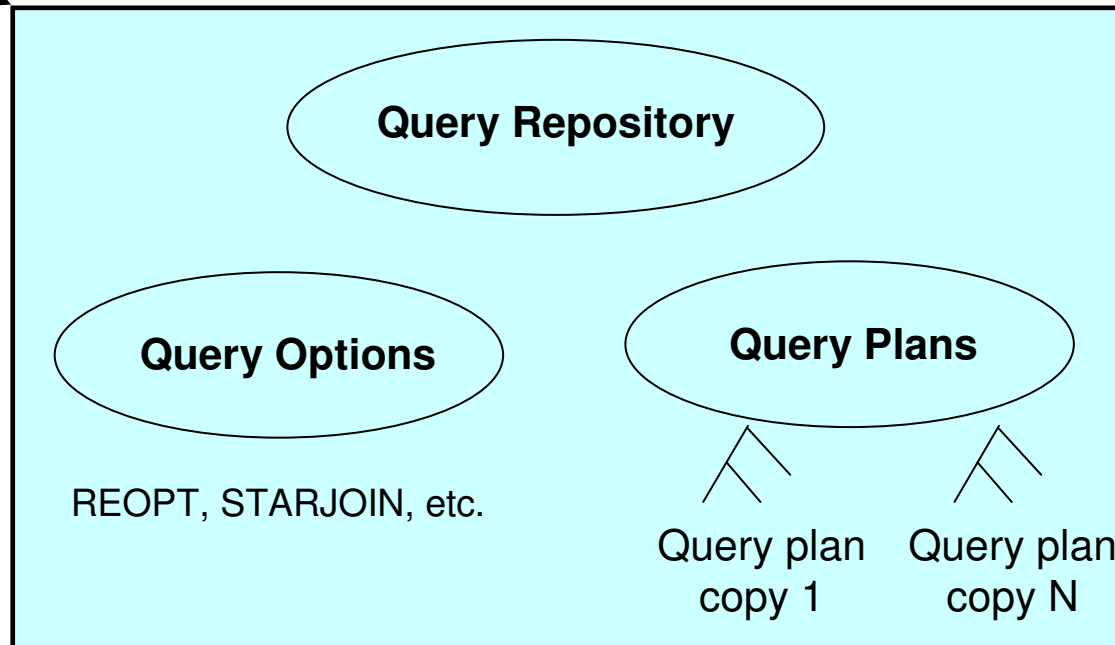


REBIND ... SWITCH(ORIGINAL)



Plan Stability – Beyond DB2 9

BIND/REBIND



Options & Overrides

PREPARE

One query, many query plans

- Historical (in DB2 X)
- Parametric (eventually)

Plan Stability – Beyond DB2 9

Provide an unprecedented level of stability of query performance achieved by stabilizing access paths:

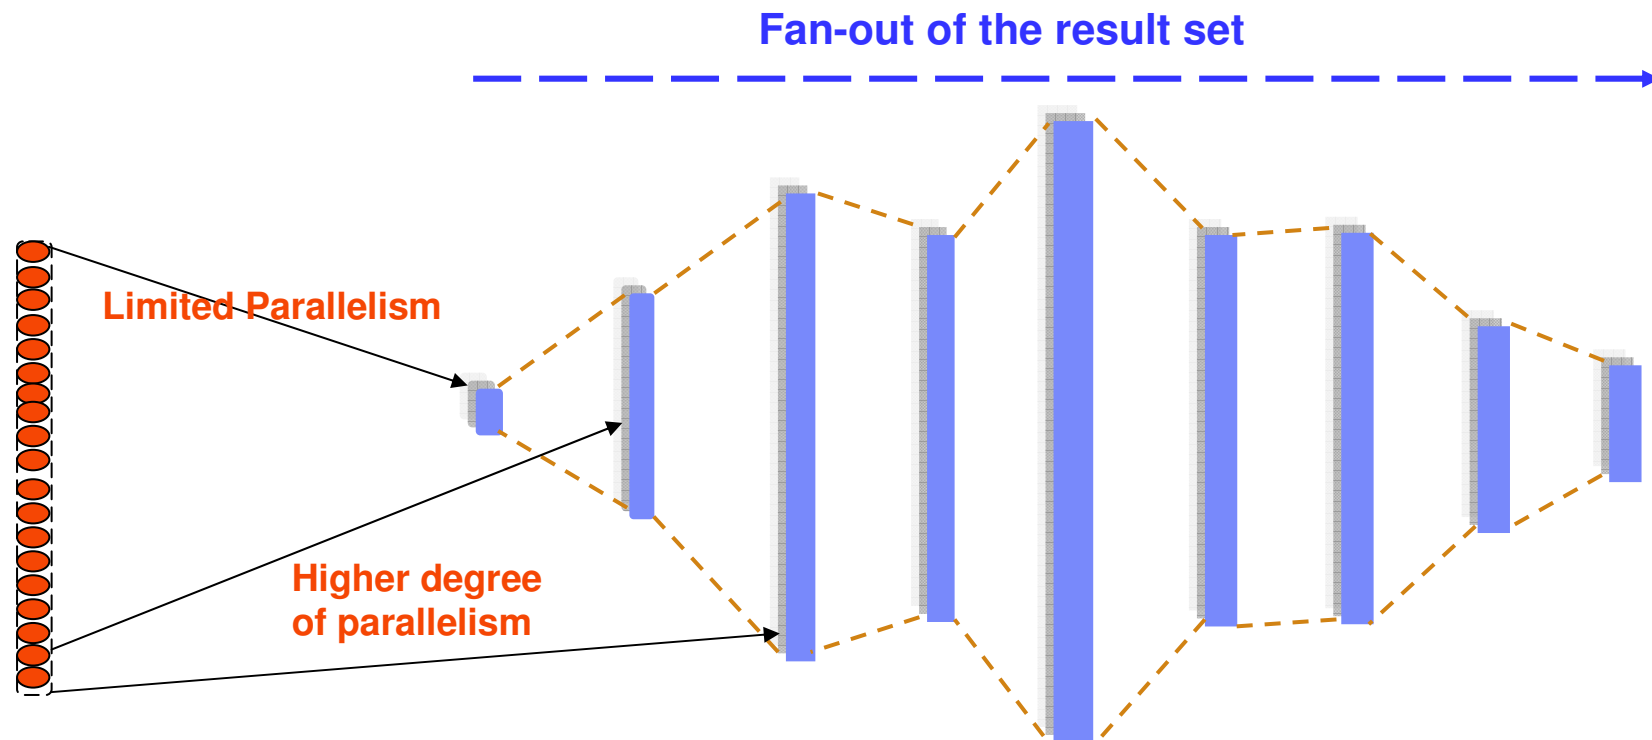
- ▶ Static SQL
 - Relief from REBIND regressions
 - Remove the fear of REBINDing!
 - Enable REOPT support
- ▶ Dynamic SQL
 - Remove the unpredictability of PREPARE
 - Extend Static SQL benefits to Dynamic SQL

Support

- Access path repository
- Versioning
- “Fallback”
- “Lockdown”
- Manual overrides
- Hints: easily influence access paths without changing apps
- Per-statement BIND options

Query Parallelism – Effectiveness

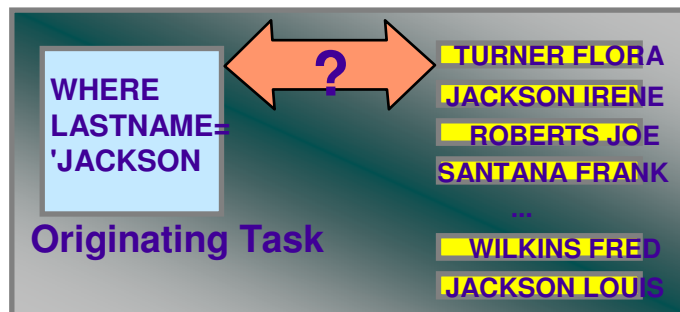
- Opening more opportunities for query parallelism
- Higher degree of parallelism through record partitioning



- Better scalability through dynamic partition of working set

Query Parallelism – Efficiency


- Dynamic working set
 - ▶ More workload partitions than number of current tasks
 - ▶ Available task takes the next available partition
 - ▶ Better workload balance



```

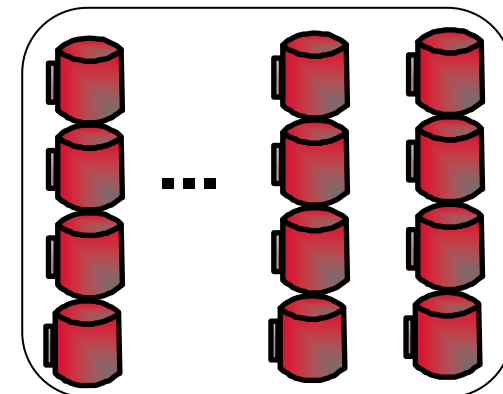
SELECT CUSTOMER, REGION
FROM CUST
WHERE LASTNAME = 'JACKSON'
ORDER BY REGION;
    
```

Parallel Tasks

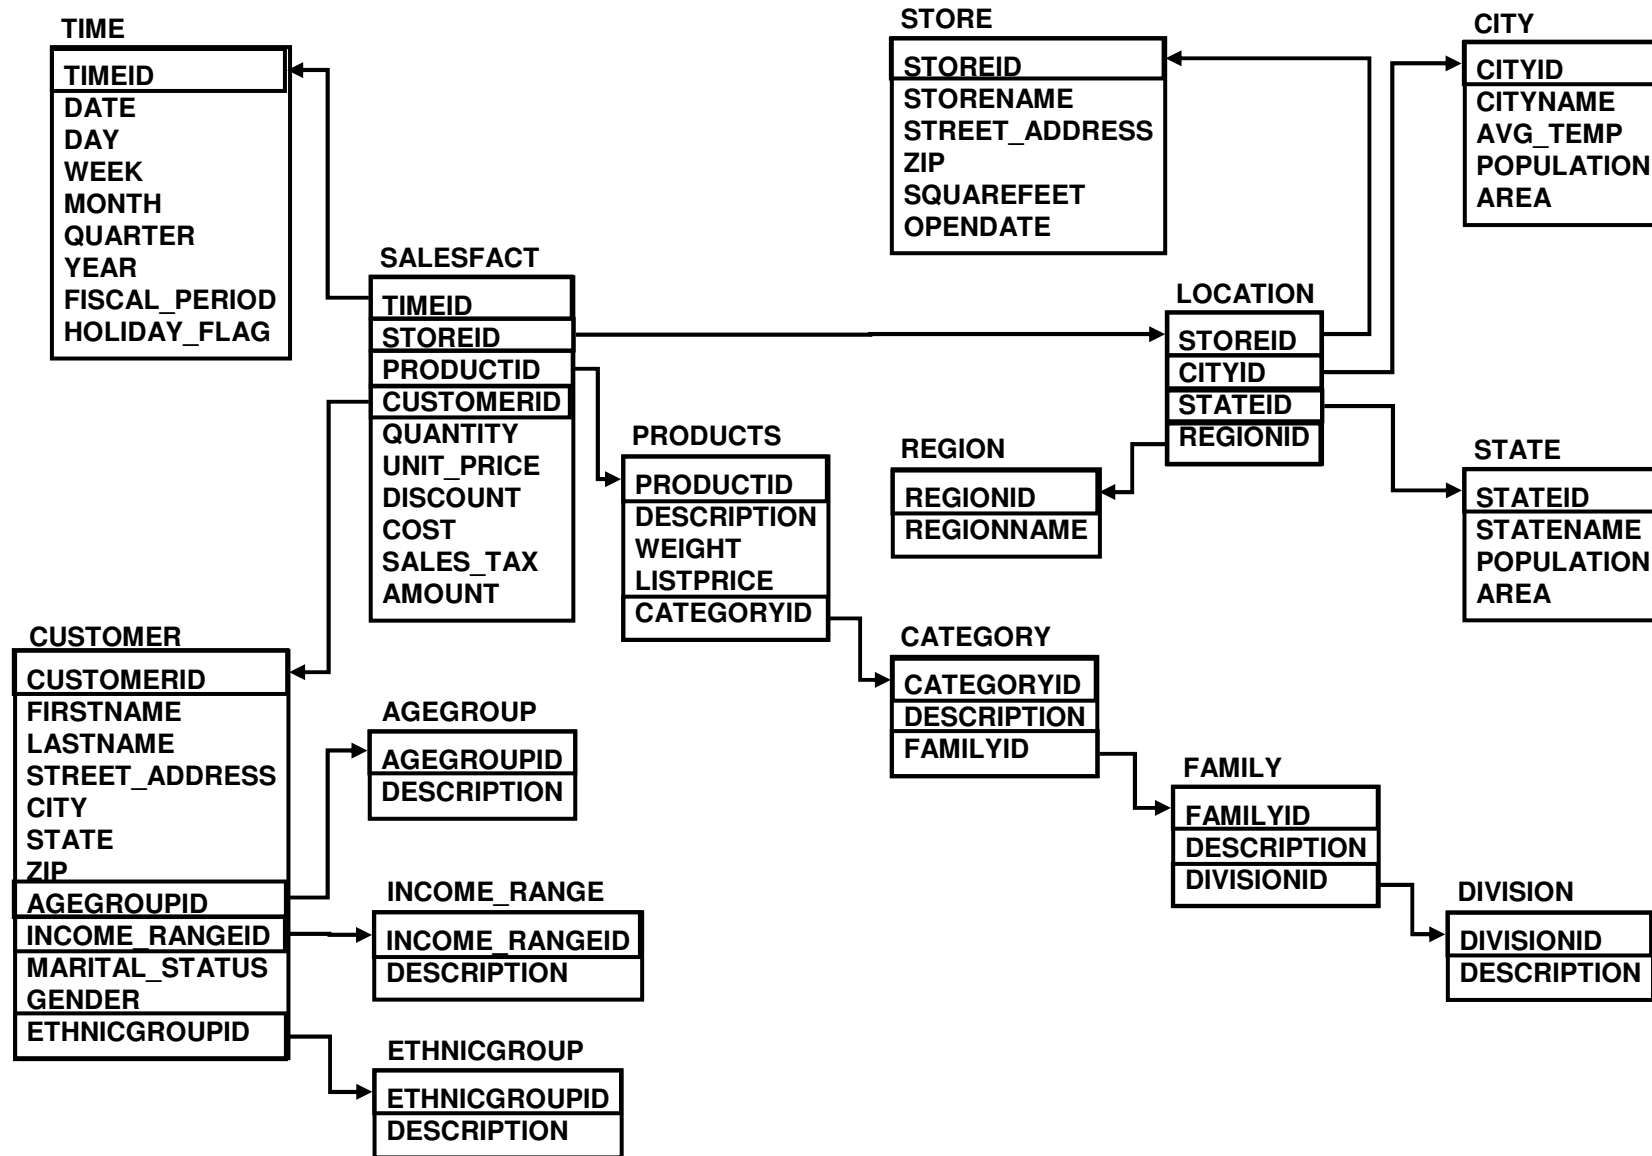


Degree of Parallelism = 4

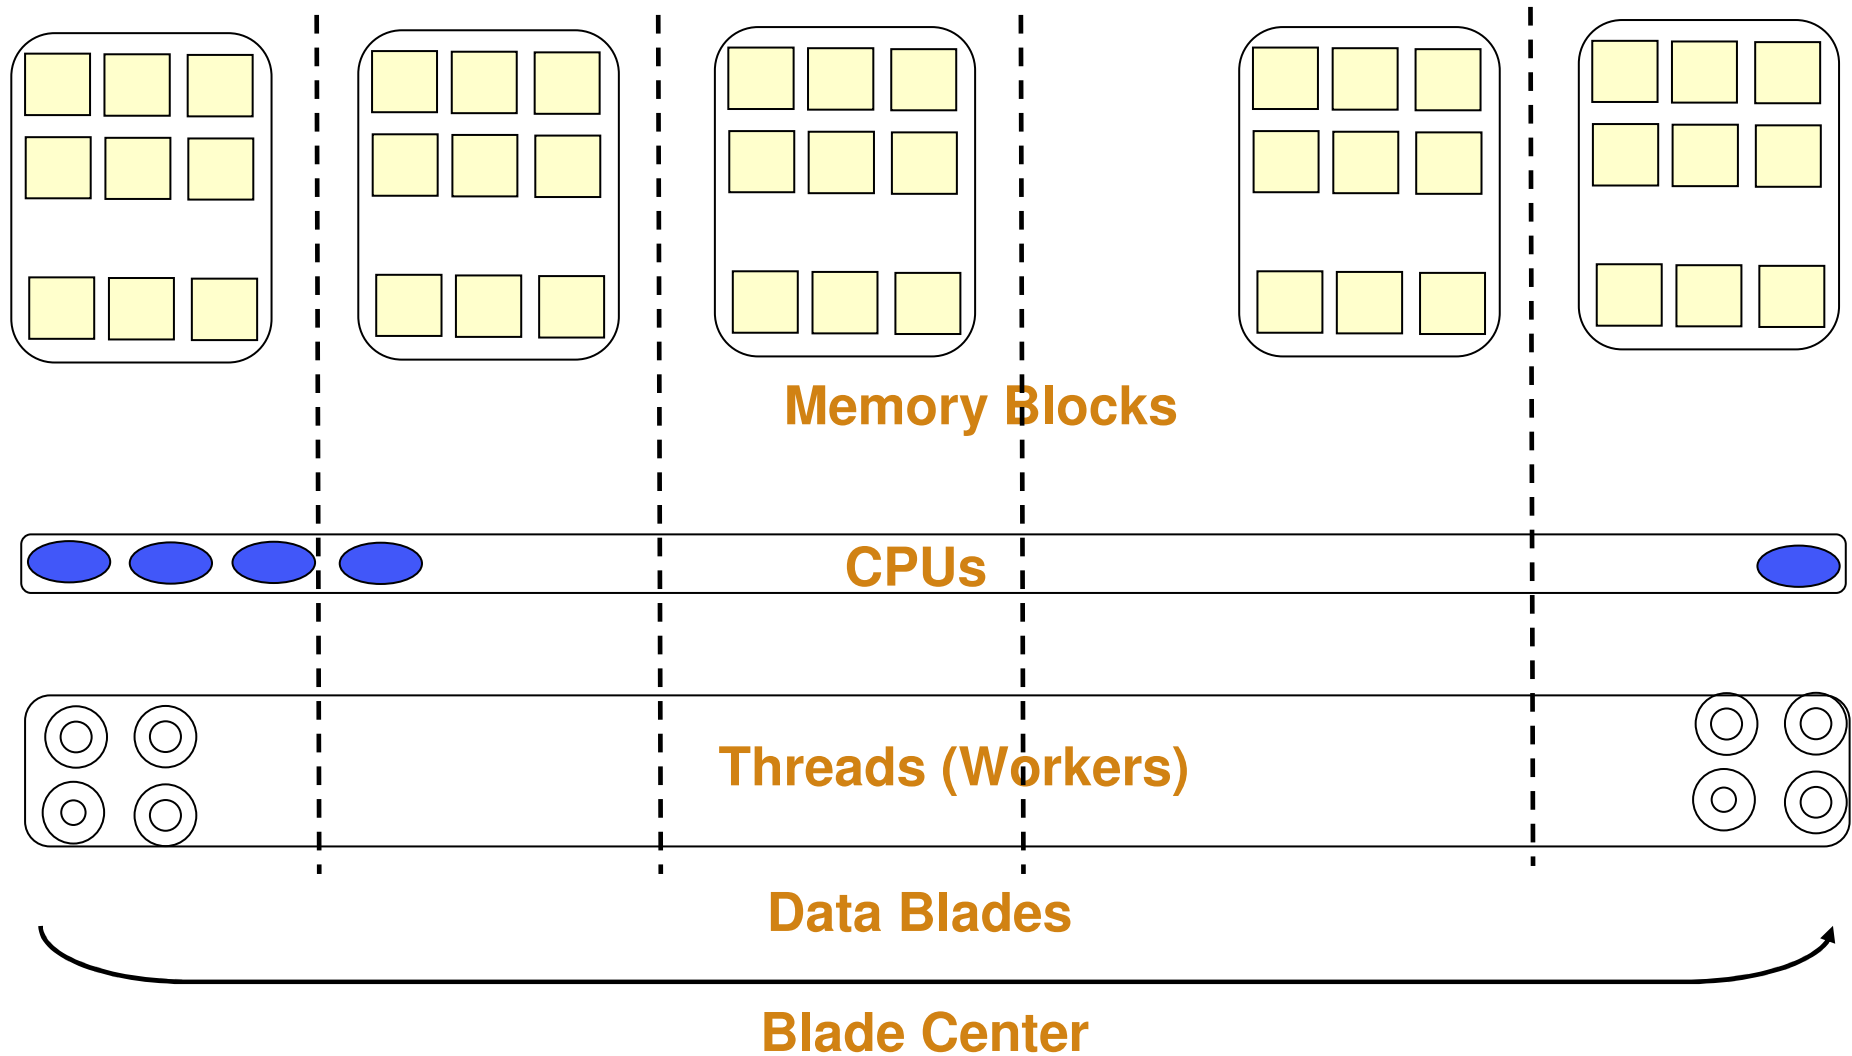
Workload Partitions



Exploitation of Massive Memory and Parallelism

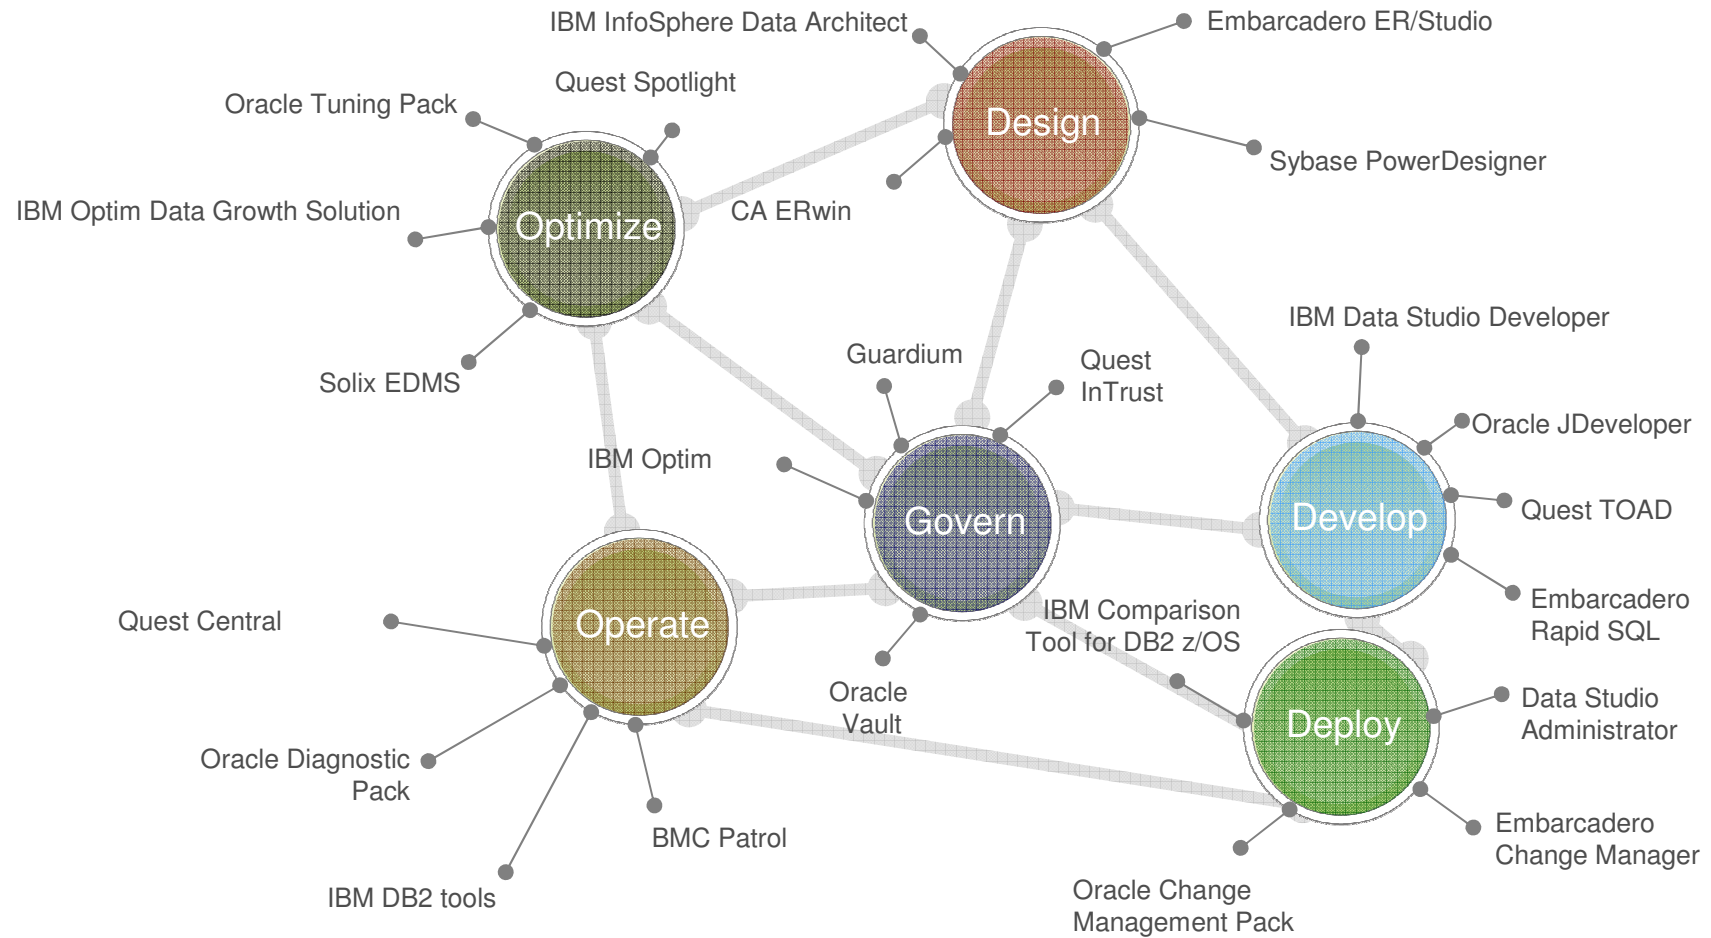


Exploitation of Massive Memory and Parallelism



What do Businesses Have?

A Collection of Disparate, Single-Purpose Products



What do Businesses Need?

An integrated environment to span today's flexible roles

- **Manage data throughout its lifecycle**
 - ▶ From design to sunset
- **Manage data across complex IT environments**
 - ▶ Multiple interrelated databases, applications and platforms
- **Facilitate cross-functional collaboration**
 - ▶ Within IT
 - ▶ Among Line of Business, Compliance functions
 - ▶ Across disparate skill sets
- **Optimize business value**
 - ▶ Respond quickly to emerging opportunities
 - ▶ Improve quality of service
 - ▶ Reduce cost of ownership
 - ▶ Mitigate risk



Introducing Integrated Data Management

An integrated, modular environment to design, develop, deploy, operate, optimize and govern enterprise data throughout its lifecycle on multiple platforms

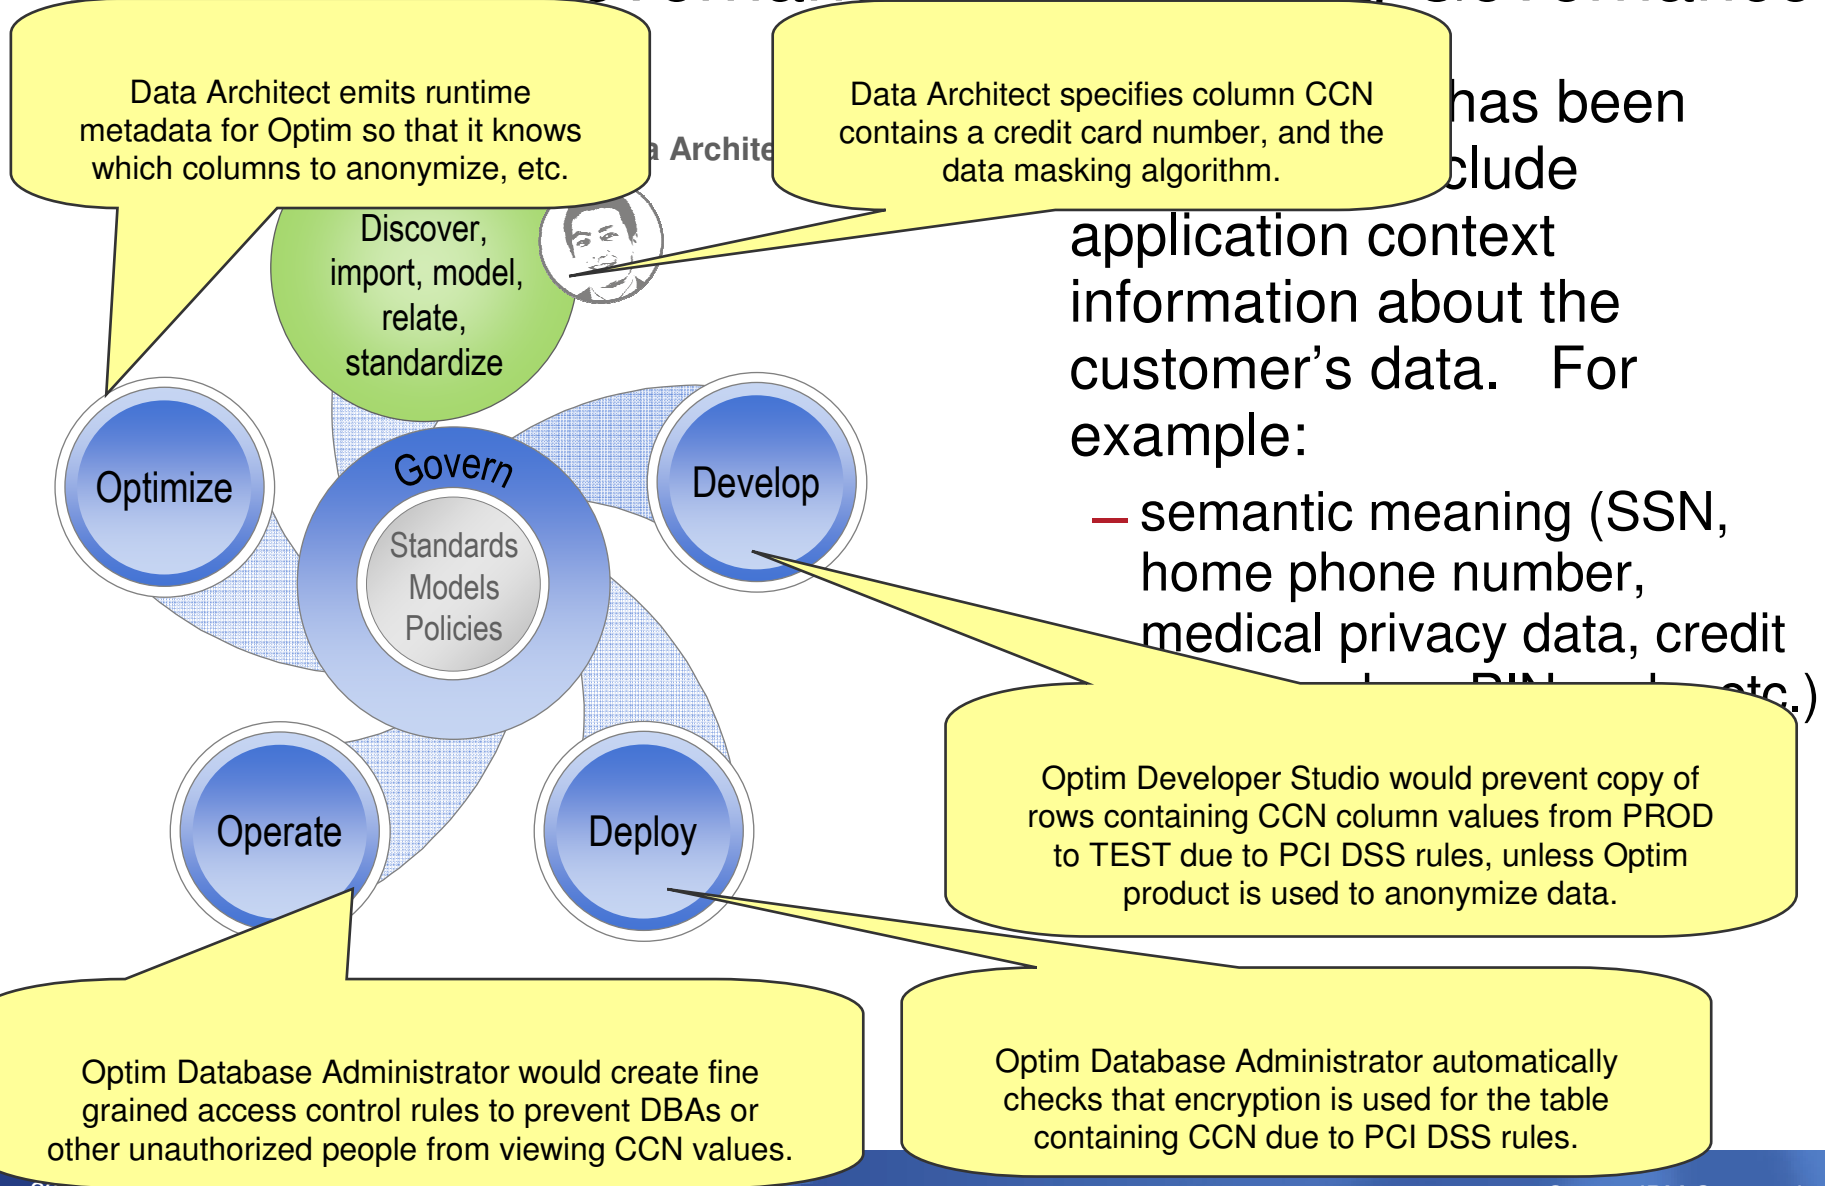


Enabling organizations to more efficiently and effectively

- **Respond to emergent, data-intensive business opportunities**
- **Meet service level agreements for data-driven applications**
- **Comply with data privacy and data retention regulations**
- **Grow the business while driving down total cost of ownership**



Model-driven Governance – Automating Governance



THANK
YOU

The image features the words "THANK YOU" in large, 3D, light blue letters. Each letter is filled with a different photograph of a person. The 'T' shows a man in a white shirt and orange tie. The 'H' shows a woman in a green top. The 'A' shows a man with a green face. The 'N' shows a woman with a blue patterned top. The 'K' shows a man with glasses. The 'Y' shows a man in a white shirt. The 'O' shows a man in an orange shirt. The 'U' shows a woman in a green top.