

Smarter Workload Optimization – IBM Software And POWER Unleashed




Analytics and OLTP On POWER – DB2 Engine Optimized for Smarter Computing

IBMDiscoveryDays2011

Copies of Today's Presentations:

<http://www.ibm.com/developerworks/offers/techbriefings/details/power.html>



Efficient Data Processing And Analytics To Handle Exploding Demand

IBM Software


Designed to optimize IBM servers and storage for analytics and data management



IBM Hardware Capabilities

IBM Power Systems, IBM Storage Systems – plus purpose-built innovations like TurboCore, DB2 pureScale, Smart Analytics, and Easy Tier

Explosive Growth of Data



IBM Analytics and OLTP

Drive Business Success
Seize Opportunity!

Data Processing Workloads Have Different Characteristics

ONLINE TRANSACTION PROCESSING (OLTP)



- Up to thousands of users
- Queries with quick response
- Smaller databases (up to terabytes)
- Mostly random I/O

Data Warehouse & Analytics



- Fewer users
- Complex queries and reports
- Large databases (up to petabytes)
- Mostly sequential I/O

Different Requirements Demand Different Optimizations For Each Type Of Workload

ONLINE TRANSACTION PROCESSING (OLTP)



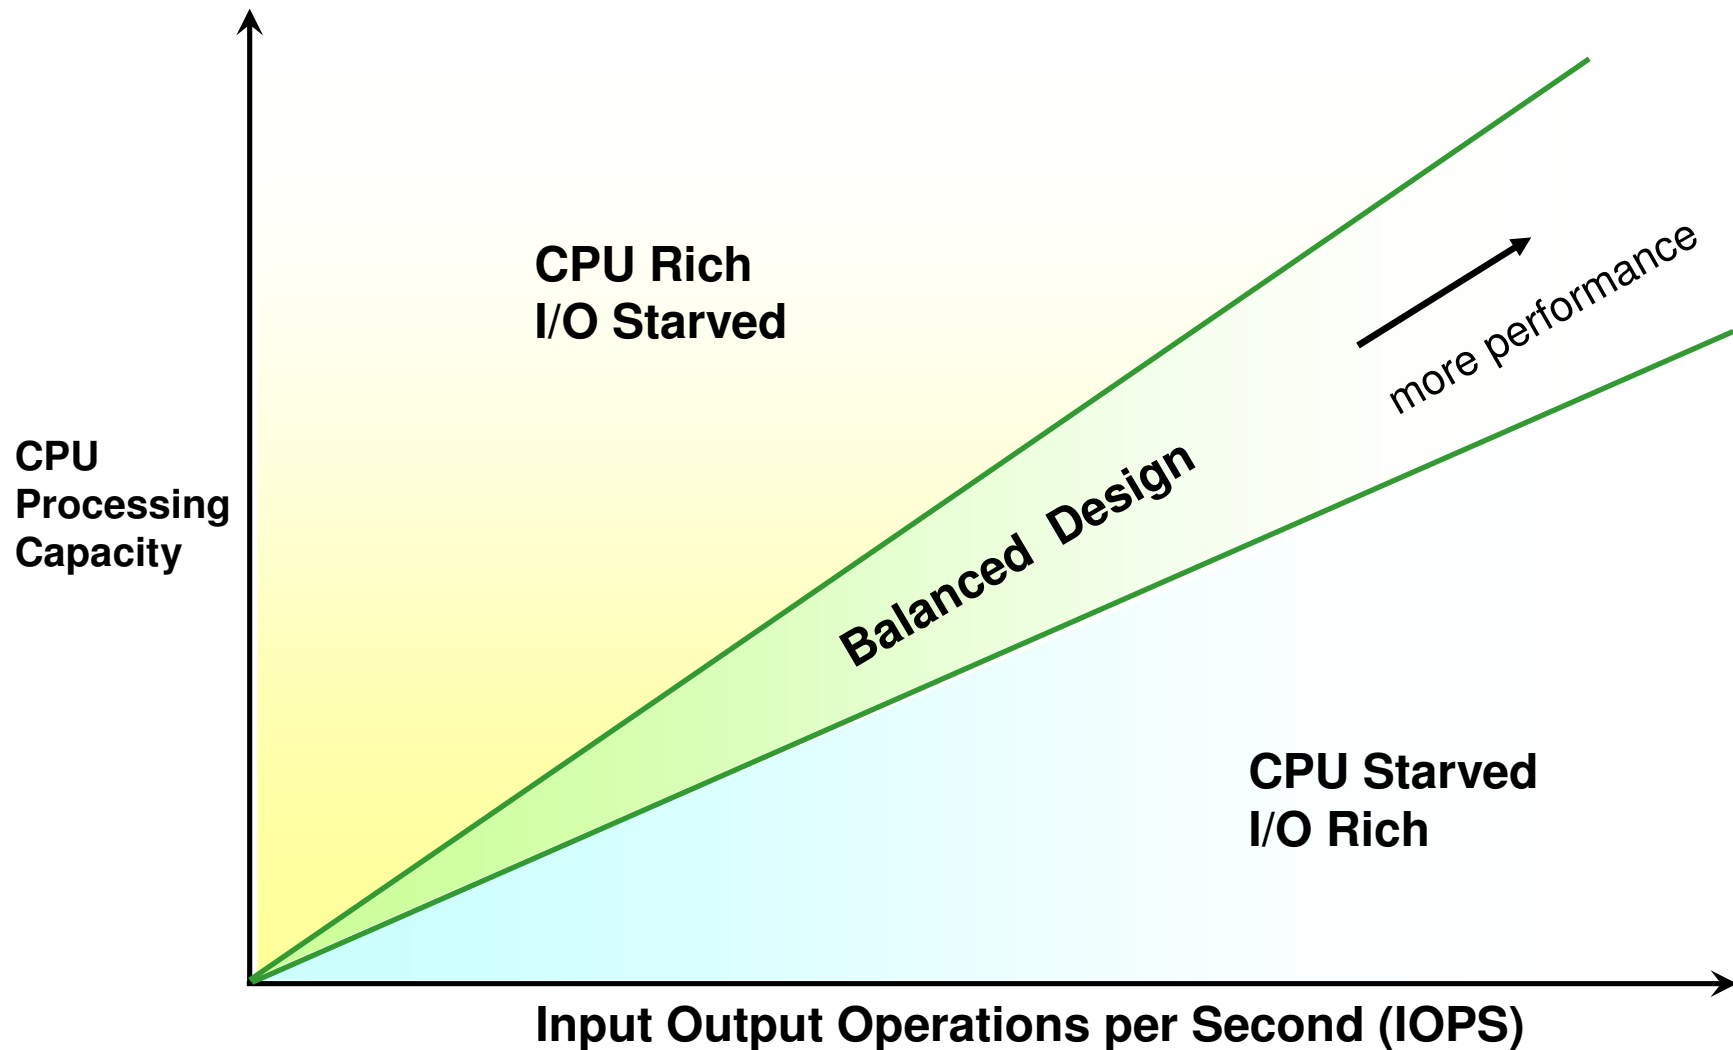
IBM Power Systems, DB2
and Storwize V7000

Data Warehouse & Analytics



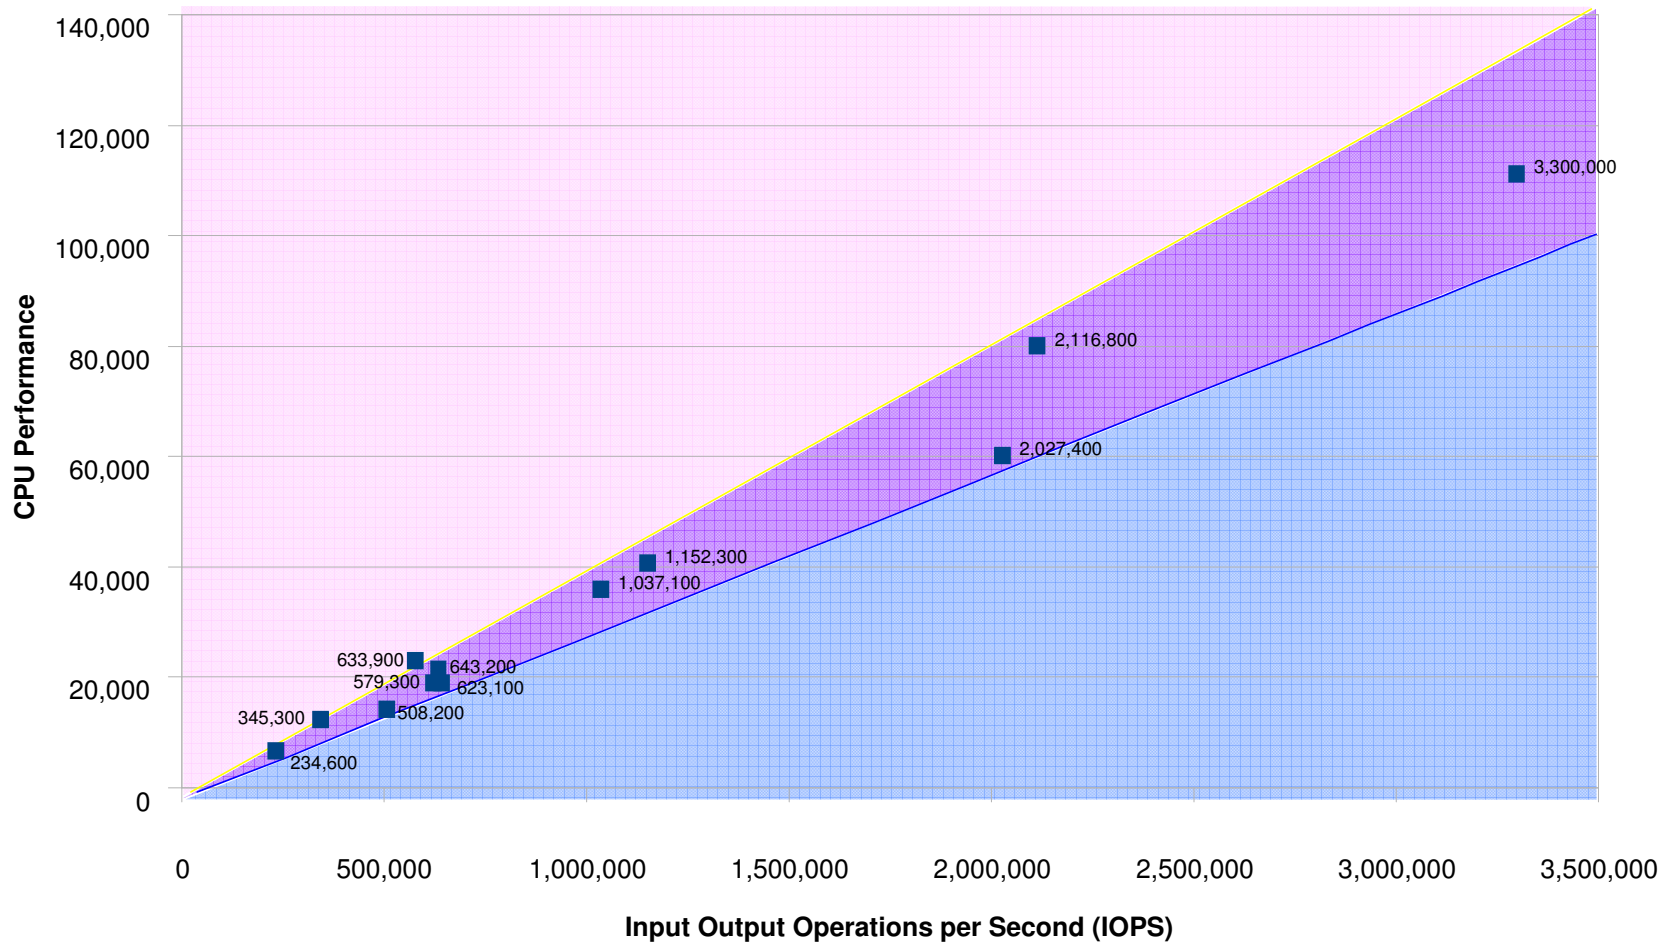
IBM Smart Analytics System 7700

Database OLTP Performance Is Bound By Two Key Capabilities



Attributes Of Balanced Servers, Top 30% Of TPC-C Benchmark Servers

Balanced Systems - Top 30% TPC-C
CPU and IOPS used



DB2 Is Optimized For POWER Architecture

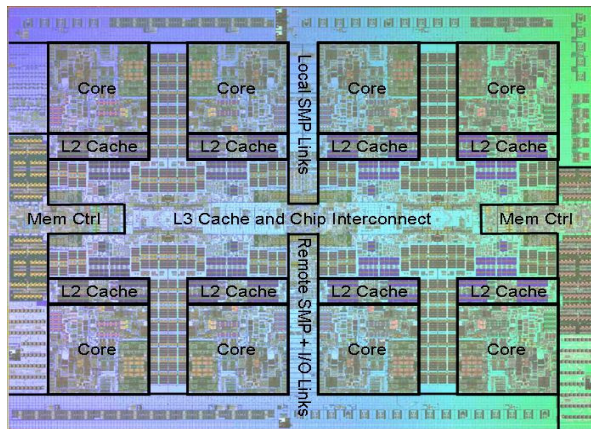

- DB2 can scale by effectively utilizing the multiple threads available in Power SMP servers
- DB2 is tuned for the POWER architecture
 - ▶ Supports AIX 64KB, 16MB and 16GB large page sizes for buffer pools
 - ▶ Leverages POWER Hardware Decimal Floating Point
 - ▶ DB2 is integrated to work with AIX Workload Management (WLM)
- Storwize V7000 Solid State Storage with Easy Tier delivers more I/O bandwidth

Deep integration of software/hardware/storage is needed to take advantage of all available resources

POWER7 Hardware Can Be Optimized For Data Intensive Workloads



MAX CORE

- 8 cores share common 32MB of L3 cache
- Clock rate 3.8 GHz
- CPU intensive workloads have up to **20%** better price/performance running like this




TURBO CORE

- 4 cores share same 32MB of L3 cache (turn off 4 cores)
- Clock rate 4.1 GHz (thermal reduction)
- Data intensive workloads have up to **7%** better price/performance running like this



Optimized Database Performance – DB2 On Power 780 With TurboCore

**DB2 efficiency leads in
Performance Per Core which drives down cost.**



IBM DB2 can do more than 8x more transactions per core than Oracle

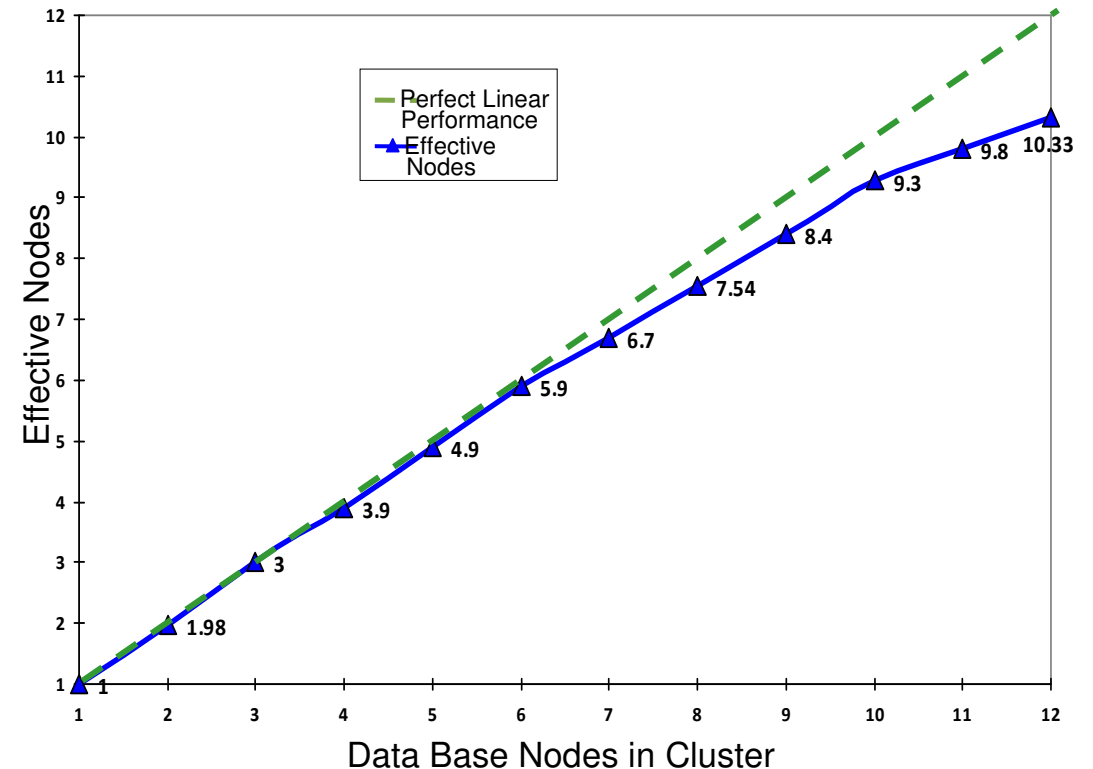

8 to 1

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC). Data current as of 2/10/2011
All results available at http://www.tpc.org/tpcc/results/tpcc_results.asp?print=false&orderby=submitted&sortby=desc,

DB2 Clustering Provides Even More Processing Capacity

- DB2 pureScale is a cluster of DB2 servers supporting one logical database
- Design based on proven mainframe Parallel Sysplex design to achieve near linear scaling

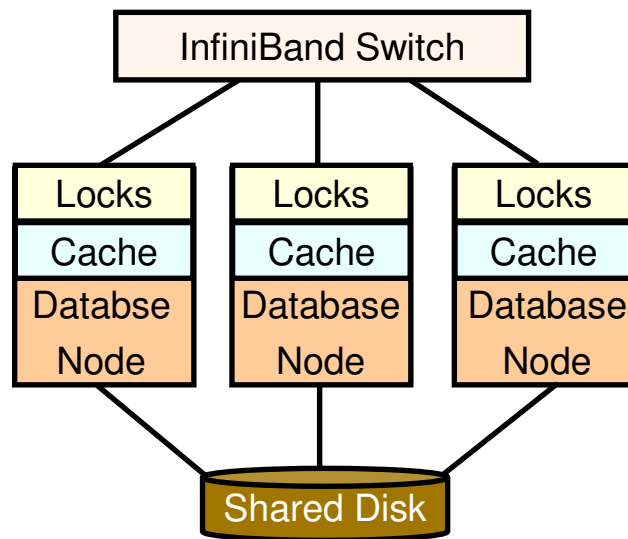
High-speed centralized lock and cache management




Source: IBM Software Group Internal Study

Without Application Changes And Extensive Database Tuning, Oracle RAC Does Not Scale

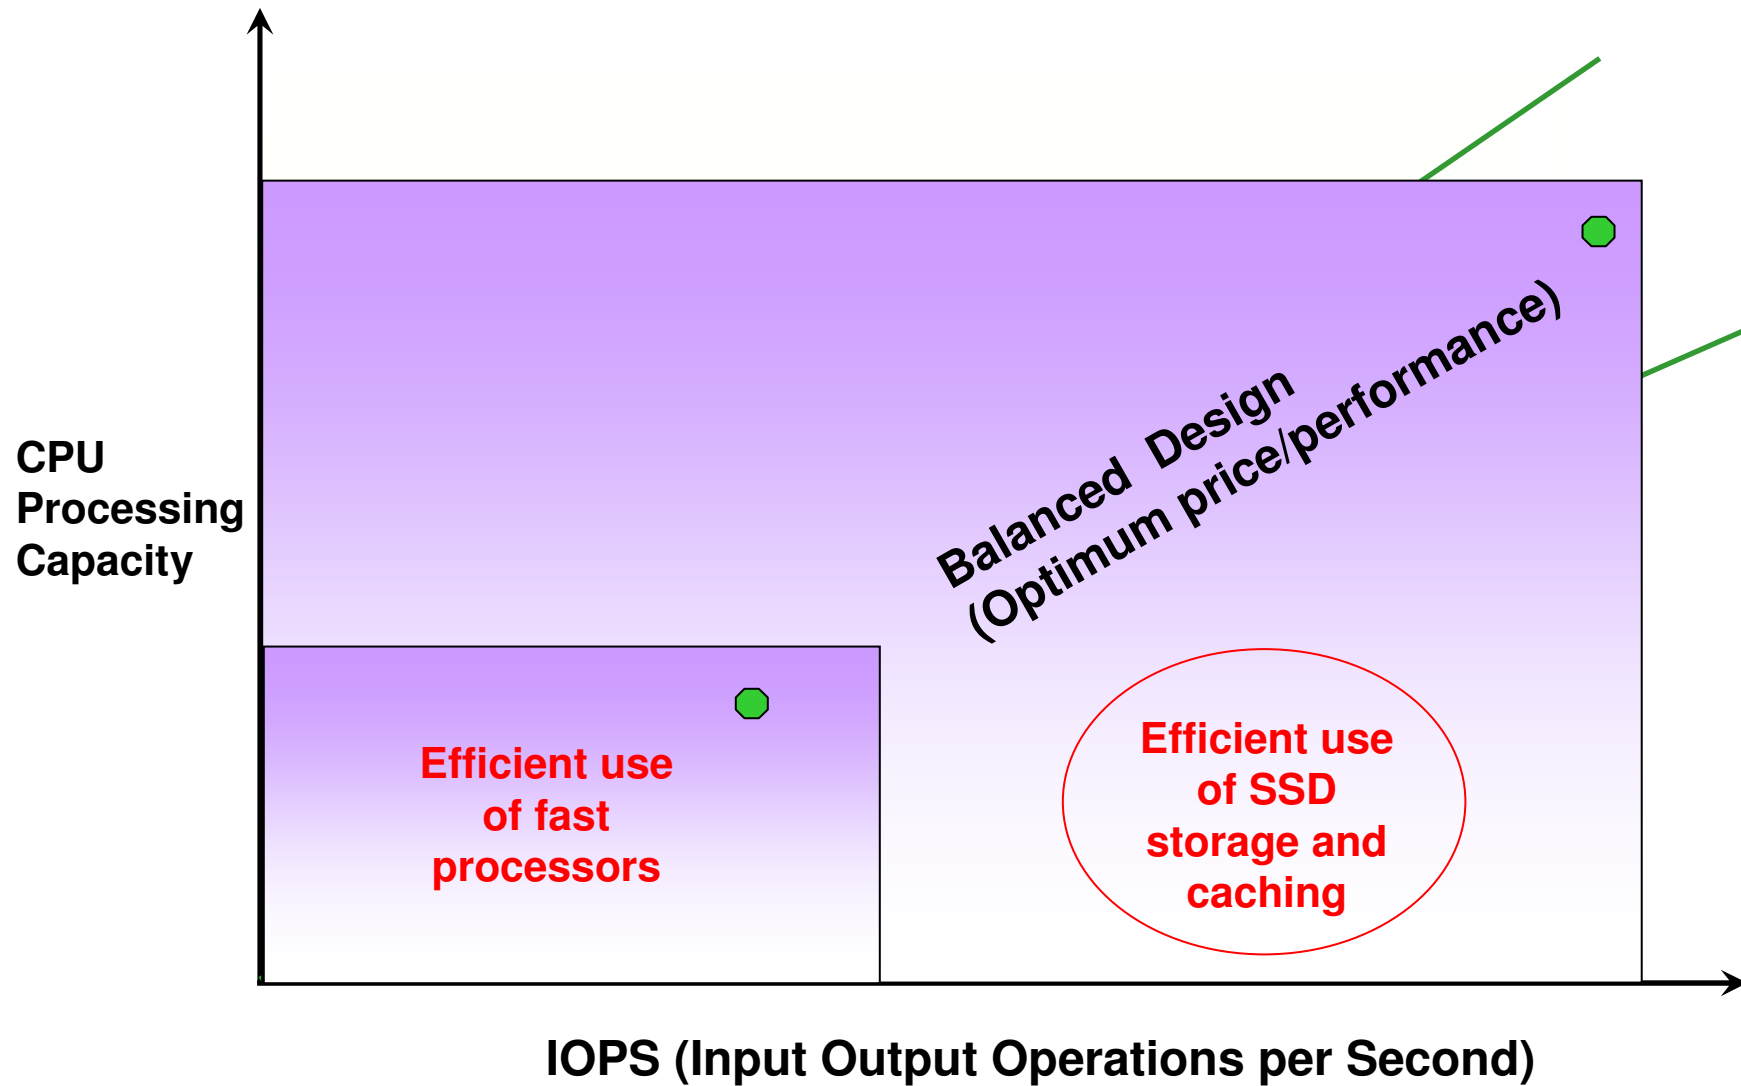
Oracle RAC Distributed Design



Inefficient distributed locking and buffer management limits scaling



More Processing Power Requires More IOPS



Solid State Drives Are Revolutionizing Storage Performance

Hard Disk Drives

Solid State Storage

SATA /SAS interface

SAS interface

SAS interface

PCIe interface

200 IOP's

300 IOP's

45,000 IOP's

100,000 IOP's

\$0.73/GB

\$3.66/GB

\$88.00/GB

\$25.93/GB

\$7.5/IOPS

\$3.7/IOPS

\$0.59/IOPS

\$0.08/IOPS

SLC Durability

MLC Durability



IBM 2TB SATA
3.5" LFF HDD
Cost: \$1499.00



IBM 300GB 10K SAS
2.5" SFF Slim-HS HDD
Cost: \$1099.00

IBM 300GB
2.5" SFF SAS SSD
Cost: \$26,399

IBM 320GB HIGH IOP MS
CLASS SSD PCIe ADPT
Cost: \$8,299

Easy Tier In Storwize V7000 Automatically Optimizes Use Of SSD


- Migrates data extents between SSD and HDD in same pool
 - ▶ Automatic hotspot detection
- Virtualized SSD is shared across all workloads using the pool
- More cost effective use of SSD versus ad hoc dedicated assignment
- Transparent to applications, no code changes required



Example: Complex database transactional workload

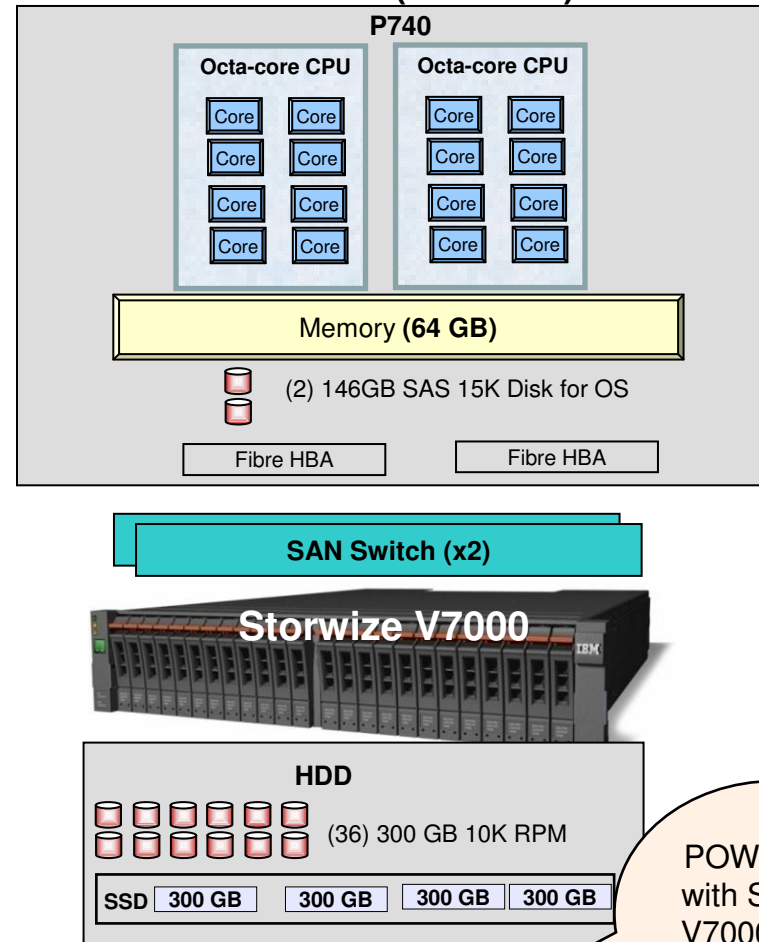
System Comparison Of Competitor Versus IBM POWER 740 + DB2 + Storwize V7000

**Competitor Database Machine
(2 DBMS Nodes, 3 Storage Devices)**



3YR TCA = \$2,230,481

**IBM POWER 740
16 cores (3.55 GHz)**



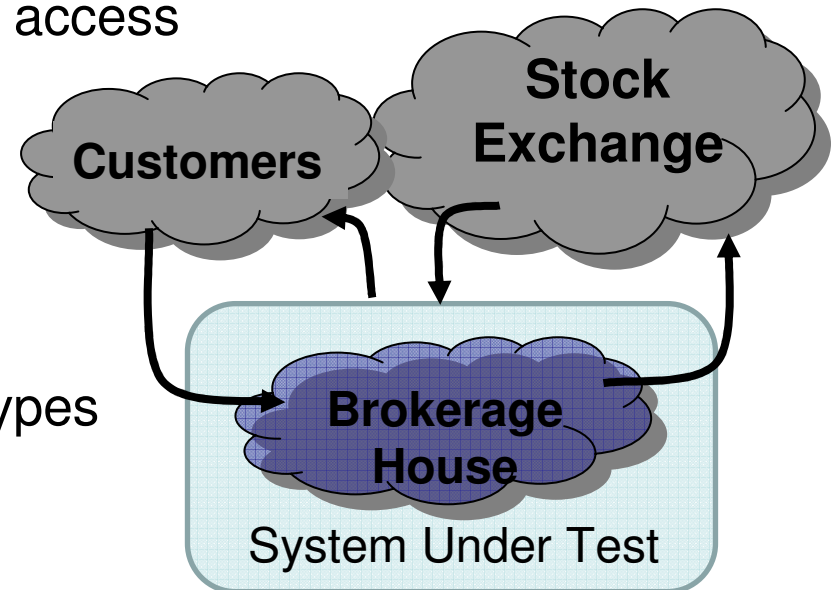
3 YR TCA = \$877,254

POWER 740
with Storwize
V7000 Costs
61% Less

3 year total cost of acquisition includes hardware, software, service & support. Based on US list prices, prices will vary by country


Brokerage On-Line Transaction Processing Workload

- More realistic OLTP workload models brokerage requirements
 - ▶ Based on TPC-E but not a full implementation – so tpc-e “like”
- Typical requirements
 - ▶ Sub-second response times
 - ▶ Multiple systems with simultaneous access
 - Customer/Client Access
 - System updates (ingest stream)
 - 3rd Party Access
- Data Application Architecture
 - ▶ Multiple tables, indexes, and data types
 - ▶ Requires referential integrity
 - ▶ Complex transactions
 - ▶ Multiple interacting subsystems



IBM Power 740 + DB2 + Storwize v7000 Delivers Better OLTP Throughput And Response Time


DB2 on Power 740 with Storwize V7000 vs. Competitor Transaction Throughput



(Higher is Better)

**1.7x
more throughput**

DB2 on Power 740 with Storwize V7000 vs. Competitor Transaction Response Time




(Lower is Better)

**2.0x
faster response time**

Performance numbers may vary based on workload profiles.

DB2 on Power 740 Delivers Even More Dramatic Savings

DB2 on Power 740 with Storwize V7000 vs.
Competitor Cost Per Transaction per second



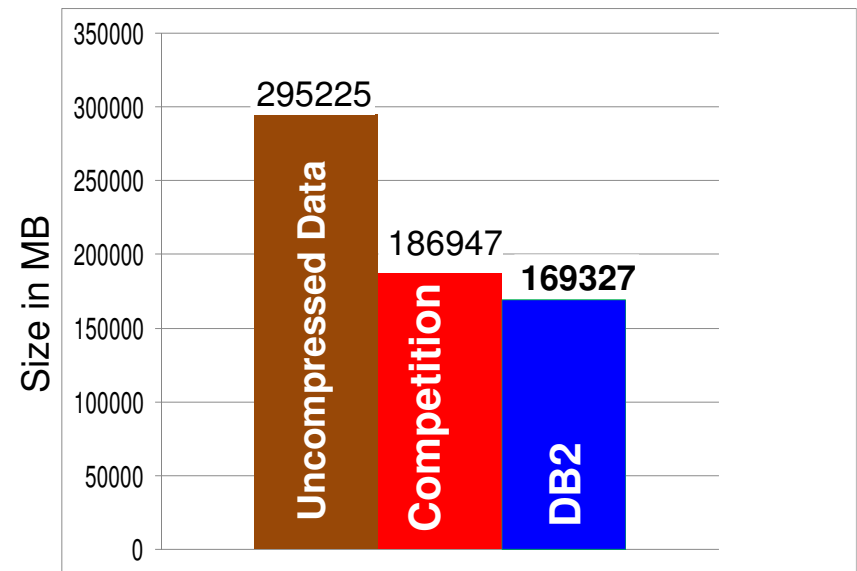
(Lower is better)

77%
lower cost
per transaction
per second
than competitor

Performance numbers may vary based on workload profiles.
3 year total cost of acquisition includes hardware,
software, service & support.
Based on US list prices, prices will vary by country.

... And DB2 Compression Gives An Additional 20% Performance Boost

- DB2 saves more space than the competition
- Compression can help performance by reducing I/O data traffic
- Enabling DB2 compression improved performance an additional 20% at no extra cost



DB2, Power, And Storwize Beat The Competitor In OLTP

Evaluation Criteria	IBM	Competitor
DB2 and DB2 pureScale are optimized for Power architecture	Yes	No
DB2 pureScale scales much better than Oracle RAC at a lower cost of acquisition	Yes	No
DB2 on Power with Storwize V7000 provides better throughput and response than competitor	Yes	No
DB2 on Power with Storwize V7000 provides a lower cost per transaction (61% lower) than competitor	Yes	No
DB2 compresses data better than competitor at no additional cost	Yes	No

Characteristics Of Business Analytics Workload

ONLINE TRANSACTION PROCESSING (OLTP)



- Up to thousands of users
- Queries with quick response
- Smaller databases (up to terabytes)
- Mostly random I/O

Data Warehouse & Analytics



- Fewer users
- Complex queries and reports
- Large databases (up to petabytes)
- Mostly sequential I/O

Using Cognos Business Analytics To Improve Business Results

If we can identify our risky mortgage assets, we can work to remove them from our books.



Service Oriented Finance CIO

We can identify risky mortgage customers by watching their activities in other business areas

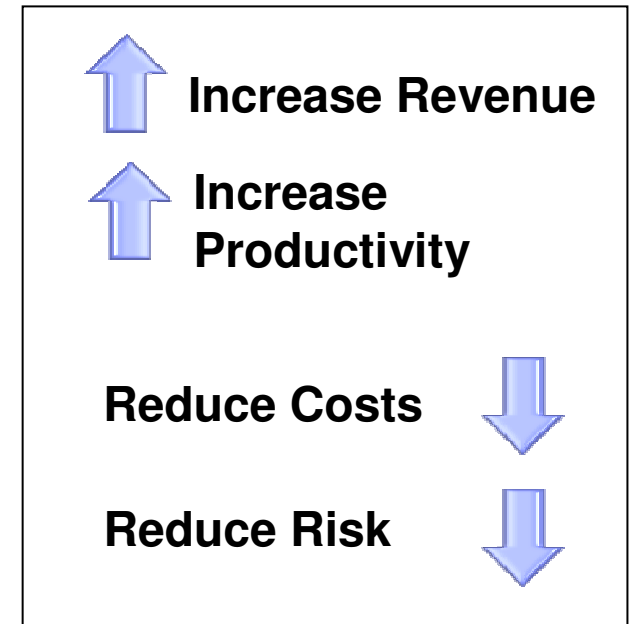
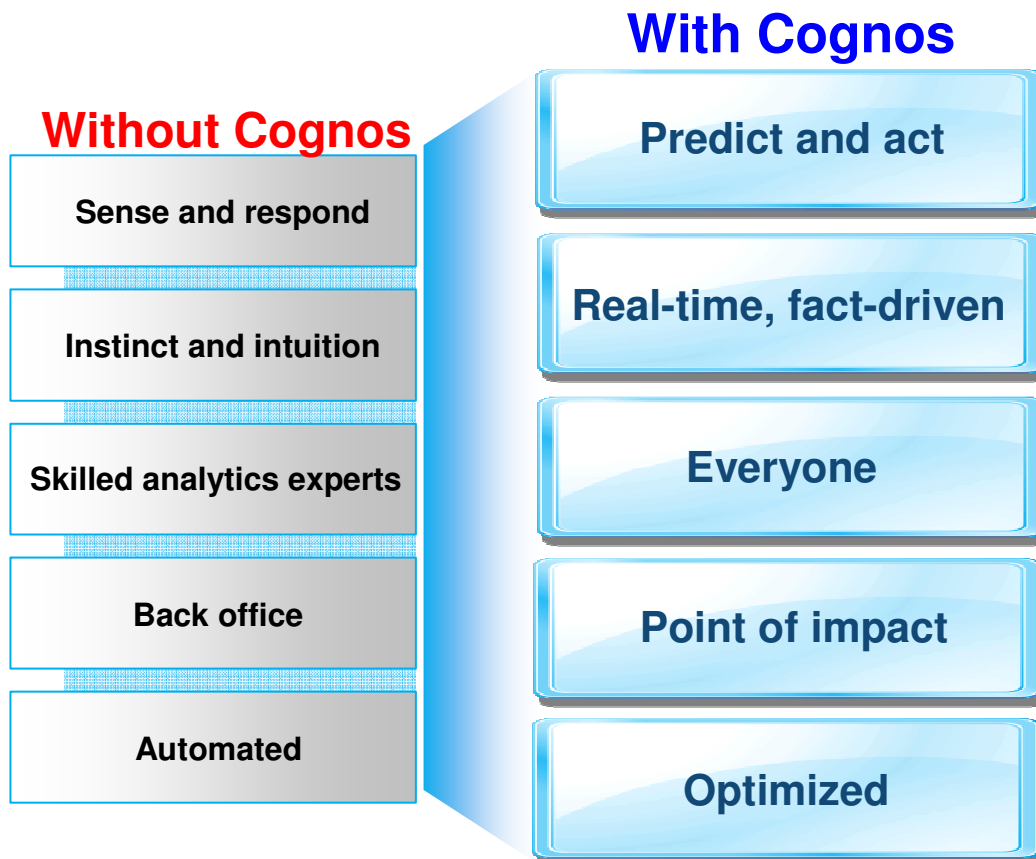
- ▶ Bounced checks
- ▶ Missed credit card payments
- ▶ Missed loan payments.



Mortgage Line of Business VP

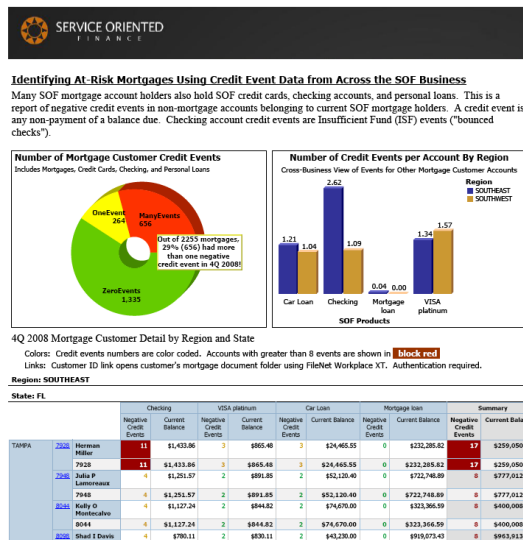
Use Cognos To Unlock The Useful Insights In Your Data

- Predictive insight
- Consistent measurement of business performance
- Anticipate and explore new opportunities



DEMO: Identify Risky Mortgage Accounts Using Cognos Business Intelligence

1. Show report generated in Cognos Report Studio in PDF format
2. Report identifies high-risk mortgages by looking at negative credit events in customers other SOF accounts (CC, Checking, etc...)
3. Report uses both structured and unstructured data (link to mortgage data stored in FileNet)



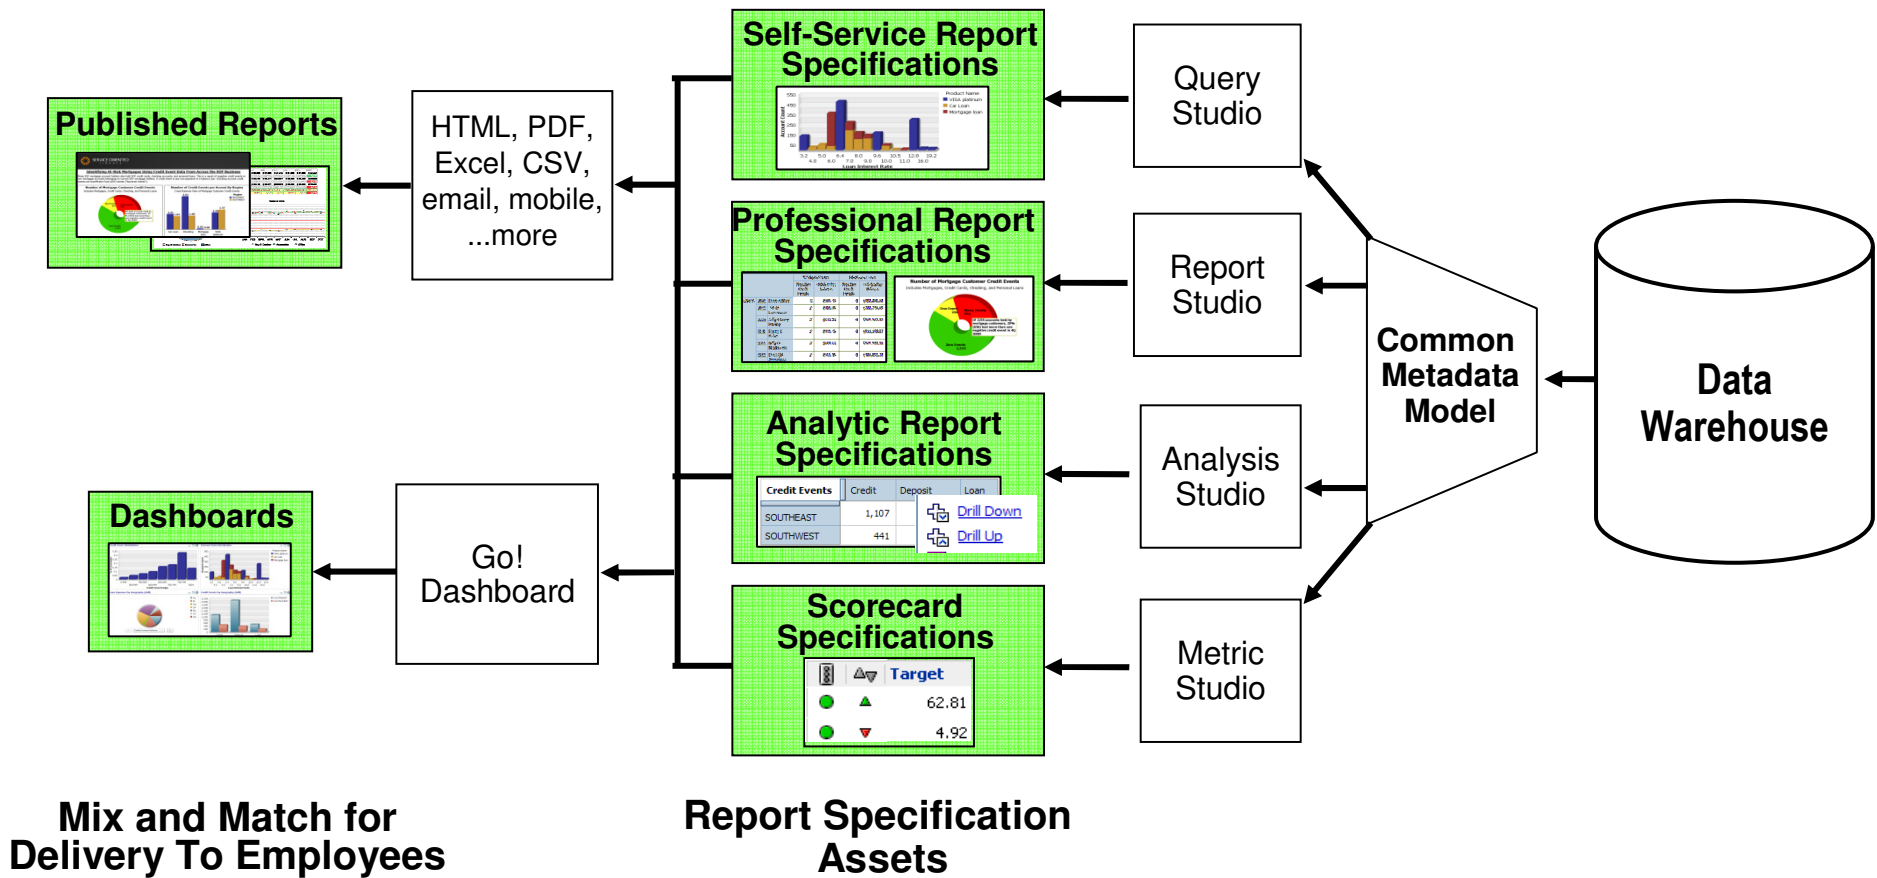
That report is just what the executives wanted!



Mortgage Line of Business VP

- At risk customers are identified and contacted to refinance
- Risky mortgages can be sold

Business Analytic Assets Created By Cognos



IBM Smart Analytics System 7700 – A Complete Workload Optimized Analytics Solution

- A pre-integrated and tuned solution containing Cognos and InfoSphere Warehouse software, Power hardware, and IBM storage

Analytics

- ☑ Optional Cognos 8 BI
- ☑ Cubing Services
- ☑ Optional SPSS 19
- ☑ Text Analytics & Data Mining

Data Warehouse Software

- ☑ InfoSphere Warehouse
- ☑ Advanced Workload Management
- ☑ Tivoli System Automation

Hardware/OS and Storage

- ☑ AIX 6.1
- ☑ IBM Power 740
- ☑ IBM System Storage DS3500
- ☑ SSD standard in all configurations



Faster Time To Value

Deploy in days instead of months
Eliminates installation, integration,
configuration, and tuning requirements

Costs less than building it yourself

Less staff and expertise required to
implement, tune, and maintain

Unique offering

Modular for flexibility and growth
No competitor offers such an
integrated analytics solution

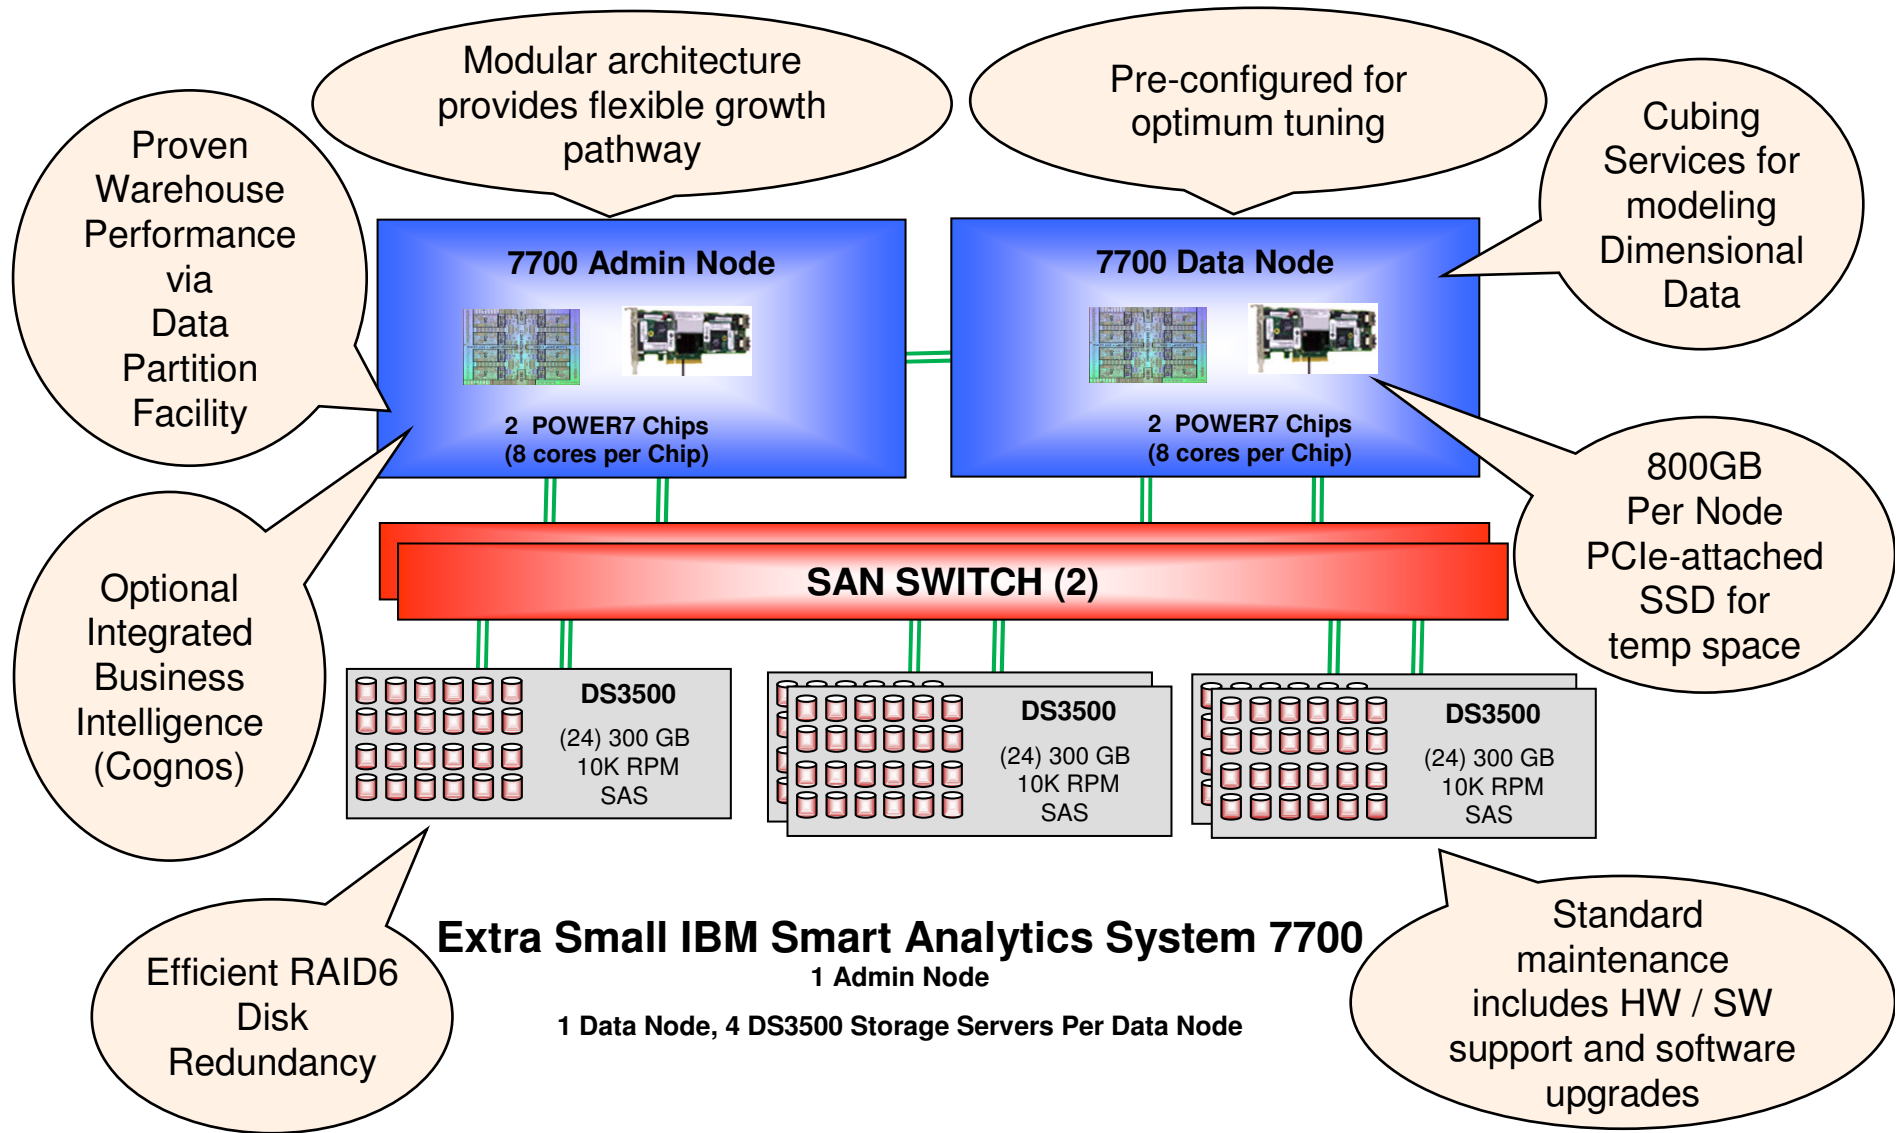
Find out more: <http://www.ibm.com/smart-analytics-system>

IBM Smart Analytics System 7700 Configurations

7700 Capacity Sizes	XS	S	M	L	XL	XXL
# Data Modules	1	2	3	6	10	20
Spinning Disk (TB)	28.8	57.6	86.4	172.8	280	560
Space available for index, temp, logs, system (TB)	9.35	18.8	28.1	56.3	98	196
Solid State Device for temp (standard) (TB)	0.7	1.4	2.1	4.2	7	14
Additional Solid Device for temp (optional) (TB)	Up to 4.2	Up to 8.4	Up to 12.6	Up to 25.2	Up to 40.2	Up to 80.4

Each data module is a **Power 740 server** with **16 cores**, 128GB memory, 96 300GB disk drives

IBM Smart Analytics System 7700 Capabilities





System Comparison Of Competitor With Smart Analytics 7700 R1 XS


Competitor V2 Database Machine 1/4 Rack Size (2 DBMS Nodes, 3 Storage Cells)

IBM Smart Analytic System 7700 R1 Extra Small Size (1 Admin Module, 1 Data Module)

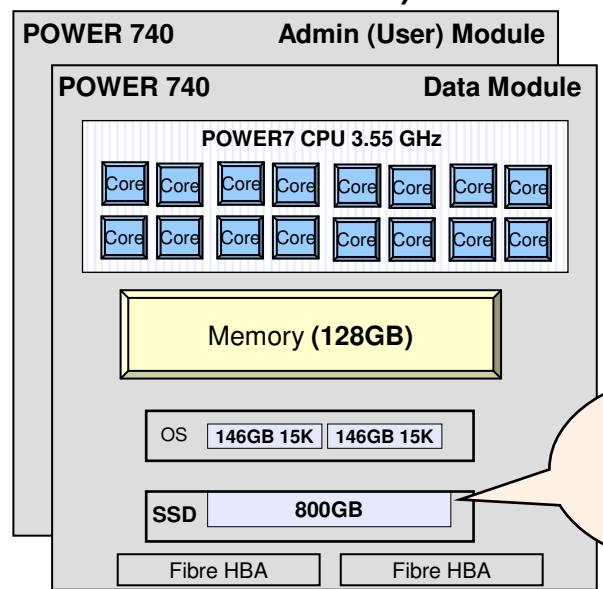
Oracle RAC DBMS




Sandfly Storage Cells



SSD used as cache



SSD used for temp space



3YR TCA = \$2,857,500

3 YR TCA = \$2,817,500

Introducing BI Day Benchmark

■ BI Day Benchmark

- ▶ Based on a typical Cognos business analytics workload
- ▶ Typical star schema data warehouse of 80 tables, ranging from 100GB to 10TB
- ▶ Serial or concurrent test modes
- ▶ Represents real world workloads

■ Industry benchmarks like TPC-H have lost relevance for customers

- ▶ Thousands of hours have been spent tuning TPC-H benchmarks (Real customers can't afford to do this)
- ▶ Large benchmark results are achieved using unrealistic and impractical hardware configurations
- ▶ TPC-H has only 8 tables and is not a star/snowflake schema

BI Day – Serial Test Mode

1 user executes
complex reports

...then... executes
intermediate reports

...then... executes
simple reports




1 User



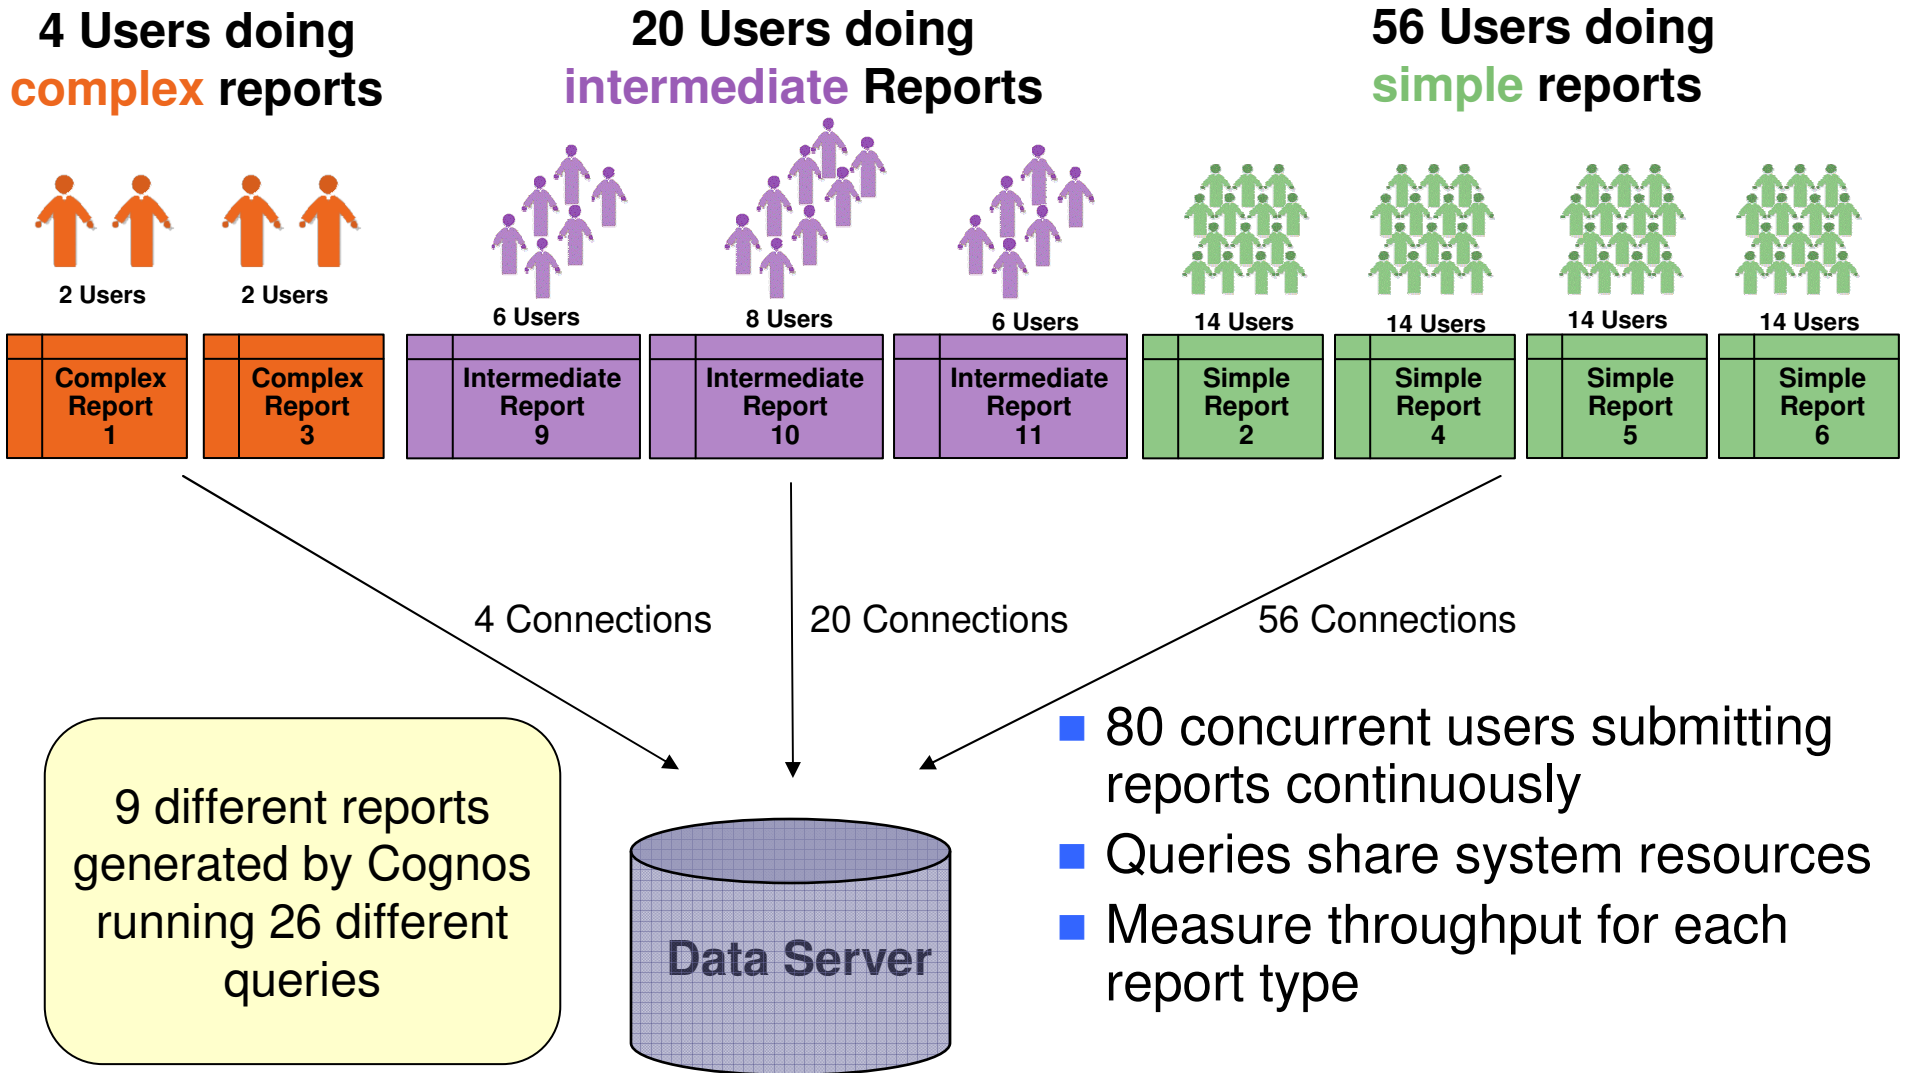
9 different reports
generated by Cognos
running 26 different
queries

Single
Connection



- Single user serial execution test
- Each query gets all system resources
- Measure elapsed time for each report


BI Day – Concurrent Test Mode



Note: Distribution of complex, intermediate, and simple workloads based on Forrester Research, Profiling the Analytic End User for Business Intelligence, 2004

Competitor Performance for Complex Reports Degrades When There Is Concurrent Work Going On


Competitor Complex Report Performance
Serial Execution versus Concurrent Throughput - 1TB



Performance numbers may vary based on workload profiles.

IBM Smart Analytics System 7700 Delivers More Throughput Concurrent Operational Reports


3.9X More Simple Reports



Reports Per Hour at 10 TB data size

(Higher is Better)


2X More Intermediate Reports



Reports Per Hour at 10 TB data size

(Higher is Better)

1.9x More Complex Reports



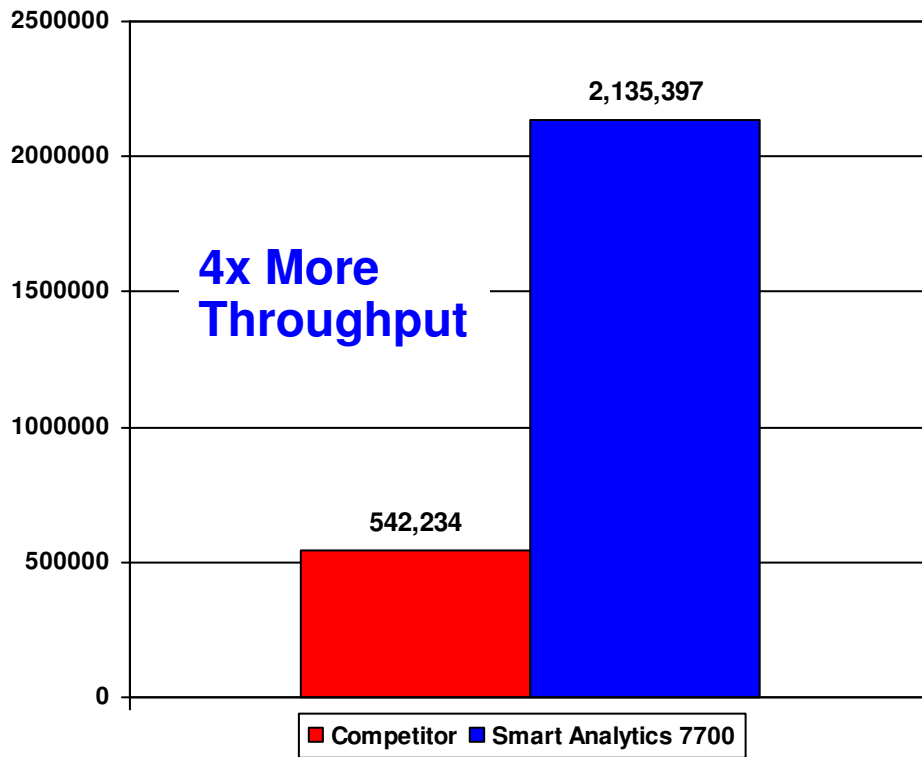
Reports Per Hour at 10 TB data size

(Higher is Better)

Performance numbers may vary based on workload profiles.


Result: IBM 7700 Delivers More Aggregate Throughput And Lower Cost Per Report For Concurrent Operational Reports

Report Throughput at 10 TB



(Higher Throughput is Better)

Cost Per Report at 10 TB



(Lower Cost is Better)

Performance numbers may vary based on workload profiles. 3 year total cost of acquisition includes hardware, software, service & support. Based on US list prices, prices will vary by country.

IBM Smart Analytics System 7700 Beat The Competitor In Business Analytics

Evaluation Criteria	IBM	Competitor
OLAP, Data Mining, Text Mining included	Yes	No
Business Intelligence Software (Cognos) Integrated	Yes	No
Proven architecture for scalability	Yes	No
100% availability even in case of server failure	Yes	No
1st year software maintenance included (All IBM Software)	Yes	No
Software license costs per disk drive (Included) - \$10,000 USD per Disk Drive for Oracle	Yes	No

Service Oriented Finance Needs To Find Good Future Customers

These reports are great! Can we take it a step further and be able to identify good future customers?



**Mortgage
Line of Business VP**

What you need is a Predictive Analysis tool. SPSS is an add-on to IBM Smart Analytics System that will help you do just that.




IBM

SPSS Helps Predict Future Events

Load


Data collection delivers an accurate view of customer attitudes and opinions



Load

Predict

Predictive capabilities bring repeatability to ongoing decision making, and drive confidence in your results and decisions



Act


Unique deployment technologies and methodologies maximize the impact of analytics in your operation

Deployment Technologies



DEMO: Automate Analysis Rules For Identifying Risky Customers Using SPSS Statistics

1. Capture data from Data Warehouse on DB2 into SPSS Statistics
2. Pre-process the data to create new attributes for quantifying negative credit events across different product lines and create a risk flag for mortgage
3. Run Comparison of Means and Decision Tree to discover rules for characterizing risky customers



Load Data



Platform

Deployment Technologies

We can now identify the conditions that lead to risky mortgages!



**Mortgage
Line of Business VP**

Summary – IBM Provides The Best Analytics And OLTP Solutions

- Different types of workloads have different data requirements which require different optimizations
- Oracle only offers Exadata as a “one size fits all” solution
- DB2 PureScale on Power Beats Exadata for Online Transaction Processing (OLTP) workloads
- IBM Smart Analytics System 7700 Beats Exadata for Business Analytics workloads
- Both provide a more scalable platform at a **lower cost** than Exadata
- IBM Smart Analytics System 7700 provides the most comprehensive Business Analytics solution in the industry
 - ▶ Includes Cognos and SPSS