



# A Case Study of Real-World Porting to the Itanium Platform

**Jeff Byard**

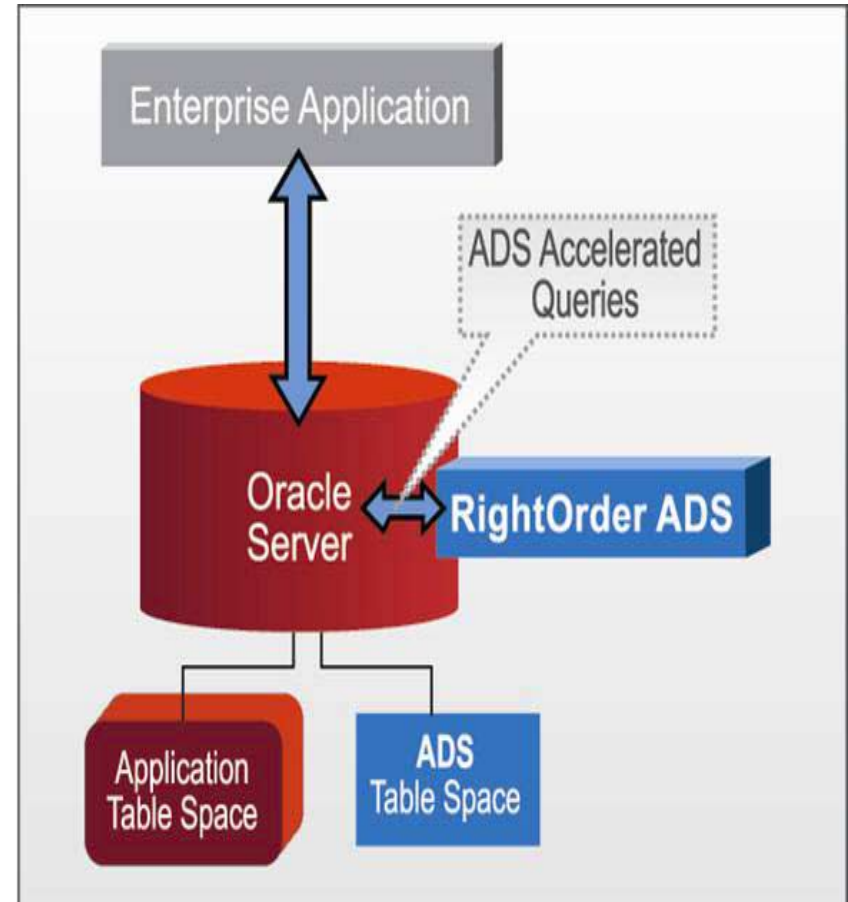
VP, Product Development  
RightOrder, Inc.

# Agenda

- RightOrder ADS Product Description
- Porting ADS to Itanium 2
- Testing ADS on Itanium 2


# RightOrder ADS product description

- Query accelerator for enterprise applications running on Oracle databases
  - Utilizes Index Fabric™ technology
  - Deployed as an Oracle Cartridge user defined index (ODCI)
- Transparent, simple deployment:
  - No change to existing applications
  - No change to schema / data
  - No additional hardware required
- Substitute and add to Oracle indexing capabilities with a faster, scalable, and richer in functionality index



# Typical performance improvement

- RightOrder ADS accelerates the most resource consuming queries for a typical business intelligence / data warehouse application



# Porting ADS to Itanium 2

- ADS is a server-side, high-performance, mission-critical product
- ADS is mostly Java (some PL/SQL)
- Itanium 2 JVM performance and stability are critical
- Tested version 1.4.1\_03 of the Itanium 2 JVM

# Components of ADS

- **Server**

- Primary component; evaluates query constraints and joins using ADS indexes
  - Communicates with Cartridge via TCP/IP and RMI
  - Communicates with Oracle proper via JDBC
- All Java

- **Cartridge**

- Interface between ADS server and Oracle
- Java and PL/SQL

# Components of ADS (contd.)

- **Proxy**

- Intercepts traffic between the client(s) and the Oracle listener, at the packet level
- All Java

- **Installer**

- Graphical InstallAnywhere utility

- **Scripts**

- Supporting scripts / tools
- Unix shell scripts and PL/SQL

# Java on Itanium 2: Works out of the box!

- Since ADS is almost entirely Java, no changes were required to:
  - ADS server
  - ADS cartridge
- No changes were necessary for any of the supporting scripts / tools
- InstallAnywhere 6.1 supports HP-UX on Itanium 2



# Changes done for Itanium 2

- Platform-specific functionality
  - ADS proxy – intercepting packets on a 64-bit OS
- Platform-specific optimization and tuning
  - Sorting during index creation was optimized for 64-bit platform
    - Do 64-bit comparisons (`long`) instead of 32-bit comparisons (`int`)
  - JVM garbage collection
    - `-XX:+UseParallelGC -XX:+UseAdaptiveSizePolicy`

# Testing plan for ADS on Itanium 2

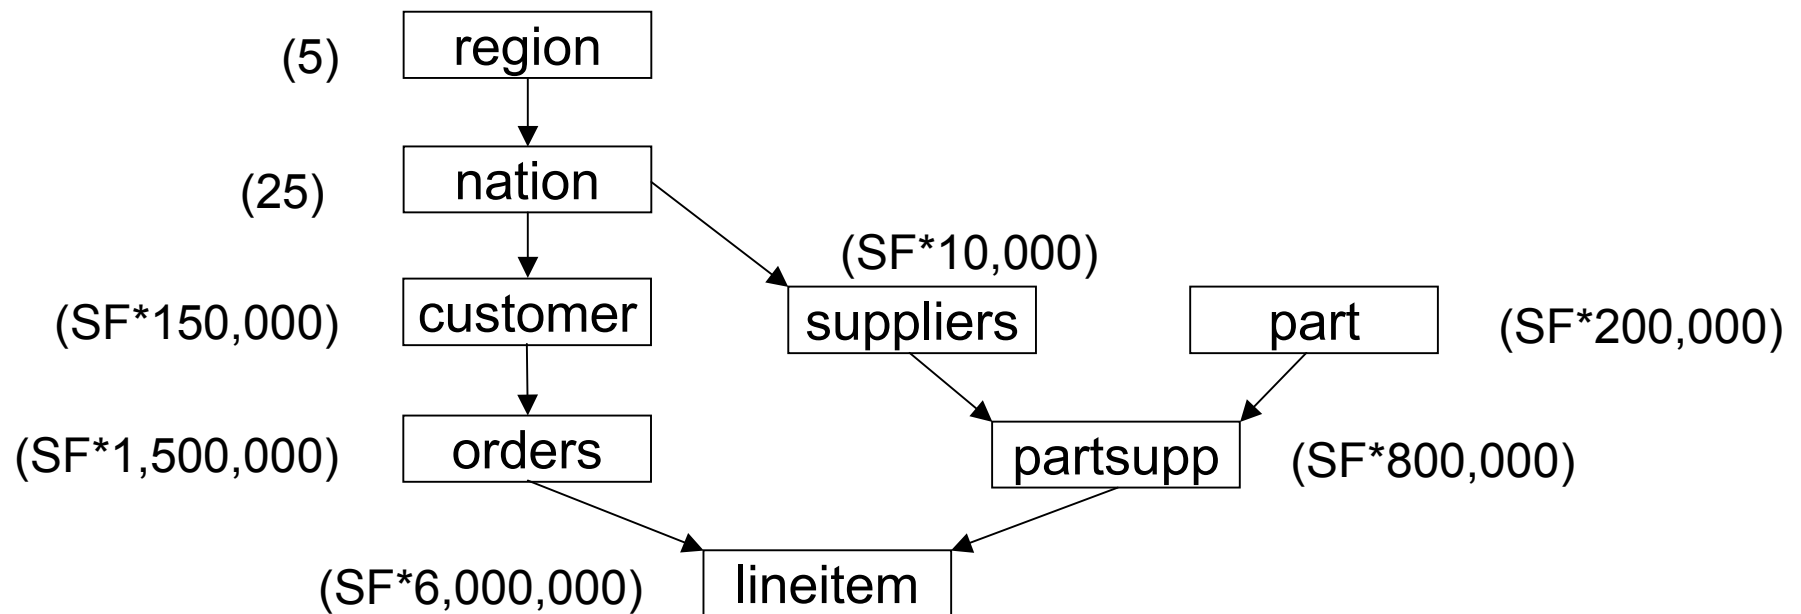
- Performance testing (ADS main “feature”)
  - Index build performance
  - Query performance
    - Single user
    - Multi-user
  - Increasing amounts of data
    - Scalability
- Functionality testing
  - ADS proxy platform-specific changes

# Testing ADS on Itanium 2: A “TPC-R-like” benchmark

- TPC-D, TPC-H, and TPC-R
  - A “family” of similar decision support benchmarks
- Tested:
  - 1GB, 10GB, 30GB, 100GB, and 300GB
  - Single-user and Multi-user
  - Cold cache and warm/hot cache
  - Time, I/O, and “consistent gets” statistics

# Schema and database size

- Several data sizes tested, up to 300 GB of raw data (= 2 TB database)



(#) – number of rows for each table

SF - Scale Factor, which is 1 for 1 GB, 10 for 10 GB, 30 for 30 GB, and so on.

**The largest table had 1.8 billion rows in the 300 GB run!**

# Porting and testing system specifications

- Earlier tests were done on HP PA-RISC platform
- HP provided an Itanium 2 machine at the Cupertino Global Solution Center
  - System specifications:
    - rx5670
    - 4 CPUs (Itanium 2 1.3GHz)
    - 16GB RAM
    - HP-UX 11.22
  - RDBMS and ADS were both installed on this machine
- Testing/Porting time frame: July to Sept '03

# Porting and testing system specifications (contd.)

- Disk and cache sizes:
  - ~2 TB of total usable storage (after mirroring, striping, standbys, OS, etc.)
    - 2 – VA7410 (45x36GB, 2 controllers, 2GB cache on each)
  - RDBMS cache
    - 200 MB for the 1 GB runs
    - 1 GB for the 10 GB runs
    - 4 GB for the 100 GB and 300 GB runs
  - 50 MB of ADS cache

# Queries

- ADS is used to accelerate some queries (typically the harder / slower ones) while leaving others untouched.
- There are 22 TPC-R queries. The same queries were tested (of these 22) on Itanium 2 as on PA-RISC.

# Run design

- Single-user runs
  - Cold cache
  - Warm cache
- Multi-user runs
  - Each user had a different ordering of queries.
    - User 1 ran: Q5, Q7, Q20, Q10, Q8, Q19m, Q8m
    - User 2 ran: Q19m, Q10, Q8m, Q8, Q5, Q7, Q20
    - Etc.



# Itanium 2 index build times

Raw Data Size	RDMBS Indexes (1 large and several small indexes, PKs, analyze)	ADS Indexes (6 join indexes, PKs)
1 GB	5m57s	25m44s (4m17s per)
10 GB	0h41m	4h34m (0h46m per)
100 GB	5h36m	35h25m (5h54m per)
300 GB	26h03m	78h00m (13h00m per)

- Goal: to build more complex (join) indexes in a comparable or lesser amount of time
- Result: Met or exceeded goal on the Itanium 2 platform

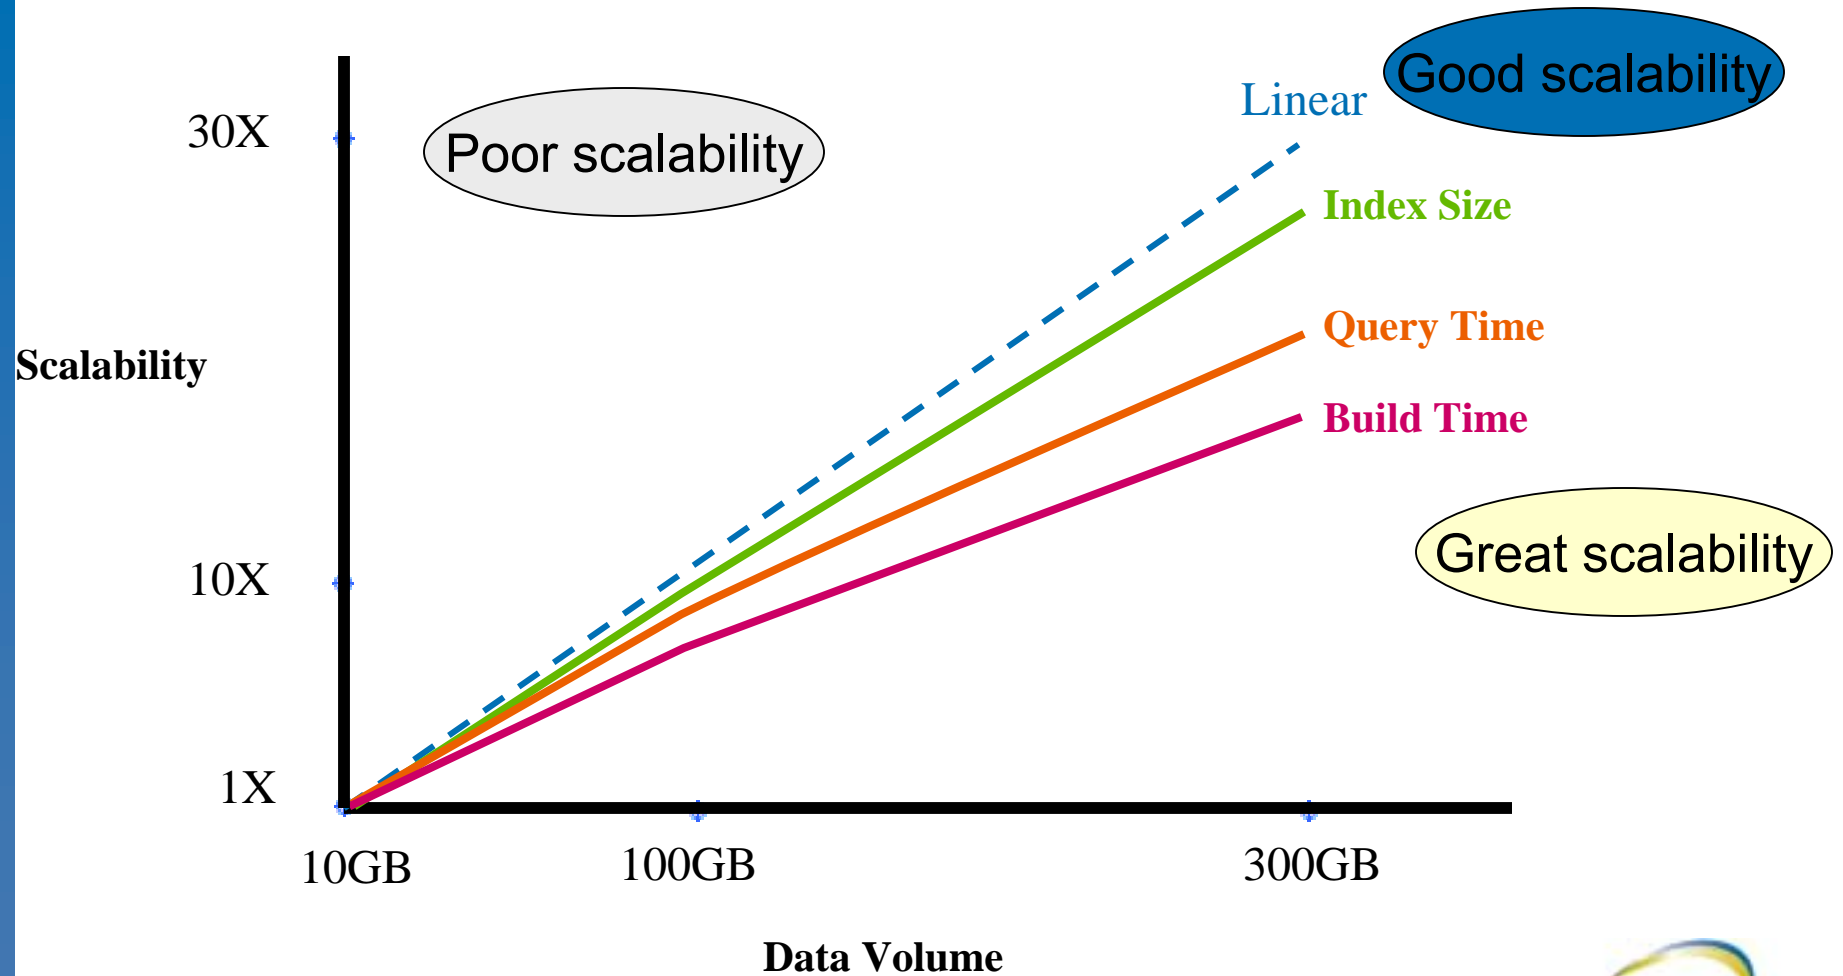
# Itanium 2 query results: 300 GB, single user, cold cache

Query	Itanium 2 Without ADS	Itanium 2 With ADS	ADS Benefit
5	9532 sec	2052 sec	<b>4.6x</b>
7	9488 sec	1436 sec	<b>6.6x</b>
8	5733 sec	790 sec	<b>7.3x</b>
8m	5711 sec	1429 sec	<b>4.0x</b>
10	7004 sec	1967 sec	<b>3.6x</b>
17	10823 sec	2058 sec	<b>5.3x</b>
19m	3703 sec	50 sec	<b>74.7x</b>
20	5353 sec	3440 sec	<b>1.6x</b>
<b>Total</b>	<b>57348 sec</b>	<b>13220 sec</b>	<b>4.3x</b>

# Itanium 2 query results: 300 GB, 10 concurrent users

Query	Itanium 2 Without ADS	Itanium 2 With ADS	ADS Benefit
5	21982 sec	4970 sec	<b>4.4x</b>
7	26399 sec	3318 sec	<b>8.0x</b>
8	14846 sec	1676 sec	<b>8.9x</b>
8m	16681 sec	2018 sec	<b>8.3x</b>
10	15269 sec	2059 sec	<b>7.4x</b>
19m	6457 sec	96 sec	<b>67.3x</b>
20	14016 sec	4757 sec	<b>2.9x</b>
<b>Total</b>	<b>115649 sec</b>	<b>18895 sec</b>	<b>6.1x</b>

# Itanium 2 ADS scalability analysis



# Summary

- Impressive query acceleration results on the Itanium 2 platform:
  - Different TPC-R queries
  - Different TPC-R data sizes (1GB, 10GB, 100GB, and 300GB)
  - Different numbers of concurrent users (1 and 10)
  - Different cache states (cold and warm)
- Excellent (sublinear) scalability from 10 GB to 100 GB to 300 GB.
- Initial runs in 3 days, entire project took 2-3 months
- **The Itanium 2 JVM, written by HP, was found to be rock solid and fast!**

“The performance of HP-UX 11i v2 on the Intel Itanium 2 -based architecture of HP Integrity servers, coupled with RightOrder’s ADS Performance Enhancer, provides customers with a complete solution to establish and maintain increased levels of cost/performance ratios in their data centers.”

**Melissa Laird**

General Manager, Solutions Enabling Division, Intel Corporation





# HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

