z/TPFDF V1.1

TPF Users Group Fall 2008
SDO z/TPFDF Data Access Service
Performance analysis

Name: Aleksandr Krymer
Venue: Database Subcommittee

AIM Enterprise Platform Software
IBM z/Transaction Processing Facility Enterprise Edition 1.1.0

Agenda

- **Availability**
- **General considerations**
- **Measurements**
- **Request types**
- **Read requests**
- **z/TPFDF Data collection comparison**
- **Analysis**
- **Conclusion**

Availability

- **SDO Access to z/TPFDF Databases is available!**
  - PK60030 and co-requisite PJ32720
  - Client side is available online:
    - http://www.ibm.com/software/htp/tpf/download/ztpfsdo.htm

General Considerations

- **Request stages:**
  - Client side (Java application)
  - WebSphere MQ
  - Server side (z/TPFDF central database routines)
- **Impact expected**
  - Difference in I/O operations on z/TPFDF – not really!
  - XML parsing on z/TPF – even cheaper than B2B scanner!
  - Network costs
  - Client-side costs (generally negligible)
- **Protective mechanisms are available**
  - Limit the number of ECBs, CPU time and I/Os
  - Limit response size – also places limit on z/TPFDF operations

Measurements

- **z/TPFDF and SDO Java drivers**
  - Running on the same machine, but communicate over the network
  - Java driver is running on Linux, Java 5
  - z/TPFDF driver was not written specifically for this test
- **Use same z/TPFDF databases on VPARS system**
- **10 runs each, roundtrip time measured in milliseconds**
  - Wall clock time
- **z/TPFDF data collection is run separately, thus execution times are not affected**
- **No other activity**

Request types

- **The data can be accessed in**
  - single-message and multiple-message requests
  - single top-level index file
  - single top-level index file and multiple detail subfiles
  - multiple top-level index files

Read requests

- **Multiple top-level index files across multiple interleaves, search by a single field**
  - Records have 3 versions, search field displacement is different for each version

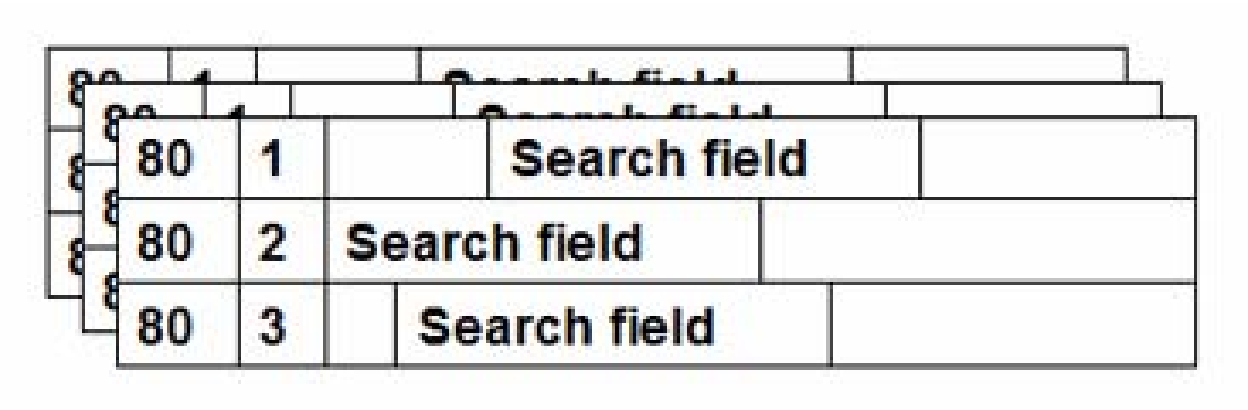

Read requests (cont.)

- **Multiple top-level index files across a single partition, wildcard search on a single field; retrieve multiple detail files 1 level deep**

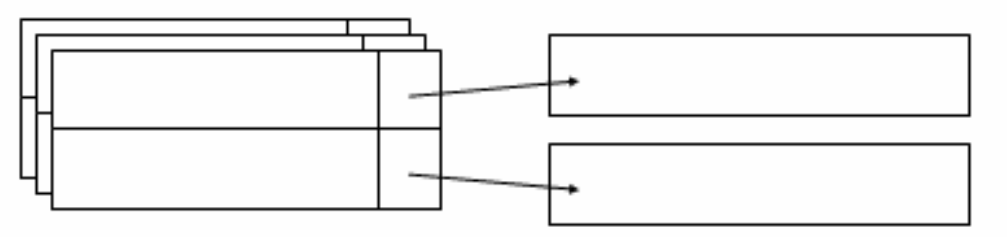

- **Single top-level index file, single index record, retrieve multiple detail files 2 levels deep**

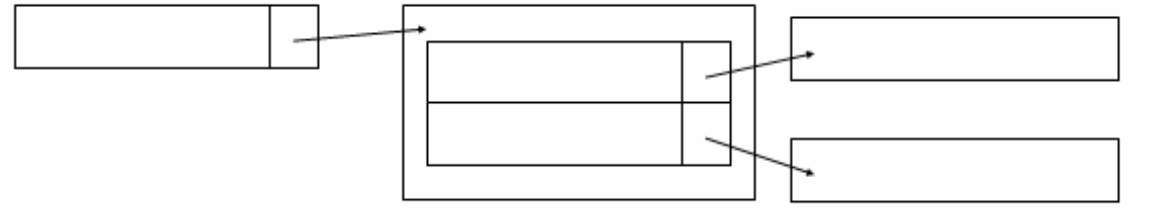

Read requests (cont.)

- **Average roundtrip times in milliseconds (wall clock):**

#	z/TPFDF	SDO
1	2450	2395
2	388	530
3	6	131

- **Network latency is almost the same in the 2nd and 3rd test cases; negligible in the 1st one**

z/TPFDF data collection comparison

- **The number of PFINDs is identical in all test cases**
- **In all test cases the number of DBREDs is the same or smaller for SDO-based driver**
- **In all test cases the number of DBOPNs/DBCLSs is the same or smaller for SDO-based driver**
- **The number of DBKEYs can vary in both directions**

Analysis

- **Longer and larger queries hide network latency effect**
  - Delays are similar on medium and short queries
- **Number of z/TPFDF API calls is very similar**
  - SDO driver causes additional filesystem I/Os
- **No significant CPU utilization increase**
  - WebSphere MQ, C++ code, XML parsing (even without B2B parser)

Conclusion

- **Additional cost for z/TPF system – very low to none**
  - Extra I/O – close to none
  - Extra CPU – a little
    - C++ code is discounted
- **Read queries do not have a negative impact on z/TPFDF I/O performance**
  - SDO applications are often as efficient as handcoded z/TPFDF applications
  - See Glenn Katzen's presentation "Application Development using SDO Access to z/TPFDF – Advanced Features" at Application Development Subcommittee for efficient SDO-based z/TPFDF application design
- **The largest impact on throughput is the network**
  - Can be mitigated by placing an application server and z/TPF on the same physical machine or close within the network

Trademarks

- IBM and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries, or both.
- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Other company, product, or service names may be trademarks or service marks of others.

- Notes
- Performance data contained in this document was determined in a controlled environment; therefore the results which may be obtained in other operating environments may vary significantly. No commitment as to your ability to obtain comparable results is any way intended or made by this release of information.
- Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
- This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

- Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

- Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

- This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.