Performance Topics

Robert Blackburn Ph.D.

IBM z/TPF

April 12, 2016

VFA vs DASD CPU consumption

- Historically considerable discussion of CPU benefits
  - records in VFA or DASD
- General agreement
    - VFA lowers message response time
- Measurement result:
  - VFA has significant CPU savings
- Customer measurements differing
    - Strongly suspect confounders
      - E.g. application change dominating effect signal

VFA / DASD Measurement

	Msg/sec	util	ITR		% gain
VFA	5399	.439	12298		26%
DASD	5365	.552	9719		

Test design – AIR1

- Need total logical accesses of AIR1 messages ~ equal
  - got rid of DUPs to lessen variability
- Actual AIR1 logical IO numbers very close
  - All VFA = 15.45 VFA / msg
  - All DASD = 15.18 DASD / msg
- Slight MIPS decrease for VFA message
    - CPUMF used
- Pathlength reduction dominated
    - Roughly 15 x (1700-400) / 85000 = 23%

Summary

- Often new potential VFA candidates
  - unknown hit ratio characteristics
- Err on the side of including records as VFA candidates
  - Non LC customers
- Include a plausible set of records
  - measure the result
- Some customers very accurately modeled by ‘square root rule’
- Others almost unaffected by VFA size changes
  - In the 1G -3G range

DF cache - One number per file

- TPF computes
  - Median usage depth of subfiles over time by DF file
    - suspect little variation over peak time
- Customer sets this value - by DF file
- Predictive information in previous use of fixed subfile is near 0
    - Applications change and dominant effect
    - Say ECB(722) did 42 DASD access
    - Best prediction for future ECB(722) is median
      - not f(42 DASD)

Data Relationships

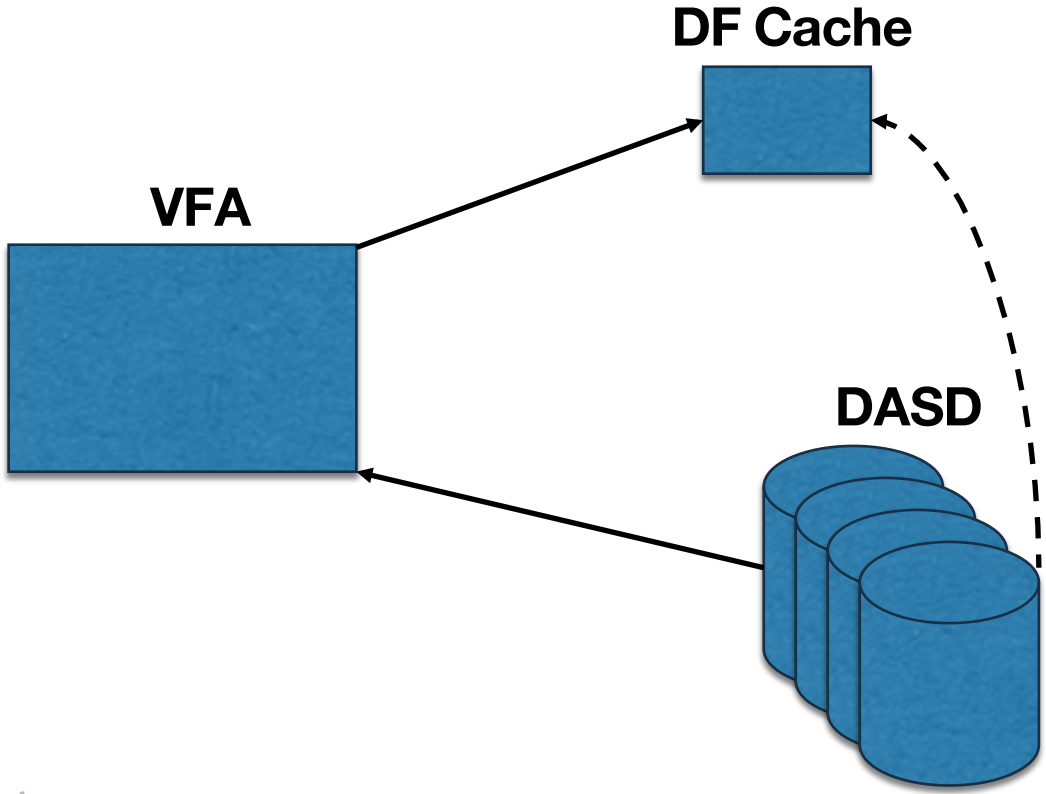

DF Cache

- Let b = say, 5 = best estimate of usage number for DF file
  - Applies to all subfiles in this file
    - e.g. PNR and specific customer
- Let n = actual access number for each subfile
  - Random variable (N=n) across the file
- Statistical costs of two kinds
  - if n > 5 then (n – 5) more VFA access must be done
  - If n < 5 then wasted (5 - n) moved records in ECB
  - Roughly costs are equally painful (symmetry)

Hit the one in the middle

- Assume throw a 11 sided dice, 0 to 10
- You get to choose one number
  - absolute deviation is cost function
- Realization of 6 dice rolls might be 7,4,0,10,5,3
  - With median = 5 as predictor
    - Deviations 2,1,5,5,0,2
    - sum = 15
  - With 10 as predictor
    - Deviations 3,6,10,0,5,7
    - sum= 31

Best Predictors

- Absolute distance cost => median
- Squared distance cost => mean
    - Worthwhile thinking about which is appropriate
- Consider standard exponential exp(-x)
  - Mean =1
  - Median = .693
- Conclusion:
  - DF cache -- absolute distance
  - Means and medians can differ significantly

Sizing DF Cache

- Number different files going into DF Cache
  - E.g. PNR , Inventory etc
  - Records accessed per day
- Small – try 100M
- Larger – try 500M
- TPF has castout and other measures of efficiency

Conjecture - on defect arrival rates

- Suspect P(defect)
  - significantly increases when new code is loaded
  - welcome customer data
- Tends to be off peak
  - some component of failure correlated with MIPS used
    - Suspect small
- Off peak has extra MIPS available and higher rate system failure
  - Therefore turn on all suggested traces

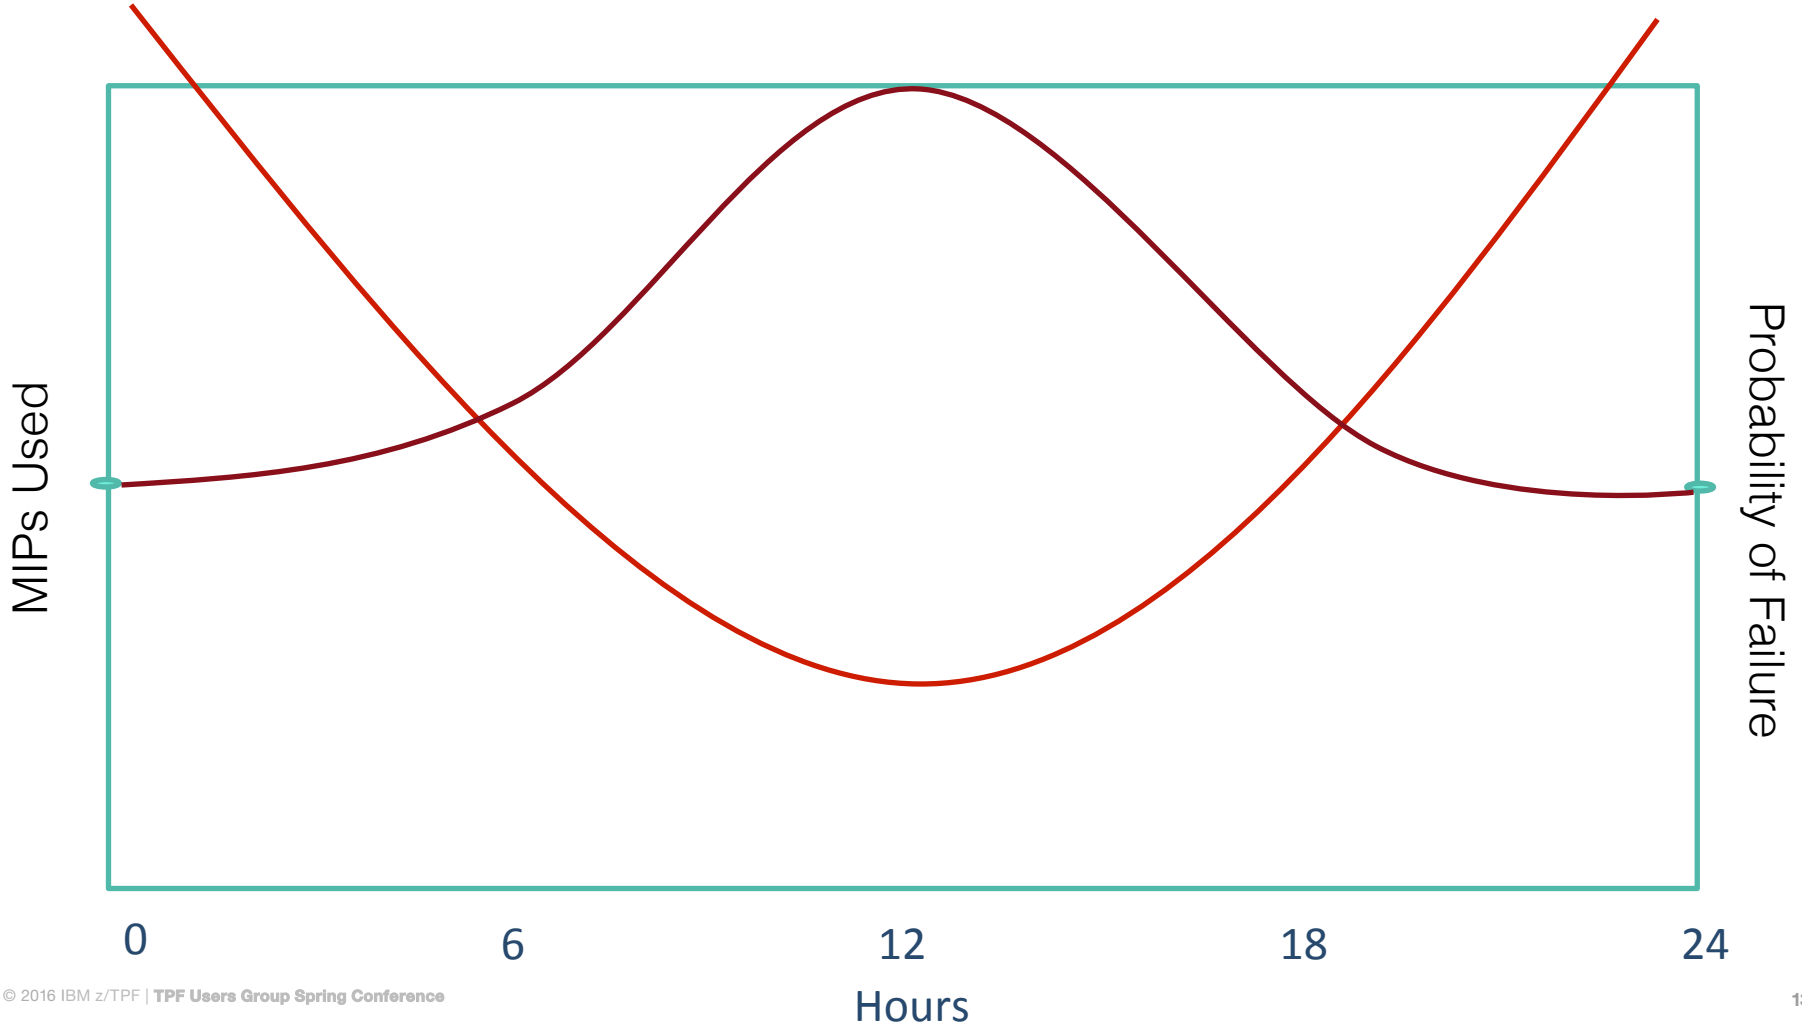

Traces of interest – off peak time

- C function trace -extended version
- ECB macro trace with regs
- Enter / Back
- ECB Heap trace
- DF enter trace
- Socket trace
- Stack validation
- NOT recommended in production
  - Block check mode
  - Heap check mode

Lab recommendation

- Run with traces on even at peak times
  - High CPU utilization defects can be severe
    - Difficult to diagnose and solve without trace data

- However if aberrant workload causes CPU overload
  - Turn off some traces
  - Try to keep C function and ECB trace active if possible

Trademarks

- IBM, the IBM logo, ibm.com and Rational are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Notes

- 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.