z/TPF EE V1.1
z/TPFDF V1.1
TPF Toolkit for WebSphere® Studio V3
TPF Operations Server V1.2

IBM Software Group

TPF Users Group Spring 2006

Storage Allocation with z/TPF

Robert Blackburn

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

Introduction

- General storage recommendations for z/TPF systems
  - test
  - production
- Preallocation methods to optimize performance
  - models
  - distributions
  - data collection reports

Don't be afraid of this

- TPF customers have solved similar problems for years
  - Utilization to run
    - CPU
    - DASD
    - CH/CU
    - SAPs
  - VFA candidates
  - File resident programs

Storage for test systems

- In the lab test systems grew from 250M to 1G
- For test systems use factor of 4
- z/TPF does't touch large storage until used
  - VFA blocks
  - CRPA
  - 1M frames
- Experiments
  - z/TPF touched roughly 1/4 of its storage
    - depends on how memory was allocated
  - TPF4.1 touched all its storage
  - z/TPF vs TPF 4.1 working sets about equal (4/4=1)
- 3 memory sizes for test systems
  - Often test systems have mean ECBs < 10

Storage for production systems

- Assume production system of 2G
  - .5G for VFA
  - 1.5G for not-VFA
- Not-VFA multiply by 3
  - 1.5G x 3 = 4.5G
- VFA 20% increase
  - .5G x 1.2 = .6G
- Sum is 5.1G we round to 6G
- Production systems use memory factor of 3

Initial z/TPF configuration

- When customers convert to z/TPF there will be extra storage
  - TPF4.1 < 2G and late model minimum CEC > 8G
- Be generous in initial z/TPF memory allocations
- Optimize using z/TPF data collection
  - TPF4.1 does not have some required reports
- Key point: have a huge number of 1M frames
  - Much more important than putting memory to VFA
  - System will be protected even with suboptimal preallocation

CDFs with different means

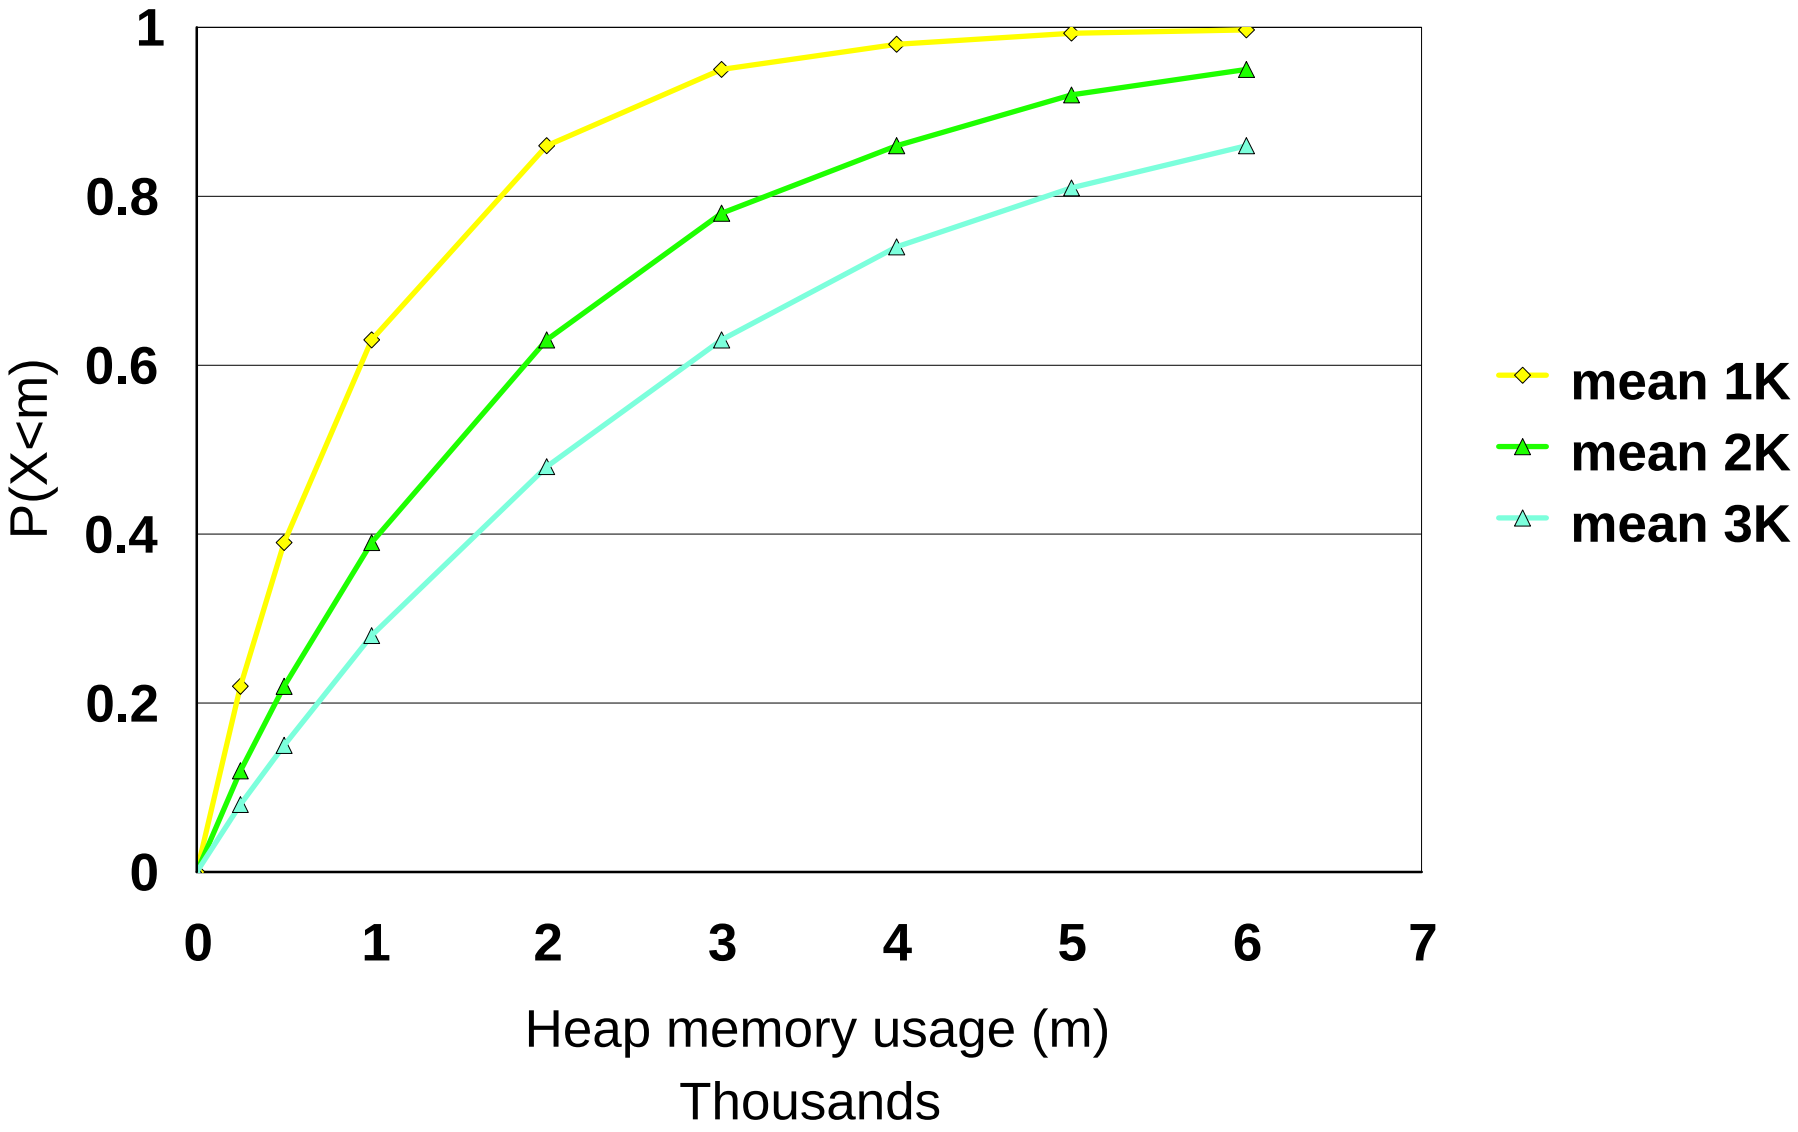

Determine 31-bit heap preallocation

- Each ECB gets fixed amount of heap - call it P
- If ECB exceeds P then gets 1M frame
- Memory vs memory and memory vs performance
- Assume mean heap of 20K
  - If P=0 then a.s. each ECB gets 1M
  - If P=1M then wasted memory--few ECBs need > 100K
- Use TPF ECB 31-BIT HEAP AREA USAGE REPORT
  - Pmf of max heap use by ECBs over their life
    - from this create cdf F(x)
  - choose x such that .93 < F(x) < .98

31-bit heap calculations

- define r = mean active ECB / defined ECBs
- let r = 300/1000
- cost is 1000x + (1-F(x)) 300 (1M)
  - g(x) = x + (1-F(x))r M
- for exponential 1-F(x) = exp(- lx)
- minimizing g(x) we get x=2.7 times mean
- F(2.7 mean) = .933

Determination of A1 to A4 sizes

- Use the ECB HEAP REQUEST SIZE REPORT
  - Set A4 = 3 times mean heap
  - Pmf shows 0 to 2A4 in increments 2A4/64
- Set
  - F(A1)=.2
  - F(A2)=.4
  - F(A3)=.6
  - F(A4)=.8

Determination of A1 to A4 counts

- Set buffer count = 0 so all ECBs miss
- Compute m= mean heap per ECB
- Counts = (m + 3 sqrt(m))/4
  - counts will vary from the mean
  - using Poisson estimate
- Key point - not concerned about single ECB greatly exceeding counts
  - not a ruin problem like insurance/dam construction
  - inherently a ratio problem
  - system performance problem but not allocation problem

Application Stack

- If ECB exceeds preallocation gets 1M frame
  - Memory vs memory has been solved
  - x such that .92 < F(x) < .96
- Virtual limit
  - too small then too high % of ECBs perish
  - too large then steal virtual space < 2G
  - choose y such that F(y) > .995
  - Use PREALLOCATED APPLICATION STACK STORAGE USAGE REPORT

Thread Stack

- No preallocated amount - backed by real as needed
- Virtual storage issues as Application Stack
- Virtual = Tstack x Maxthread
  - If > Tstack then death
  - If > maxthread get negative returns - so can control
- Use THREAD STACK STORAGE USAGE REPORT
- Use THREAD NUMBER PER PROCESS REPORT
- for both y such that F(y) > .995
- weak upper bounds - control tail of process to get tighter

ECB Private Area

- If exceed prealloc by even 4K then build page/segment tables up to virtual limit
  - limit in 4.1 is 1M so good start point
- More accuracy/change use ECB PRIVATE AREA STORAGE USAGE REPORT
  - set prealloc at 6 times mean ECB private usage
  - if have cdf then x such that F(x)=.99
  - examine long tail--bad ECBs or error conditions

2G Virtual limit

- 31-bit system heap
- ECB 31-bit heap
- Thread stack
- Application stack
- ECB private area(<32M)
- 31-bit CRPA

Trademarks

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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.