z/TPF V1.1

TPF Users Group Fall 2008

Performance Topics

Robert Blackburn Ph.D.

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

Topics

- **z/TPF initial 1M allocation**
- **z/TPF low utilization effect**
- **z/TPF performance APARs**
- **PR/SM usage virtual**

At previous TPFUG it was stated

- **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**

- **Common question: What is a ‘huge number’ of 1M frames?**
- **With several LPARs in a CEC this becomes more important**

z/TPF Initial Allocation of 1M Frames

- **Use TPF4.1 'ECB Heap Area Usage Summary Report'**
- **Set ECB preallocated (P) at the 90% quantile**
  - thus on average 10% of the ECBs will need a 1M frame
- **Let R = (Defined ECB) x .1**
  - Worst case essentially all ECBs become active
- **Let C= number 1M frames needed for core resident programs**
- **Initial Allocation = 2(R+C)**
- **Can quickly optimize as discussed in previous talk and reduce 1Megs as needed**

z/TPF Low Utilization Effect

- **Low utilization effect**
  - ITR(low) < ITR(high)
- **Measured on z9 2094-701**
  - z/TPF has .5% (<1%) utilization at 0 throughput
  - TPF4.1 has .15% utilization at 0 throughput
  - Difference is 1.4 MIPS at idle
- **On large z990, z9 and z10 with high utilizations this effect is trivial**
- **On older smaller CECs or subcapacity large CECs the effect can be significant and we need adjustment**
- **Lab is working to reduce this effect to near TPF4.1 levels**

Typical Graph of Low Utilization Effect

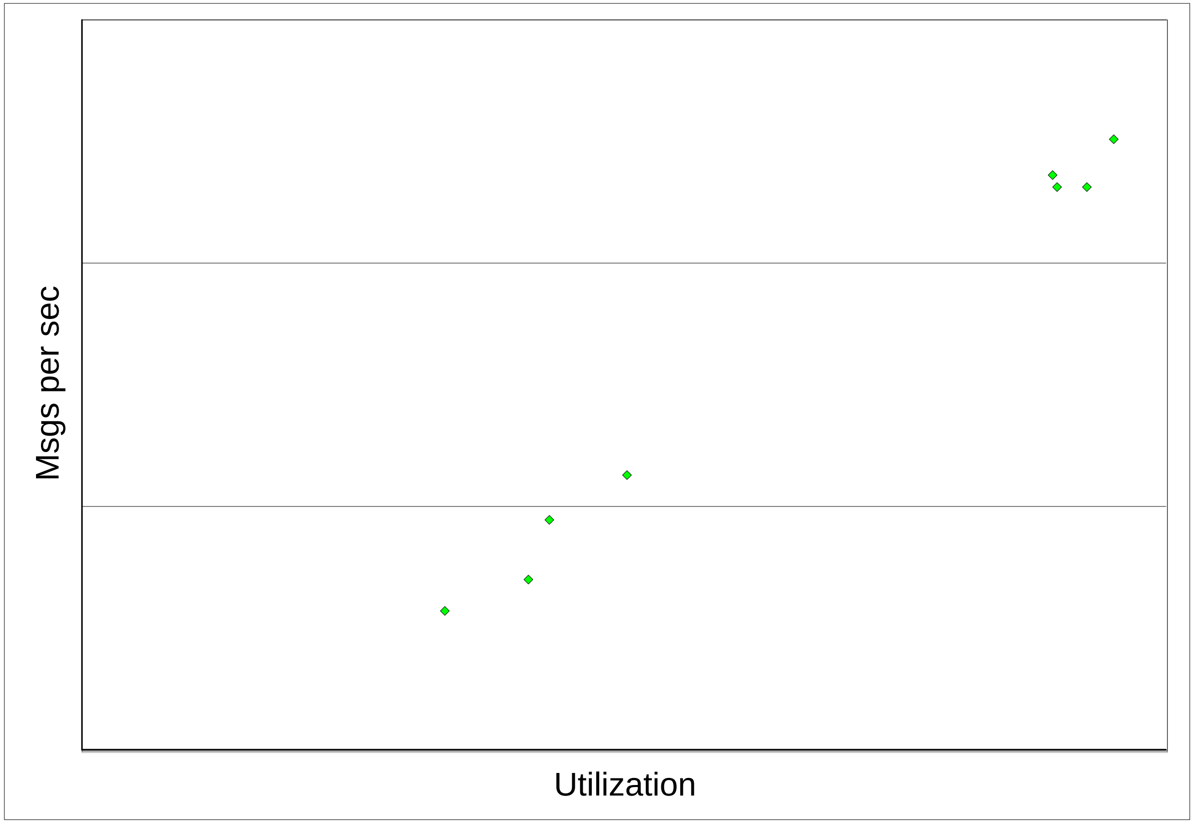

Adjust for low utilization

- **Assume the throughput graph is of the form m = a(u – b) where b>0**
  - where u = utilization, m = message rate
  - a and b to be determined empirically
    - e.g. by a graph similar to one on previous page
- **For our example we set a=100, b=.1**
- **At u=.4 data collection will report m=100(.4-.1) = 300**
- **Raw ITR = m/u = 300/.4 = 750**
- **But real throughput at u=1 will be**
  - 100(1-.1)=900
- **Adjusted ITR = m(1-b) / (u-b) gets the true ITR**
  - 300(1-.1)/(.4-.1)=900
- **Basically just pull (b) out of the utilization but then only scale up to (1-b) since the effect still exists at very high utilization**

z/TPF Performance APARs and in-flight APARs

- **Increased use of common bit**
  - Two new CRPAs
    - No COW provision
    - Large memory saving
      - reduced page/segment table construction
- **Sweeper for low usage programs(64 bit)**
  - Memory saving
- **VFA 381 and 1K rotation inside 4K**
  - Cache benefit
- **Prefix page separation I and D cache lines**
  - Cache + coherency manager gain
- **Reduce low utilization effect**

Routing Weights – previous TPFUG talk – priority queue

- **CEC has 3 real CPs**
  - TPF LPAR1 has 3 shared CPs
  - TPF LPAR2 has 1 shared CP
- **Deliberately LPAR1 has weights set higher than it ever uses**
  - RESULT: a priority queue with LPAR1 high priority
- **LPAR1 p= .6**
- **LPAR2 p=.8**
  - Input list queues of 1200
- **CEC utilization**
  - (3x.6 + .8)/3 = .86
  - Why is LPAR2 acting as if its utilization=1?
- **Assume independence of 3 virtual CPs in LPAR1**
  - P(all 3 CPs busy) = $.6^3$ = .22
  - P(at least one CP available) = 1-.22 = .78
  - This is very close to actual LPAR2 util of .8
  - Key point is that LPAR1 doesn't 'feel' LPAR2 - preemptive
- **Thus LPAR2 can not reach p=1 with only 1 defined virtual CP**

Routing weights – non priority queue

- **Now consider when no LPAR is favored**
- **CEC has 3 real CPs**
  - TPF1 has 2 shared CPs – weight = 100
  - TPF2 has 2 shared CPs – weight = 100
  - TPF3 has 1 shared CP – weight = 100
- **Result is TPF3 p=.92 --- at input list size of 1000**
  - TPF1 p=.23, TPF2 p=.45
  - Total CEC is (2(.23+.45)+.92)/3 = .76
- **Look at 2(.23+.45) /3 = .45**
  - Not TPF3 busy component
- **P(3 CPs busy) = .45**3 = .09**
- **Since TPF3 is further ahead of its weight than the other LPARs it will not preempt**
- **1-.09 = .91 close to .92 actual utilization of TPF3**

Shared PR/SM utilization guidelines

- **Consider p=.9 as an upper bound for the utilization to run a TPF shared LPAR**
  - e.g. for dyadic at p=1 move to a triadic so p drops to .66 and then just pay increased V/R costs
- **This is not limited to just a single CP LPAR**
- **Essentially PR/SM doesn't see a very long input list – just that an LPAR is ready or not to run**
- **Every time an LPAR gives up control and another grabs 10 mill slice is lost potential if trying to run at p=1**
- **Altered intuition: at p=1 an open queueing sytem has mean queue length = infinity**
  - this bothers some people but is mathematically provable

Trademarks

- IBM, z9, and z10® are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.
- Other company, product, or service names may be trademarks or service marks of others.

- 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.