

z/VM 7.2 1Q21 Performance Highlights

Brian K. Wade, Ph.D.

Email: bkw@us.ibm.com

Version 2021-06-10.1

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

Db2*	FlashCopy*	IBM (logo)*	OMEGAMON*	z13*	z/Architecture*	zSeries*
DirMaint	FlashSystem	IBM Z*	PR/SM	z13s	zEnterprise*	z/VM*
DS8000*	GDPS*	LinuxONE*	RACF*	z14	z/OS*	z Systems*
ECKD	ibm.com	LinuxONE Emperor	System z10*	z15	zSecure	
FICON*	IBM eServer	LinuxONE Rockhopper	XIV*	z10 BC		
				z10EC		

* Registered trademarks of IBM Corporation

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries. IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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.

Linux is a registered trademark of Linus Torvalds 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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

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

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other product and service names might be trademarks of IBM or other companies.

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.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

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.

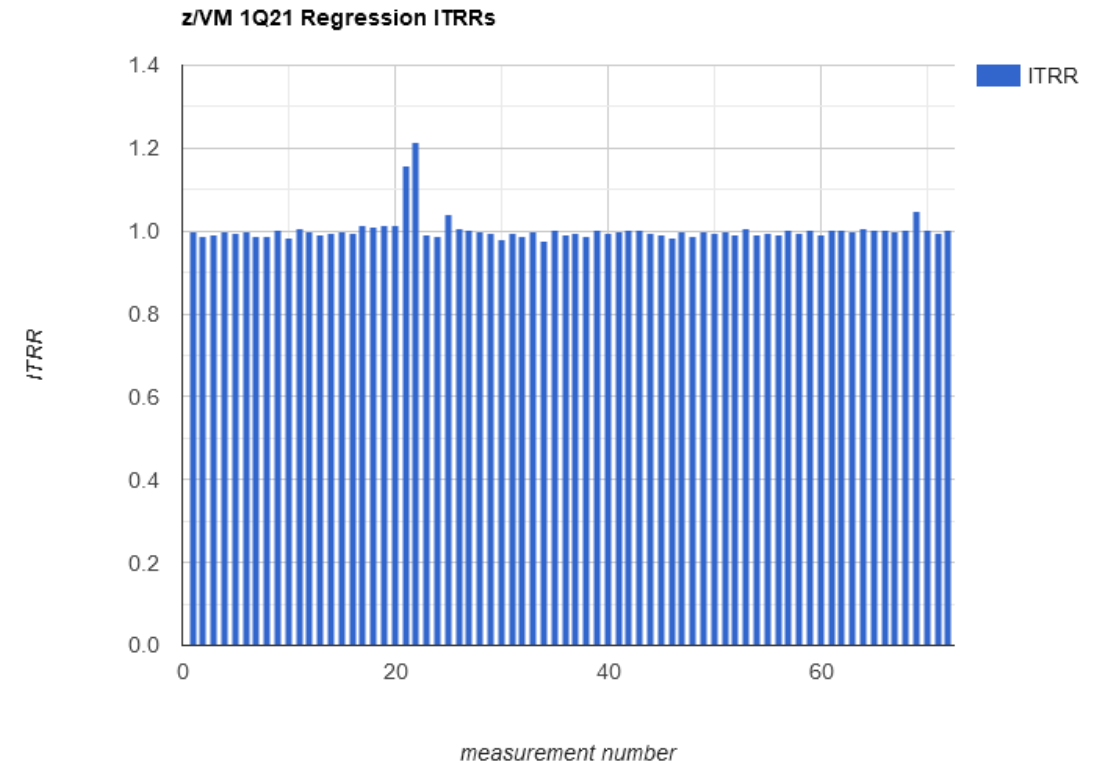
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Agenda and Your Takeaways

Agenda item	APAR number	Your takeaway
1Q21 regression results		all good
4 TB exploitation	VM66173	it scales
z/HPF I/O for dumps	VM66431	dump elapsed time -45%
Guest instruction simulation improvements	VM66467	transaction response time -25% steal percent -24%
TCP/IP exploitation of OCSP	z/VM 7.2 base	CPU time per connection -20%
Coping with run variation		We're coping!

Regression Report Card

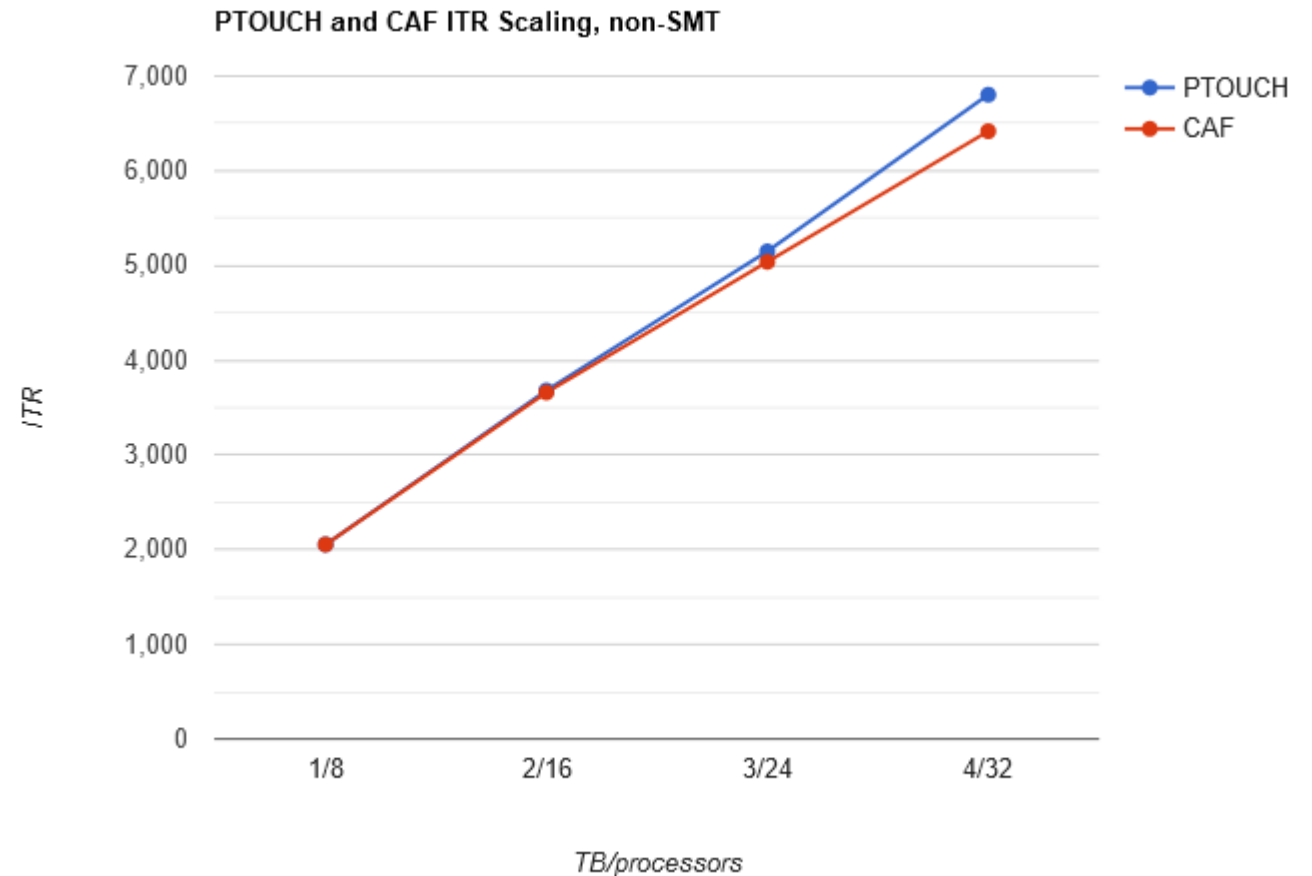
- We run about 70 regression experiments for every z/VM release or bundle of NF APARs
 - These are workloads that **do not intentionally exploit new function**
 - They are meant to find performance problems a client might find in simply going from z/VM x to z/VM $x+1$
- For each workload we collect:
 - "External" transaction rate (ETR): application work accomplished per elapsed second
 - "Internal" transaction rate (ITR): $(ETR / (\text{system busy})) =$ a measure of CPU efficiency
- For each workload we compute regression ratios: $ETR(\text{new})/ETR(\text{old})$ and $ITR(\text{new})/ITR(\text{old})$
- For z/VM 1Q21 compared back to z/VM 4Q20:
 - $ETRR (\mu, \sigma) = (1.004, 0.029)$ smallest ratio was 0.969
 - $ITRR (\mu, \sigma) = (1.004, 0.033)$ smallest ratio was 0.978
- Analysts: Bill Guzior, Dave Spencer, Xenia Tkatschow, Brian Wade, Dave Wierbowski



The two outliers are workloads that very heavily stress guest page instantiation and release. The 4 TB line item improved that code path. The ITR difference was caused by reduced path length.

4 TB Exploitation: NF APAR VM66173

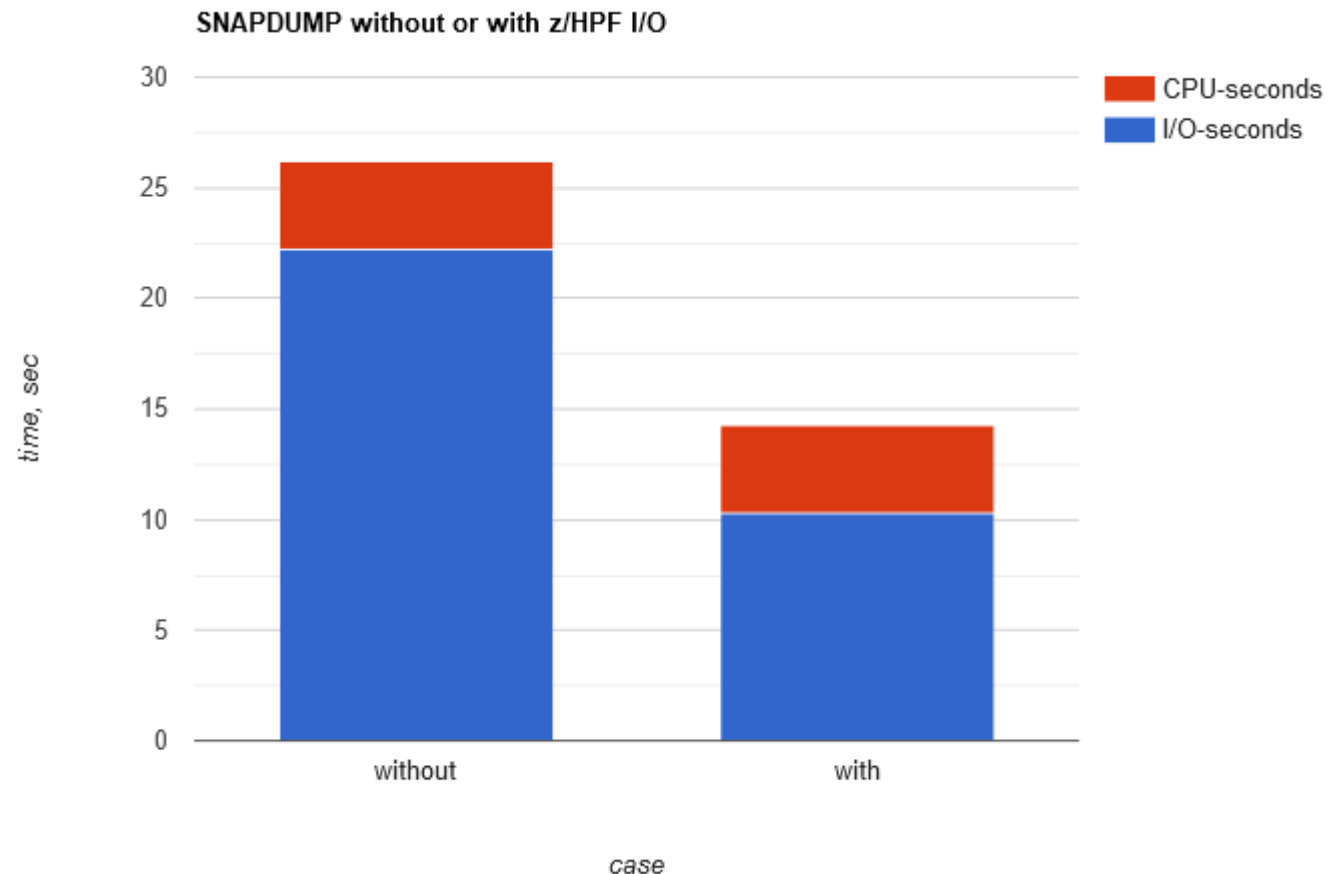
- Exploiting larger hardware is far more than just not abending
- We ran a variety of workloads:
 - Low, medium, high nest pressure
 - Some Linux, some CMS
 - Some memory-rich, some memory-constrained
- What we typically do is run an *arc* of scaling measurements
 - Stepwise ramp-up of workload, processors, and memory
- Net: if the Z hardware can be expected to scale the client's workload, z/VM also can be expected to scale it
- <http://www.vm.ibm.com/perf/reports/zvm/html/1q14tb.html>
- Analysts: Dave Spencer, Xenia Tkatschow, Brian Wade



NB: **CAF** is the scaling curve predicted by the SYSIB 1.2.2 Capability Adjustment Factor coefficients

z/HPF for Dumps: NF APAR VM66431

- Changed SNAPDUMP, etc. to use z/HPF (transport-mode) I/O instead of ordinary command-mode I/O
- **Result:**
 - I/O time in the dump was cut by about half (I/O is twice as good)
 - In our measurement, dump elapsed time was cut by 45%
- <http://www.vm.ibm.com/perf/reports/zvm/html/1q1dmp.html>
- Analysts: Bill Guzior, Brian Wade



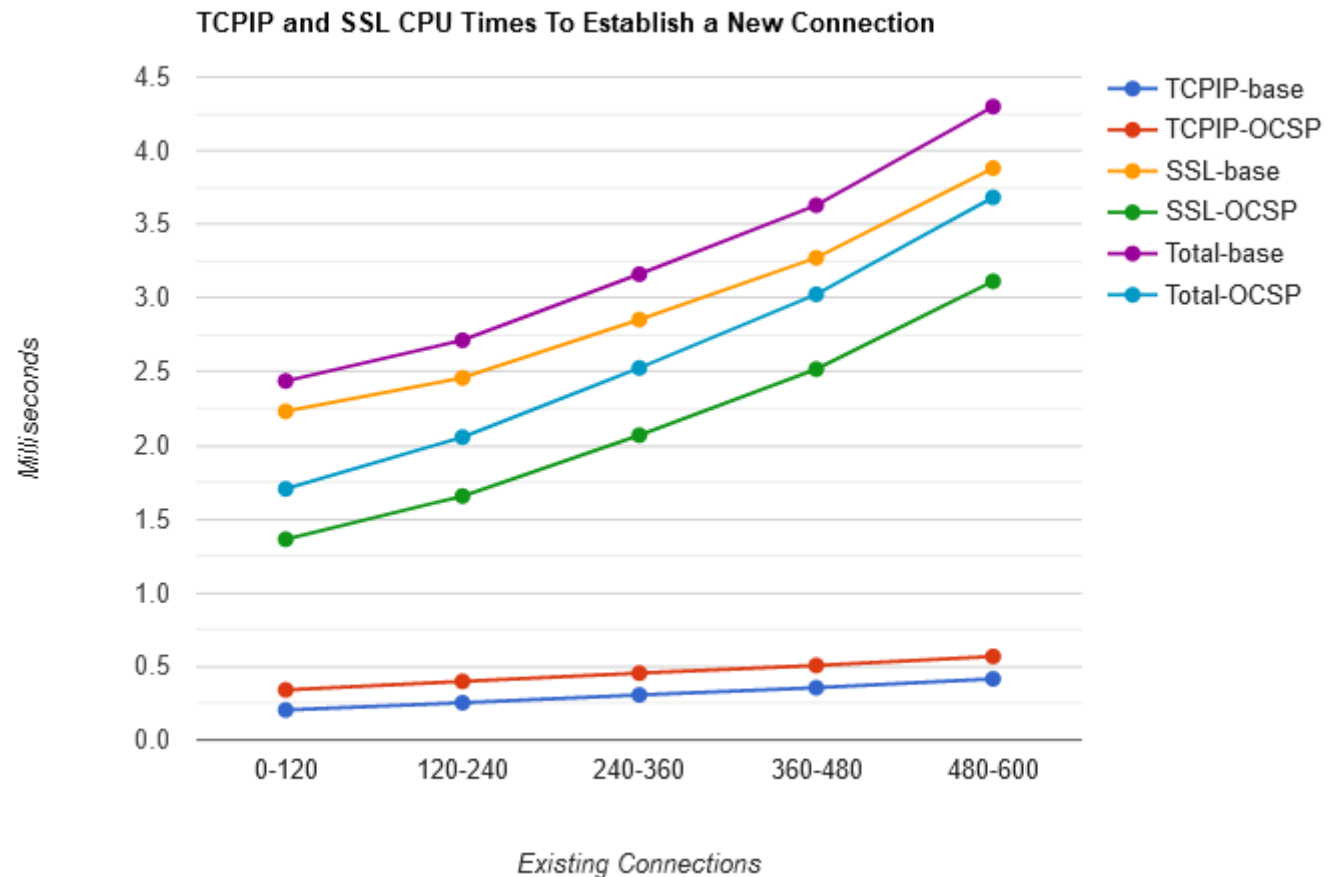
Guest Instruction Simulation Improvements: NF APAR VM66467

- Improved z/VM's simulation of an instruction guests use when interacting with a PCIe adapter
- z/VM CP now talks to the real hardware less often
- Results: in our workload,
 - 25% decrease in transaction response time
 - 24% decrease in guest steal time
- <http://www.vm.ibm.com/perf/reports/zvm/html/1q1rpc.html>
- Analysts: Marc Beyerle, Thorsten Winkler, Brian Wade



TLS/SSL Server OCSP Exploitation, z/VM 7.2 Base

- Supports Online Certificate Status Protocol (OCSP) and exploits Certificate Distribution Points (CDPs)
- Offloads the validation of client certificates from the z/VM SSL server to an external validation server
- Reduces CPU time needed to establish a new connection with a client
- Overall a 20% decrease in the CPU time needed to establish a new connection
- <https://www.vm.ibm.com/perf/reports/zvm/html/1q1ssl.html>
- Analyst: Xenia Tkatschow



1Q21 Honorable Mentions

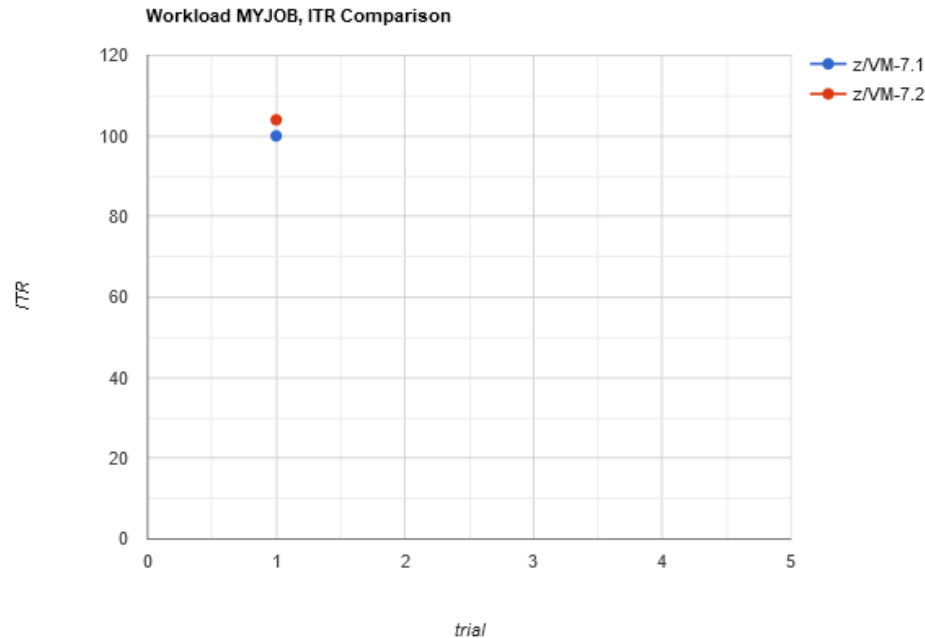
- Fast z/VM Dump Distiller
 - VM66430
 - Decreases time needed to run DUMPLOAD
 - The improvement will vary according to the character of the dump
- z/XC Architecture Support
 - VM66201 (CP), plus VM66425 (CMS), plus VM66489 (Perfkit)
 - After a drought of a few years, we can once again run Perfkit in an XC-mode virtual machine
 - This means we can once again run Perfkit in a guest that is also using SFS Dircontrol in Data Space
 - This made my day!

Coping with Run Variation

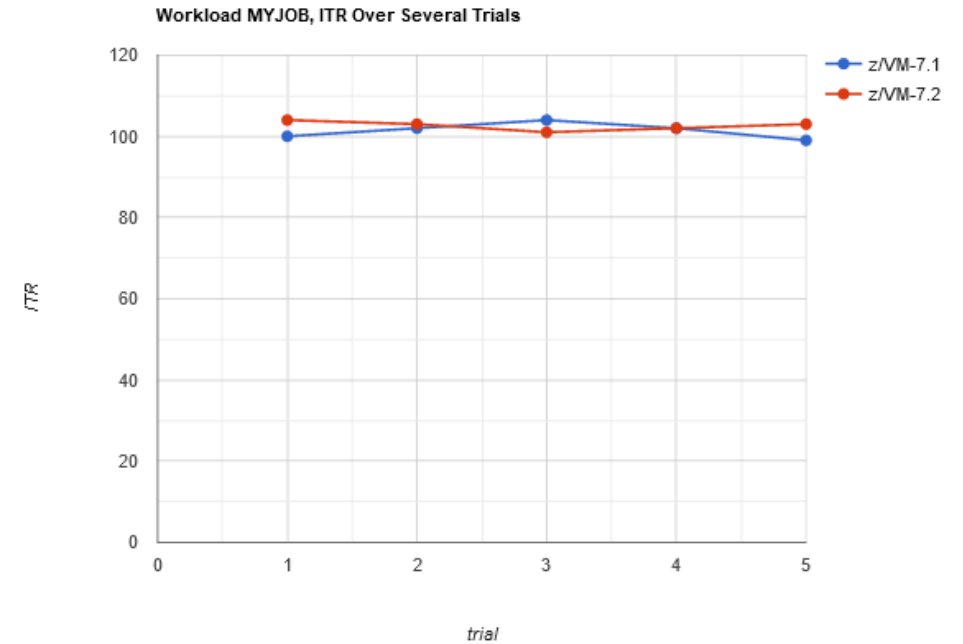
- We want everything to repeat precisely, but it doesn't
 - This is the natural way of the world
 - Especially hits us in measurements using networking or I/O
- To be successful as analysts we need to be on the lookout for variation and cope with it when it happens
- There are mathematical techniques we can use to discern between run variation and true change

We Can Be Misled

- One trial of each experiment showed us a 4% increase
- More trials showed us this is an experiment that naturally varies
- Student's *t* showed us the difference in the means is **not significant**
- We can conclude **nothing**



z/VM 7.1 100
z/VM 7.2 104
change +4%




z/VM 7.1 mean 101.4
z/VM 7.2 mean 102.6
change +1.2%

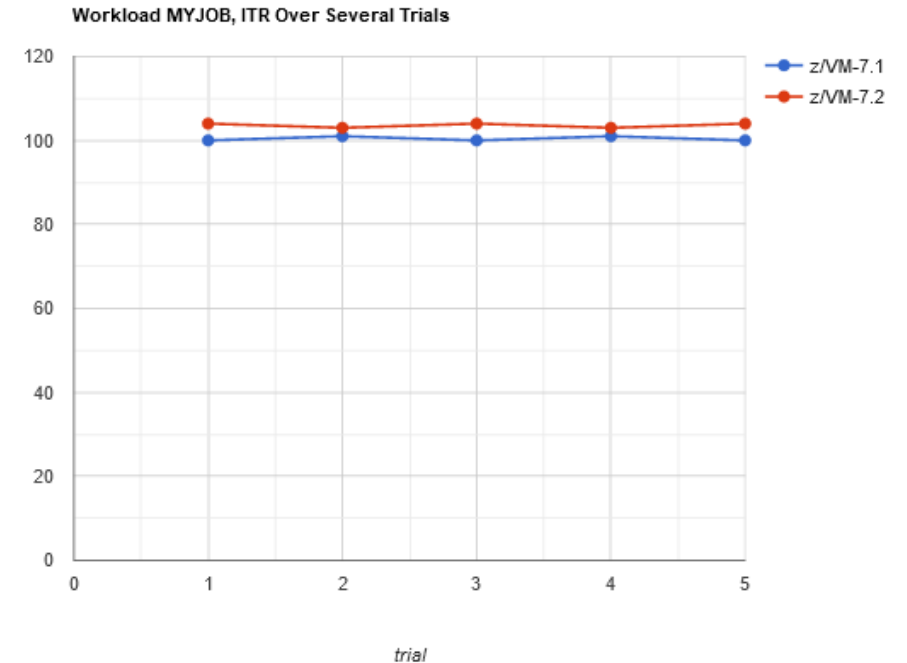
but $cl=0$
the difference is not meaningful

Detecting Meaningful Change

- One trial of each experiment showed us a 4% increase
- More trials showed us this is an experiment that naturally varies
- Student's t showed us the difference in the means is significant
- We can conclude there was a change



z/VM 7.1 100
z/VM 7.2 104
change +4%



z/VM 7.1 mean 100.4
z/VM 7.2 mean 103.6
change +3.2%

and $ci=99\%$
the difference is meaningful

Repeating Is Not Necessarily Painful

- Maybe your run takes only ten minutes and it's automated
 - So you just repeat it a few times
- Maybe your run is like this:
 - It runs for ten minutes, and
 - It emits measurement data at 1-minute intervals, and
 - It also collects a sample of ETR at 1-minute intervals
 - In this case your run is really ten 1-minute experiments
- If your ETR is something tracked in monitor sample records, you're golden
 - For example, throughput of a vswitch uplink port

Our Networking and I/O Analyst: Dave W

- His runs are very heavily automated
- Each run takes about five minutes
- So, for each workload, he
 - Collects N trials of, say, 4Q20
 - Collects N trials of, say, 1Q21
 - Uses student's t to check whether the difference of the means is statistically significant
- Is the difference of the means statistically significant at $\alpha \geq 95$?
 - Yes: investigate
 - No: move on
- This helps him not to be misled by natural run variation



[This Photo](#) by Unknown Author is licensed under [CC BY](#)

Summary



Regression performance our clients expect



Clients should be able to run a 4 TB partition just fine



Nice improvement for
dumps

More and more relevant as clients'
systems become larger



Nice improvement for guest PCIe workloads



Nice improvement from OCSP



We are coping with run variation