Capture ratios for SMF

Barton Robinson, CTO & Founder
barton@velocitysoftware.com

Velocity Software Inc.
196-D Castro Street
Mountain View CA 94041
650-964-8867

Velocity Software GmbH
Max-Joseph-Str. 5
D-68167 Mannheim
Germany
+49 (0)621 373844

Agenda

- Who Am I
- Why do I care
  - Bogus data in = bogus data out….
- Where is data from
  - Platform standard interfaces
- What did I learn
  - PRSM,z/VM,z/VSE,Linux,z/OS,CICS,DB2
- When do we meet in person?

VELOCITY SOFTWARE

Who Am I

Washington System Center (3 years)

- VM Performance responsibilities
- Author "HPO Tuning Guide"

Senior Manager IBM Systems Evaluation Lab

Founder / CTO Velocity Software (33 years)

Performance Management Development for:

- VM/XA, VM/ESA, z/VM
- Linux, WAS, Oracle
- Network data
- z/VSE
- Now, z/OS, CICS, DB2

VELOCITY SOFTWARE

My Focus

Providing Correct Data for System Performance Mgmt:

- Capacity Planning
- Performance Analysis
- Chargeback/Accounting
- Operational Alerts

Business decisions are (hopefully) made based on data

- Better decisions are made on **correct data**…
- Validate the data (Challenge very old "traditional wisdom")
- Understand what is missing – and how much

- When Linux first virtualized, Linux reported CPU incorrectly by up to 2 ORDERS OF MAGNITUDE…. (conclusion: mainframe bad)
- SMT does not exactly add up…. (topic for different presentation)

VELOCITY SOFTWARE

What is (my) definition of capture ratio?

Objective is to know where 100% of resource is used

- System management time ("Physical" overhead)
- Workload management time ("logical" overhead)
- Workload
- IDLE time
- Uncaptured (hopefully zero)

Does platform instrumentation provide 100%?

- PRSM / LPAR: yes
- z/VM: yes
- Linux: yes
- VSE: yes
- z/OS???? DB2?? CICS??

What is the overhead of the platform?

VELOCITY SOFTWARE

Why care about capture ratios?

Capture ratios validate the data and instrumentation

CPU Data has multiple data sources – do they agree?

- If not, what was missed? Validate the instrumentation …
- PRSM / LPAR – Assigned time vs Operating System reported utilization
  - z/OS smf 70 – what fields show true system overhead?
  - z/VM monitor sytprp – provides measured system overhead

LPAR (HMC data) provides instrumentation for:

- Physical Overhead
- Assigned time
  - Logical Overhead
  - Virtual Assigned time (The Real Work)
- Non-captured time at next level, not reported – about 1%
- But the analysis was very interesting???

VELOCITY SOFTWARE

How to Get Capture Ratio

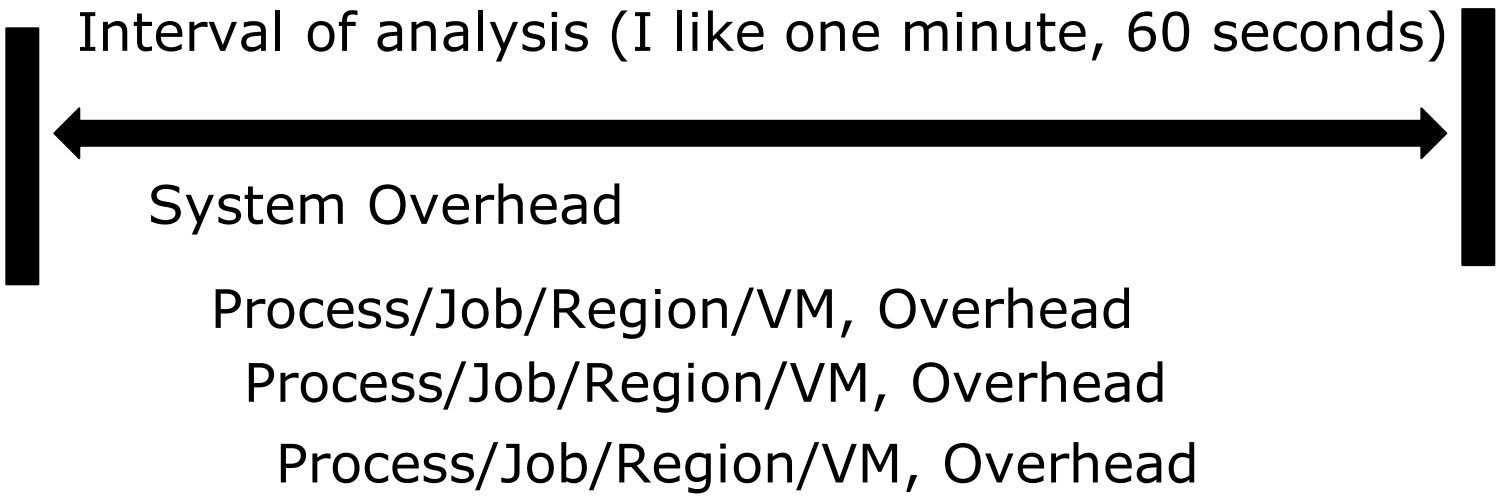

Capture Ratio –

- Accumulate all interval process data available
- Compare to system total values
- (Inconsistent intervals make life difficult)

Watch for my analysis of CICS, DB2, MQ for next year

VELOCITY SOFTWARE

Capture Ratio Process

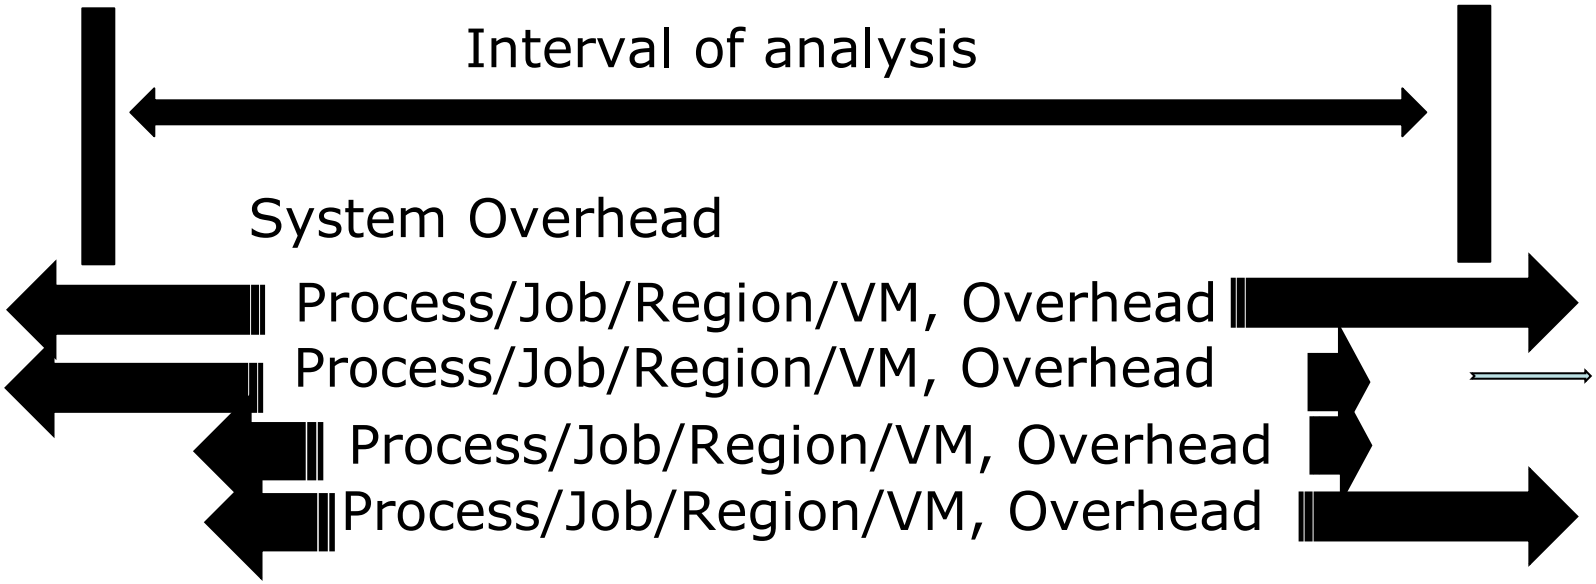

Workload Capture Ratio – Always 4 types

- Work that starts prior to interval and ends after
- Work that starts prior to interval and ends during
- Work that starts within interval and ends during
- Work that starts within interval and ends after

VELOCITY SOFTWARE

Compare Multiple Data Sources

CPU Data has multiple data sources

Do they agree? If not, what was missed?

- z/VM:
  - LPAR data (SYTCUP, SYTCUM)
  - z/VM System CPU (sytprp)
  - z/VM User / Virtual Machine CPU (USEACT,USELOF)
  - Hardware PRCMFC (SMF 113)
- Linux (virtualized linux cpu data was bogus…)
  - Virtual machine data
  - Kernel cpu / irq cpu
  - Process data
- VSE
  - Virtual machine data (normally)
  - System data
  - Partition data

VELOCITY SOFTWARE

z/OS Data Sources

z/OS Data source is SMF records

Very highest level is from PRSM / LPAR:

- LPAR data Physical Time (SMF 70)
- LPAR Logical Overhead (SMF 70)
- LPAR Logical Time (SMF 70)
- Hardware SMF 113 – measures cycles consumed

z/OS Data Comparisons:

- LPAR Logical Time (smf 70)
- z/OS Job data (SMF 30)
- Hardware (SMF 113)
- CICS (SMF 110)
- DB2 (SMF 100/102)

NOTE: All metrics in absolute terms of one CPU

VELOCITY SOFTWARE

Transaction Analysis

CICS / DB2 For example

- Many Transactions
- Long running "system" transactions
- Some transactions cross interval start or end

Objective

- Capture all transaction CPU (100%) in interval
- What transaction ID is consuming the CPU
- Group for capacity planning
- Group transactions for chargeback
- What is missing and why? Iterative process…

VELOCITY SOFTWARE

Building the "Capture Model"

Every platform has 5 CPU Components

- Hypervisor/OS Management Time (physical overhead)
- Work Management (logical overhead)
- Work time
- IDLE, vs steal time
- Uncaptured – Platform does not define or report

Steal Time

- Virtualized environment, underlying CPU "stolen"
- Not relevant for capture analysis, CPU not utilized

PARK Time

- Not relevant (to me) for capture analysis – CPU not utilized
- But measure unparked time and cycles consumed

Objective is to know what / who is using CPU

VELOCITY SOFTWARE

z/vm Capture Ratio model

5.2 billion cycles per second per cpu

- Where did they all go?
- Set interval = 1 minute to understand variations

For every platform, objective is to show:

- System overhead – Not related to applications
- Application associated overhead
- Application CPU
- Uncaptured – to be identified, objective is zero

If uncaptured CPU is zero (or very low)

- Platform is fully instrumented
- Data can be "trusted" for business decisions
- No "guessing" or "crystal balls"

VELOCITY SOFTWARE

Building the "Capture Model"

Every platform has 4 CPU Components plus IDLE

At end of day, what do we know for sure?

	Mgmt	Logical	Work	Uncaptured
HMC/LPAR				
z/VM				
z/VSE				
Linux				
z/OS				
CICS				
DB2				

VELOCITY SOFTWARE

Capture Ratio – LPAR Level

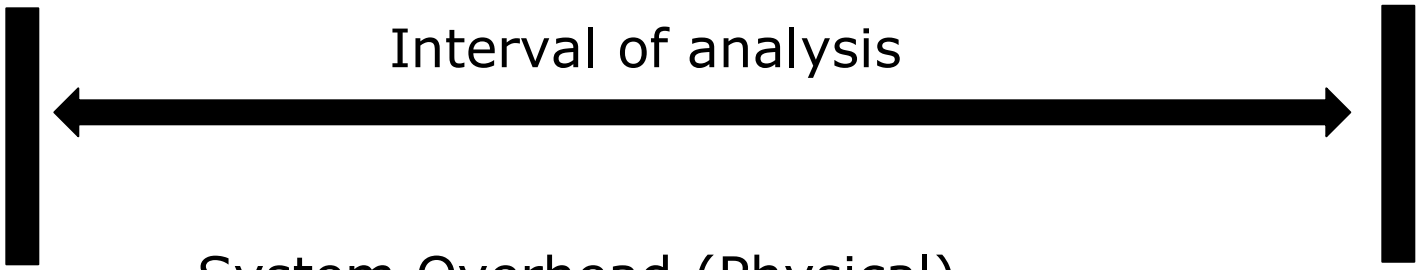

System Overhead (Physical)

LPAR Consumption VSI1, Overhead

LPAR Consumption VSI2, Overhead

LPAR Consumption VSI3, Overhead

Capture Ratio – Highest level

- Accumulate all LPAR data by CPU TYPE
- LPAR data Highest level – simplest case
- LPARs don't "start" and "end" during intervals

VELOCITY SOFTWARE

z/VM (and z/OS) LPAR / CEC Data Source

Every z has LPAR data, **One record per VCPU:**

- Assigned Time to LPARs: SYTCUP.LCUCACTM
- LPAR Time (exclude ovhd: SYTCUP.LCUCLPTM
- Add data by LPAR, by Engine Type

<--------Logical Partition------->						
		Virt	CPU	<%Assigned>		
Name	Nbr	CPUs	Type	Total	Ovhd	
VSIVM5	05	2	CP	84.2	0.0	(VSE, z/OS)
VSIVM5	05	2	IFL	1.5	0.1	
VSIVC1	07	1	IFL	23.3	0.1	
VSIVC2	08	1	IFL	0.7	0.0	
VSIVC3	09	1	IFL	0.5	0.0	
VSIVC4	0A	1	CP	2.4	0.0	(VSE)
VSIVC4	0A	1	IFL	0.5	0.1	(linux)
VSIVM1	01	1	IFL	1.3	0.0	
VSIVM2	02	1	IFL	1.2	0.0	
VSIVM3	03	1	IFL	0.4	0.0	
VSIVM4	04	2	IFL	85.5	0.4	
ZOSLP1	0E	2	CP	56.6	0.0	(z/OS)
ZOSLP2	0F	2	CP	56.7	0.0	(z/OS)

VELOCITY SOFTWARE

LPAR / CEC Capture Ratio model

Full picture of CEC, Add by CPU TYPE (z/VM Model)

- Physical Overhead : SYTCUM.LCUMGTM
- Assigned Time to LPARs: SYTCUP.LCUCACTM
- LPAR Time (exclude ovhd: SYTCUP.LCUCLPTM

Working example (LPARs for z/VM, z/OS, cloud)

Totals by Processor type:

<---------CPU------->				<-Shared Processor busy->			
Type	**Count**	**Ded**	**shared**	**Total**	**Logical**	**Ovhd**	**Phys**
CP	**2**	**0**	**2**	**200.0**	**199.8**	**0.1**	**0.1**
IFL	**4**	**0**	**4**	**116.0**	**114.0**	**0.8**	**1.1**
ZIIP	**1**	**0**	**1**	**0.6**	**0.5**	**0.0**	**0.0**

CEC Level LPAR Capture ratio – 100%

We DO Know What LPAR consumes the CPU

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low" – set an alert, high overhead happens
- Capture Ratio 100%
- We know exactly what LPAR is consuming what….

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.1%	.1%	99%+	0%	100%
z/VM					
z/VSE					
Linux					
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

CP Monitor provides **One record per CPU/Thread:**

- System CPU: sytprp.pfxtmsys (physical overhead (1%)
- User Ovhd: sytprp.pfxutime – sytprp.pfxprbtm (1-2%)
  - (Same concept as PRSM, total assigned time , logical assigned time)
- User CPU: sytprp.pfxprbtm
- IDLE: sytprp.pfxtotwt
- Steal: 100 – (system cpu + user cpu – idle)

		<--CPU (percentages)-->			
CPU	CPU Type	Total util	Emul time	User ovrhd	Sys ovrhd
-	----	-----	-----	-----	-----
0	IFL	47.0	45.9	0.6	0.4
1	IFL	50.0	48.9	0.7	0.4
2	IFL	45.5	44.4	0.7	0.4
3	IFL	47.3	46.1	0.8	0.4
4	IFL	42.5	41.0	0.8	0.7
5	IFL	53.6	52.7	0.6	0.3
6	IFL	44.3	43.3	0.6	0.4
7	IFL	56.3	55.3	0.6	0.3
		-----	-----	-----	-----
		386.4	377.7	5.4	3.4

VELOCITY SOFTWARE

z/vm Capture Ratio Hierarchy

First level Platform analysis

- PRSM / HMC (SYTCUP) vs z/VM (SYTPRC)

Compare Assigned time to reported utilization

- PRSM reports by LPAR by VCPU
- z/VM reports by VCPU

VELOCITY SOFTWARE

CP Monitor CPU vs PRSM?

- LPAR / PRSM data 100%, What does z/VM see?
- LPAR Data vs z/VM CPU Data: 99.3% (for every CPU…)
- Discrepancy likely setting up and dispatching

<PRSM / LPAR Measurements>					<---z/VM-CPU (percentages)->					VM/
VCPU Addr	CPU Type	<---%Assigned--> Total	Ovhd	Emul	Total util	Emul time	User ovrhd	Sys ovrhd	Stl Pct	PRSM CaptR
0	IFL	62.4	0.7	61.6	61.2	58.4	1.2	1.6	8.71	0.99
1	IFL	62.4	0.6	61.7	61.3	58.6	1.1	1.5	8.59	0.99
2	IFL	62.2	0.6	61.7	61.3	58.7	1.1	1.5	8.20	0.99
.....										
8	IFL	62.3	0.7	61.5	61.1	57.6	1.2	2.2	8.70	0.99
9	IFL	62.6	0.8	61.7	61.3	58.6	1.2	1.5	8.66	0.99
10	IFL	62.5	1.0	61.5	61.1	58.3	1.2	1.6	8.82	0.99
11	IFL	62.6	0.6	62.0	61.6	59.0	1.1	1.5	8.60	0.99
12	IFL	62.5	0.9	61.6	61.2	58.4	1.2	1.6	8.77	0.99
13	IFL	62.2	0.8	61.5	61.1	58.4	1.2	1.5	8.67	0.99
Total	IFL	873.5	10.5	863.1	857.2	818.6	16.4	22.2	121	0.99

VELOCITY SOFTWARE

Charge back model is NOT 100%

Data for chargeback requires "fudge factor"

- PRSM Overhead: 1% ?
- LPAR Overhead: 2%?
- LPAR Capture ratio: 1% (capture ratio 99%)
- z/VM System overhead
- z/VM virtual machine overhead
- Virtual machine real work – this is what we charge for

What does SMT do? Not a top here…

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low" – set an alert, high overhead happens
- Capture Ratio 100%
- We know exactly what LPAR is consuming what….

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.1%	.1%	99%+	1%	99%
z/VM					
z/VSE					
Linux					
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

z/VM Capture Ratio model

Compare "system data" to "virtual machine data"

One Record per CPU / Thread

- Virtual Machine "user" CPU Time: sytprp.pfxutime
- VM Problem Time: sytprp.pfxprbtm
- User Overhead: pfxutime - pfxprbtm
- System overhead: sytprp.pfxtmsys
- Idle, "steal"

One Record per Virtual Machine VCPU

- Virtual Machine CPU Time by VM: USEACT.VMDTTIME
- Problem (Virtual) Time: USEACT.VMDVTIME
- USELOF: Logoff
- Add up all the virtual machines, Compare:
- Note: (USEACT.VMATTIME_PRO – SMT Prorated – off topic)

VELOCITY SOFTWARE

CP Monitor provides **One record per CPU/Thread:**

- System CPU: sytprp.pfxtmsys (physical overhead (1%)
- User Ovhd: sytprp.pfxutime – sytprp.pfxprbtm (1-2%)
  - (Same concept as PRSM, total assigned time , logical assigned time)
- User CPU: sytprp.pfxprbtm
- IDLE: sytprp.pfxtotwt
- Steal: 100 – (system cpu + user cpu – idle)

		<--CPU (percentages)-->			
CPU	CPU Type	Total util	Emul time	User ovrhd	Sys ovrhd
-	----	-----	-----	-----	-----
0	IFL	47.0	45.9	0.6	0.4
1	IFL	50.0	48.9	0.7	0.4
2	IFL	45.5	44.4	0.7	0.4
3	IFL	47.3	46.1	0.8	0.4
4	IFL	42.5	41.0	0.8	0.7
5	IFL	53.6	52.7	0.6	0.3
6	IFL	44.3	43.3	0.6	0.4
7	IFL	56.3	55.3	0.6	0.3
		-----	-----	-----	-----
		386.4	377.7	5.4	3.4

CP Monitor, **One record per Virtual Machine/CPU:**

- Add up all the virtual machines, Compare:
- User overhead: Total assigned – total virtual
- User overhead = 383.1 – 377.7 = 5.4
- Add up all users for totals

```
<---CPU time--->
UserID    <(Percent)> T:V
/Class    Total  Virt  Rat
-------- ----- ----- ----
09:01:00 383.1 377.7 1.01

WASM8096 82.95 82.86
WWAS8042 28.56 28.24 1.01
WWAS8038 25.22 24.97 1.01
WWAS8046 24.45 24.20 1.01
WWAS8000 23.82 23.51 1.01
WWAS8005 23.48 23.15 1.01
DB2M8002 23.18 22.81 1.02
.
:
:
```

```
<--CPU (percentages)-->
     CPU    Total   Emul   User    Sys
CPU  Type    util   time  ovrhd  ovrhd
 -   ----   -----  -----  -----  -----
 0   IFL     47.0   45.9    0.6    0.4
 1   IFL     50.0   48.9    0.7    0.4
 2   IFL     45.5   44.4    0.7    0.4
     IFL     47.3   46.1    0.8    0.4
     IFL     42.5   41.0    0.8    0.7
 5           53.6   52.7    0.6    0.3
 6   IFL     44.3   43.3    0.6    0.4
 7   IFL      .3    55.3    0.6    0.3
                   -----  -----  -----
            386.4  377.7    5.4    3.4
```

VELOCITY SOFTWARE

z/VM Capture Ratio Analysis

System data by CPU (sytprp)

Virtual Machine Data (USEACT)

Virtual Machine logon (Delta "zero", vmdttime)

Virtual Machine interval (Delta start,stop vmdttime)

Virtual Machine logoff (Delta start, logoff vmdttime)

	Syst Ovhd	Emulation	User Ovhd
sytprp	3.4	377.7	5.4
USEACT		377.7	5.4

Capture Ratio from z/VM is 100.0% consistently

- **Two data sources agree**

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low" – set an alert, high overhead happens
- Capture Ratio 100%
- We know exactly what LPAR is consuming what….

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.1%	.1%	99%+	1%	99%
z/VM	< 1%	<2%	97%+	0%	100%
z/VSE					
Linux					
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

z/vm Capture Ratio Hierarchy

CEC/LPAR Management

- z/VM Reporting (CP Monitor)
  - Guest Reporting
    - CMS Virtual Machines
    - VSE: Partitions (GP processors)
    - Linux: Processes (IFL processors)
    - z/OS Guests, regions

CEC/LPAR Management

- z/OS Reporting (SMF 70)
  - Region / Job Reporting (smf 30)
    - CICS: Transactions (110)
    - DB2: Transactions (100,102)

VELOCITY SOFTWARE

zVSE Capture Ratio model

zVSE Virtual Machine Metrics Technology

- z/VM Guest machine data provided by IBM CP Monitor
- (USEACT.VMDTTIME – Traditional, non SMT)
- (useact.vmdvtime) – VSE Dispatch time

zVSE Metrics Technology – 4 data sources

- System data provided by IBM via SNMP
- Virtual Processor data provided by IBM via SNMP
- Partition data provided by IBM's SNMP
- Job Data provided by Velocity's SNMP

Note: snmp data does not perfectly align with z/VM data collection interval.

VELOCITY SOFTWARE

zVSE Capture Ratio model

Compare z/VM Virtual Machine CPU to VSE System

- Is VSE Capturing all CPU?

Compare VSE System data to Partition data

- Is data model complete?

Compare VSE Job data to VSE System data

- Is chargeback model complete?

VELOCITY SOFTWARE

zVSE z/VM Capture Ratio model

zVM Metrics Technology – accurate to microsecond

- Virtual machine data – minute by minute, exactly on minute!

z/VM Virtual Machine management CPU

- T:V Ratio, total to virtual, shows overhead (I/O, memory)
- z/VM MANAGEMENT OVERHEAD: about 2% for this guest

	<-VM-CPU time->		
UserID	<(Percent)>		T:V
/Class	Total	Virt	Rat
ZVSE61B	30.58	30.04	1.02
ZVSE61B	30.44	29.87	1.02
ZVSE61B	30.29	29.71	1.02
ZVSE61B	29.29	28.71	1.02
ZVSE61B	29.00	28.46	1.02
ZVSE61B	28.47	27.94	1.02

VELOCITY SOFTWARE

zVSE z/VM Capture Ratio model

zVSE Analysis, z/VM to VSE System

- z/VM Virtual Machine data accurate to microsecond
- Both average over 6 minutes: 29.1%
- Capture Ratio: 99.9++% (over time)
- VSE Captures CPU as provided by z/VM

	<---CPU time--->		
UserID	<(Percent)>		T:V
/Class	Total	Virt	Rat
ZVSE61B	30.58	30.04	1.02
ZVSE61B	30.44	29.87	1.02
ZVSE61B	30.29	29.71	1.02
ZVSE61B	29.29	28.71	1.02
ZVSE61B	29.00	28.46	1.02
ZVSE61B	28.47	27.94	1.02

<---CPU time--->	
NODE	<CPU
/Time	Total
zvse61b	27.3
zvse61b	30.6
zvse61b	30.2
zvse61b	29.3
zvse61b	29.2
zvse61b	28.1

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.1%	.1%	99%+	1%	99%
z/VM	1-2%	1-2%	95%+	0%	100%
z/VSE	<.1%				99.9%
Linux					
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

Capture Ratio – z/VSE

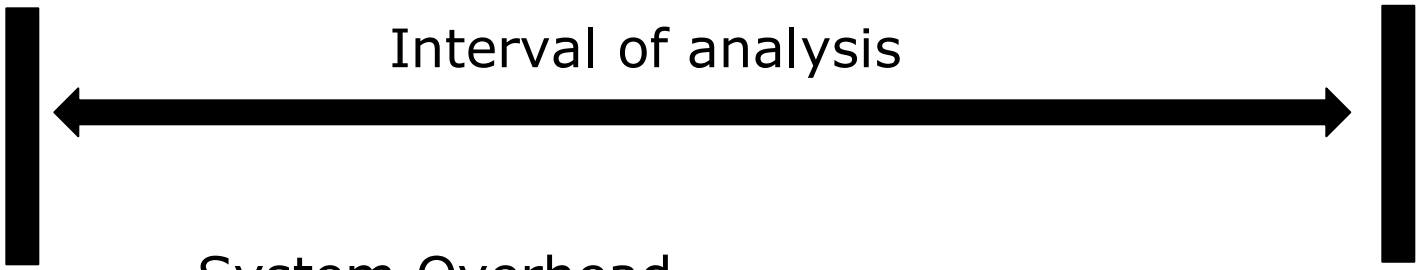

Capture Ratio –

- Accumulate all process data,
- compare to system
- Inconsistent intervals make life difficult

VELOCITY SOFTWARE

VSE Partition Metrics, total the partitions

- CPU for "work", Overhead

NODE /Time	Part ID	Job Name	Phase Name	<-CPU CPU	Pct--> Overhd
ZVSE61B		Totals		21.7	1.4
	FB	SECSERV	BSTPSTS	0	0
	F1	POWSTART	IPWPOWER	0.0	0.0
	F2	CICSICCF	DFHSIP	0.0	0.0
	F3	VTAMSTRT	ISTINCVT	0.0	0.0
	O1	CICSJA60	DFHSIP	18.7	1.2
	R1	STARTVCS	IESVCSRV	0.0	0.0
	R2	STARTMAS	IESMASNM	0.3	0.0
	T1	BSTTINET	BSTTINET	1.4	0.1
	T2	BSTTVNET	BSTTVNET	1.3	0.1
	T3	BSTTFTPD	BSTTFTPS	0	0
	Z1	DMFSTART	DFHDFSIP	0.1	0.0

VELOCITY SOFTWARE

zVSE z/VM Capture Ratio model

zVSE Analysis

- z/VM Virtual Machine data accurate to microsecond
- VSE System Data (via IBM snmp)
- VSE Partition data (via VSI snmp)
- Intervals are one minute, but not synchronized exactly
- Compare (partition total plus ovhd) to (total)
- PARTITION DATA Capture Ratio: 99.9% (over time)

NODE /Time	<CPU Utilization> Total	Mstr	Spin	<Partn CPU> TOTal	Ovhd	Capture Ratio
zvse61b	27.3	4.5	0	25.6	1.6	99.7
zvse61b	30.6	5.1	0	28.5	1.8	99.0
zvse61b	30.2	5.2	0	28.5	1.8	100.5
zvse61b	29.3	5.2	0	27.7	1.8	100.8
zvse61b	29.2	5.0	0	27.4	1.7	99.7
zvse61b	28.1	4.9	0	26.3	1.7	99.9

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	Logical (srb)	Work	Uncaptured	Capture Ratio
HMC/LPAR	.5-1%	.5-1%	98%+	1%	99%
z/VM	1-2%	1-2%	95%+	0%	100%
z/VSE	<.1%	6-8%	92-94%	0%	99.9%
Linux					
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

Linux Capture Ratio model

Linux data captured via snmp

- System CPU Data by cpu, by system:
- Process Data by process

System data provides

- IRQ, SoftIRQ, Kernal,
- Nice

Process data provides

- CPU data by process, for process and "children"
- Parent process information

Challenge in Linux when process terminates

- CPU added to parents when process terminates

VELOCITY SOFTWARE

Capture Ratios – Linux

Linux system data vs z/VM data

- Linux Includes IRQ, Krnl time (2%)
- Linux collection time 5-10 seconds prior to z/vm monitor pop

z/VM time (78%)

<---CPU time-->			
UserID /Class	<(Percent)> Total	Virt	T:V Rat
RLNX08P0	78.64	74.67	1.1
RLNX08P0	72.33	66.01	1.1
RLNX08P0	53.09	48.31	1.1
RLNX08P0	61.48	56.38	1.1
RLNX08P0	84.47	79.56	1.1
RLNX08P0	93.25	88.30	1.1
RLNX08P0	120.7	116.7	1.0
RLNX08P0	96.25	91.80	1.0
RLNX08P0	83.71	78.61	1.1

Linux time (78%)

<Processor Total	Pct Syst	User	<CPU Overh Krnl	IRQ
78.5	5.5	71.5	0.3	1.3
80.3	5.6	73.0	0.4	1.3
41.0	4.9	34.6	0.3	1.1
63.4	7.2	54.0	0.4	1.7
68.4	5.8	61.0	0.3	1.3
65.5	5.3	59.1	0.3	0.9
127.7	7.1	119	0.4	1.5
98.0	6.5	89.7	0.4	1.4
79.6	6.5	71.4	0.4	1.4

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	**Logical (srb)**	**Work**	**Uncaptured**	**Capture Ratio**
HMC/LPAR	.5-1%	.5-1%	98%+	1%	99%
z/VM	1-2%	1-2%	95%+	0%	100%
z/VSE	<.1%	6-8%	92-94%	<1%	99.9%
Linux	2%				
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

Capture Ratios – Linux

Capture ratio concept for Linux process table

- **When Linux process terminates, where does CPU go? – the**
- **Does "crond" get charged anything? No, "children"**
- **Must build process tree**

node/ Name	<-Process Id ID	PPID	<------CPU Percents----> Tot	sys	user	syst	usrt
snmpd	1919	1	0.10	0.08	0.02	0	0
crond	2116	1	0.03	0	0	0.02	0.02
seosd	2515	1	0.02	0.02	0	0	0
selogrd	2549	1	0.02	0	0.02	0	0
db2sysc	3097	3095	1.24	1.01	0.24	0	0
db2fmp	3118	3095	0.02	0.02	0	0	0
db2sysc	3248	3246	0.05	0.02	0.03	0	0
dsmc	30061	1	0.02	0	0.02	0	0

Node/ Name	<----Proces ID	PPID
init	1	1
kthreadd	2	1
migratio	3	2
crond	2116	1
crond	30034	2116
sh	30035	30034
sendmail	30086	30034
postdro	30087	30086
db2syscr	3095	1
db2sysc	3097	3095
db2syscr	3103	3095
db2syscr	3104	3095
db2syscr	3105	3095
db2vend	3107	3095
db2fmp	3118	3095
db2syscr	3246	1
db2sysc	3248	3246
db2syscr	3254	3246
db2syscr	3255	3246
db2syscr	3256	3246
db2vend	3258	3246
db2fmp	3266	3246
login	3326	1
bash	3332	3326

VELOCITY SOFTWARE

Capture Ratios – Linux

Capture ratio concept for Linux process table

- Compare "linux system data" to "Linux Process Data"
- Typically 100%....
- Collecting 1000 processes synchronously has "variation"…
- "system Time" 7-10% ?

Node/ Name	<Linux Pct CPU> Total	Syst	User	<Process Data> Total	Syst	User	Capture Ratio
RLNX01p1	1.7	0.8	0.9	1.8	0.9	0.9	1.056
RLNX02p1	1.3	1.0	0.3	1.3	1.0	0.3	1.000
RLNX03p0	2.0	0.6	1.4	2.0	0.6	1.4	1.000
RLNX04p0	10.6	1.7	8.9	10.7	1.8	8.9	1.007
RLNX05p0	9.0	1.5	7.5	9.0	1.5	7.5	1.000
RLNX06p0	11.6	1.7	9.9	11.6	1.7	9.9	1.000
RLNX07p0	18.3	2.9	15.4	18.3	2.9	15.4	1.000
RLNX08p0	78.4	6.3	72.0	79.8	6.7	73.0	1.018

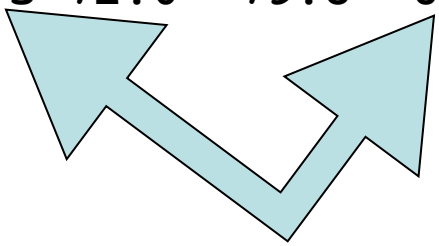

VELOCITY SOFTWARE

Linux Capture Results

Linux process data, applications

- Capture ratio 100%
- Applications identified and quantified
  - (db2 vs was vs oracle)
- System overhead is provided by standard metrics
- Process Data straight from Linux process table
- Arithmetic is complicated but doable

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.5-1%	.5-1%	98%+	0%	100%
z/VM	1-2%	1-2%	95%+	0%	99%
z/VSE	<.1%	6-8%	92-94%	0%	99.9%
Linux	2%	7-10%	90%	< 1%	99% +
z/OS					
CICS					
DB2					

VELOCITY SOFTWARE

z/OS Capture Ratio model

For every platform, objective is to show:

- System overhead – Not related to applications
- Application associated overhead
- Application CPU
- Uncaptured – to be identified, objective is zero

If uncaptured CPU is zero:

- Platform is fully instrumented
- Data can be "trusted" for business decisions
- No "guessing" or "crystal balls"

If multiple data sources, validate….

VELOCITY SOFTWARE

z/OS Capture Ratio model

At each layer, something is lost

PRSM / LPAR Layer

- PRSM Overhead (Physical overhead)
- PRSM Assigns a physical CPU (smf70pdt)
- PRSM has Logical overhead (pdt-edt)
- LPAR is dispatched (smf70edt)

z/OS Layer

- Uncaptured: Smf 113 metrics shows less cpu than smf70edt
- System overhead?
- Regions / jobs (smf 30) shows uncaptured

VELOCITY SOFTWARE

z/OS Capture Ratio model

At each layer,
something is lost

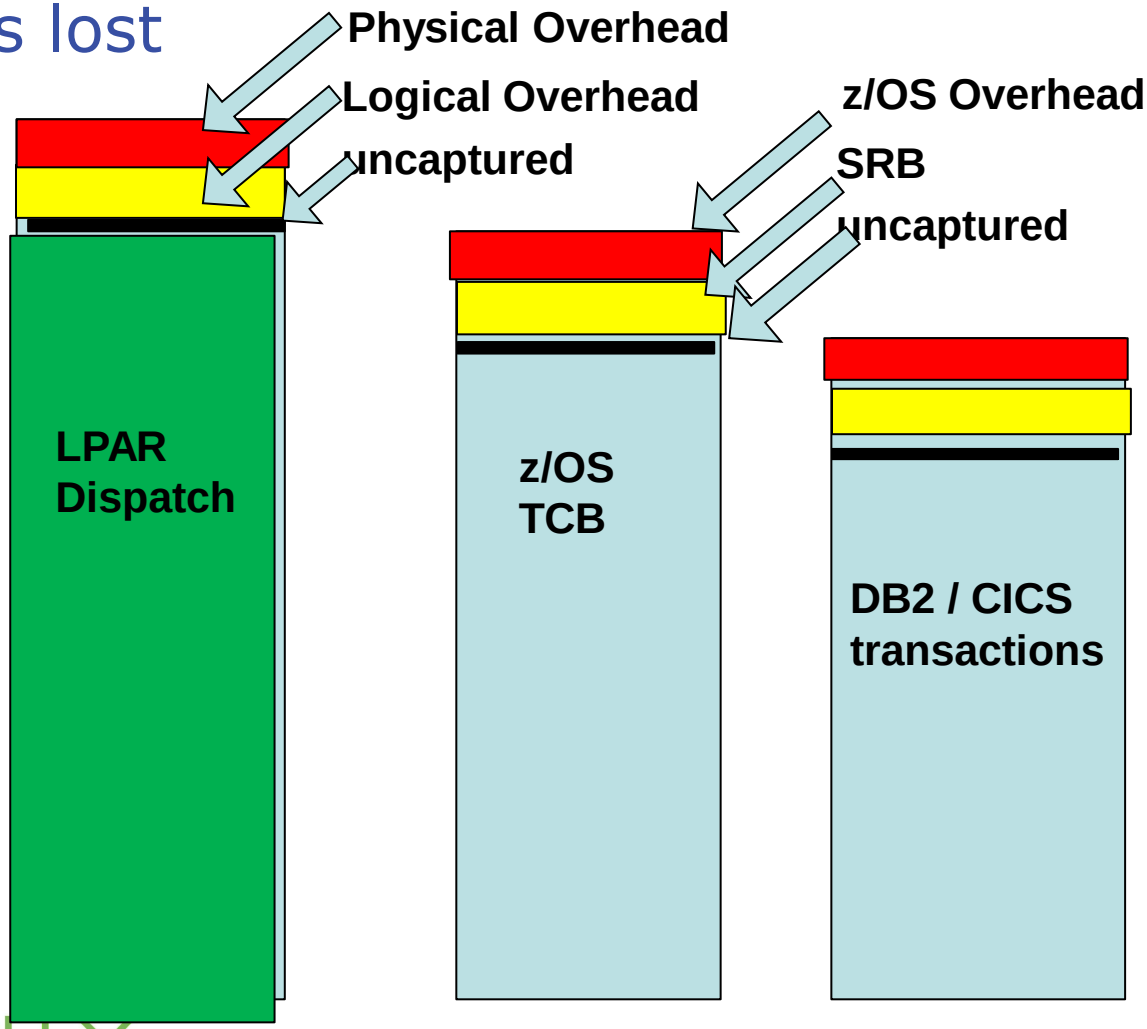

VELOCITY SOFTWARE

z/OS Capture Ratio model

SMF 70 Subtype 1

- CPU Data:
  - SMF70WAT (wait / idle time?)
- LPAR Data
  - SMF70PDT (assigned, includes logical overhead)
  - SMF70EDT (the real work)

Since z/OS CPU data comes from PRSM

- Can't compare "z/OS CPU data" to "PRSM CPU data"
- Can compare "z/OS CPU data" to hardware data
- SMF 113 comparison shows a loss of about 3%

- This would be 3% uncaptured time from PRSM to z/OS

VELOCITY SOFTWARE

z/OS Capture Ratio model

SMF 113 analysis, PRSM CPU vs Cycles consumed

<Percent>	<-------Processor------>						
		<CPUBusy>		Speed/<-Rate/Sec->			
Time	CPU	LPAR	MFC	Hertz	Cycles	Instr	Ratio
J90	0	79.7	78.5	5200M	4080M	1416M	2.882
	2	82.0	79.9	5200M	4156M	2025M	2.053
	4	67.4	65.4	5200M	3400M	2043M	1.665
	6	68.4	65.7	5200M	3417M	1345M	2.540
	.....						
	24	46.7	43.7	5200M	2273M	1466M	1.550
	26	18.2	16.6	5200M	861M	324M	2.660
	28	10.9	9.8	5200M	509M	161M	3.169
	.....						
	36	0.6	0.0	5200M	65265	465.5	140.2 ←low entitled
	38	0.1	0.0	5200M	65368	459.1	142.4
	40	39.1	37.8	5200M	1963M	1715M	1.145 <-ziip
	.....						
	48	20.6	19.9	5200M	1033M	900M	1.148
		----	----	------	-----	-----	-----
GPCPUs:		846	824	5200M	42.8G	17.8G	2.407
zIIPs		142	138	5200M	7200M	6212M	1.159

VELOCITY SOFTWARE

z/OS Capture Ratio model

SMF 30 – Next step is to compare jobs data:

SYSID	<------JOB----------> Name	JobID	Step Nbr	<CPU Percents--- Total	STD	SRB
VSI1	Totals		.	24.88	21.5	3.2
	CATALOG	CATALOG	1	0.02	0.0	0
	CICSJZ1	STC01591	1	0.05	0.0	0
	CICSJZ2	STC08736	1	20.20	19.7	0.5
	JES2	JES2	1	0.22	0.2	0
	JES2MON	JES2MON	1	0.17	0.1	0.0
	JES2S001	JES2S001	1	0.02	0.0	0
	MSTJCL00	MSTR	1	0.03	0.0	0.0
	RMF	STC08659	1	0.13	0.1	0
	SDSF	STC08666	1	0.02	0.0	0
	SDSFAUX	STC08682	1	0.10	0.1	0.0
	SMF	SMF	1	0.03	0.0	0
	SMS	SMS	1	0.05	0.0	0.0
	SYSLOGD	STC08661	1	0.03	0.0	0.0
	TCPIP	STC08662	1	0.95	0.2	0.8
	TN3270	STC08696	1	1.78	0.1	1.7
	WLM	WLM	1	0.30	0.2	0.1
	ZOSMNVM2	STC00547	1	0.12	0.1	0
	ZOSMNVM4	STC00548	1	0.13	0.1	0

z/OS Capture Ratio model

SMF 70 Subtype 1

- CPU Data:
  - SMF70WAT (wait / idle time?)
- LPAR Data
  - SMF70PDT (assigned, includes logical overhead)
  - SMF70EDT (the real work)

SMF 30 – add up all jobs during interval

- Job Data CPU =
  - smf30CPT + /* General*/
  - smf30CPS + /*CPU Time SRB*/
  - smf30ICU + /*InitiatorTCB*/
  - smf30ISB + /*InitiatorSRB*/
  - smf30IIP + /*I/O */
  - smf30RCT + /*RegionCntrl */
  - smf30HPT ; /*Hyperexstore*/

VELOCITY SOFTWARE

z/OS Capture Ratio model

SMF 30 CPU Capture

- Types 2,3,6 capture all CPU for an interval
- Must respect 4 activity type models

Break out CPU BY Function

- Job/Region CPU Data = SUM
  - smf30CPT + /* General CPU*/
  - smf30CPS + /*CPU Time SRB*/
  - smf30ICU + /*InitiatorTCB*/
  - smf30ISB + /*InitiatorSRB*/
  - smf30IIP + /*I/O */
  - smf30RCT + /*RegionCntrl */
  - smf30HPT ; /*Hyperexstore*/

If capture ratio 100%, Chargeback model is valid

VELOCITY SOFTWARE

SMF 30 – jobs data:

SYSID	<------JOB----------> Name	JobID	Step Nbr	<CPU Percents--- Total	STD	SRB
VSI1	Totals		.	24.88	21.5	3.2
	CATALOG	CATALOG	1	0.02	0.0	0
	CICSJZ1	STC01591	1	0.05	0.0	0
	CICSJZ2	STC08736	1	20.20	19.7	0.5
	JES2	JES2	1	0.22	0.2	0
	JES2MON	JES2MON	1	0.17	0.1	0.0
	JES2S001	JES2S001	1	0.02	0.0	0
	MSTJCL00	MSTR	1	0.03	0.0	0.0
	RMF	STC08659	1	0.13	0.1	0
	SDSF	STC08666	1	0.02	0.0	0
	SDSFAUX	STC08682	1	0.10	0.1	0.0
	SMF	SMF	1	0.03	0.0	0
	SMS	SMS	1	0.05	0.0	0.0
	SYSLOGD	STC08661	1	0.03	0.0	0.0
	TCPIP	STC08662	1	0.95	0.2	0.8
	TN3270	STC08696	1	1.78	0.1	1.7
	WLM	WLM	1	0.30	0.2	0.1
	ZOSMNVM2	STC00547	1	0.12	0.1	0
	ZOSMNVM4	STC00548	1	0.13	0.1	0

z/OS Capture Ratio model

SMF 70 Subtype 1

- CPU Data:

Break out CPU BY Function

- TCB
- Local SRB
- Global SRB
- Supervisor
- Program

If uncaptured CPU is zero:

- Platform is fully instrumented
- Data can be “trusted” for business decisions
- No “guessing” or “crystal balls”

VELOCITY SOFTWARE

z/OS Jobs Capture Ratio model

SMF 70, CPU Utilization

TIME/ SYSID	<--CPU--> ID	Type	Sample Count	<CPU Total
VSI1	0	GP	1	17.2
	1	GP	1	9.6
	Tot	GP	2	26.9

SMF 30, CPU Utilization capture = 92%

- Logical overhead 3.2 / 24.88 = 13%

SYSID	<------JOB----------> Name	JobID	Step Nbr	<CPU Percents-- Total	STD	SRB
VSI1	Totals		.	24.88	21.5	3.2
	CATALOG	CATALOG	1	0.02	0.0	0
	CICSJZ1	STC01591	1	0.05	0.0	0
	CICSJZ2	STC08736	1	20.20	19.7	0.5
	IOSAS	IOSAS	1	0.03	0.0	0.0
	IXGLOGR	IXGLOGR	1	0.03	0.0	0.0
	IZUSVR1	STC08689	3	0.10	0.1	0

VELOCITY SOFTWARE

z/OS Jobs Capture Ratio model

SMF 30 / 70 Ratios over time?

- Ratio consistent across LPARs 92%, 88%
- No effective measurement of system overhead???

SYSID	SYSPLEX Name	<-Physical <CPU Util> Total	Wait	Capt Ratio	<Job CPU Total
16:51:00 - 16:52:00					
VSI1	VSIPLEX	26.9	172.5	92.7	24.9
V24A	VSIPLEX	33.8	122.6	85.2	28.8
16:52:00 - 16:53:00					
VSI1	VSIPLEX	26.3	173.5	92.6	24.4
V24A	VSIPLEX	29.9	129.4	88.2	26.4
16:53:00 - 16:54:00					
VSI1	VSIPLEX	25.0	174.8	92.0	23.0
V24A	VSIPLEX	34.6	124.7	85.1	29.4
16:54:00 - 16:55:00					
VSI1	VSIPLEX	24.7	175.1	92.2	22.8
V24A	VSIPLEX	37.9	118.1	88.7	33.6
16:55:00 - 16:56:00					
VSI1	VSIPLEX	26.4	173.3	92.4	24.4
V24A	VSIPLEX	45.8	92.0	86.4	39.6

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.5-1%	.5-1%	98%+	0%	100%
z/VM	1-2%	1-2%	95%+	0%	99%
z/VSE	<.1%	6-8%	92-94%	0%	99.9%
Linux	2%	7-10%	90%	< 1%	99% +
z/OS	????	13%	85-92%	8-15%?	85-92%
CICS					
DB2					

VELOCITY SOFTWARE

SMF 110 subtype 1 (cics transactions)

- CpuTime = USRCPUT (tcb) + RLSCPUT(srb)
- Add up transactions on the fly for one minute interval

Time/ SYSID APPLID	Transaction ID	Count	<-Response Time-> Total Resp	Susp Time	Disp Time	<Disp CPU Time
VSI1						
CICSZA2	Total	1192	0.015	0.005	0.010	0.010
	CSSY	6	0.203	0.203	0.000	0.000
	STRH	149	0.005	0.001	0.004	0.004
	STR1	126	0.003	0.000	0.003	0.002
	STR2	109	0.005	0.001	0.004	0.004
	STR3	90	0.006	0.000	0.006	0.005
	STR4	81	0.010	0.000	0.009	0.009
	STR5	109	0.014	0.003	0.011	0.011
	STR6	117	0.015	0.002	0.013	0.013
	STR7	125	0.018	0.003	0.015	0.015
	STR8	148	0.023	0.006	0.016	0.016
	STR9	132	0.035	0.016	0.019	0.018

VELOCITY SOFTWARE

SMF 30 (CICS job / region cpu) – 1 minute sample

SYSID	<------JOB----------> Name	JobID	Step Nbr	<CPU Percents--- Total	STD	SRB
VSI1	Totals		.	24.88	21.5	3.2
	CICSJZ2	STC08736	1	20.20	19.7	0.5
V24A	Totals		.	28.83	25.4	3.2
	C24ASTND	STC01838	1	20.22	19.6	0.6

SMF 110 (CICS capture ratio: 19.5 / 20.2 = 97%?

Time/ SYSID/ Time	APPLID	<Transactions> Group	Count	<-Total-> <--CPU--> Secs	Pct	<-Response Time-> Total Resp	Susp Time	Disp Time
VSI1	CICSZA2	Totals	1192	11.7	19.5	0.015	0.005	0.010
V24A	C24ASTND	Totals	1156	11.6	19.3	0.037	0.021	0.016

VELOCITY SOFTWARE

z/OS Capture Ratio model

At each layer,
something is lost

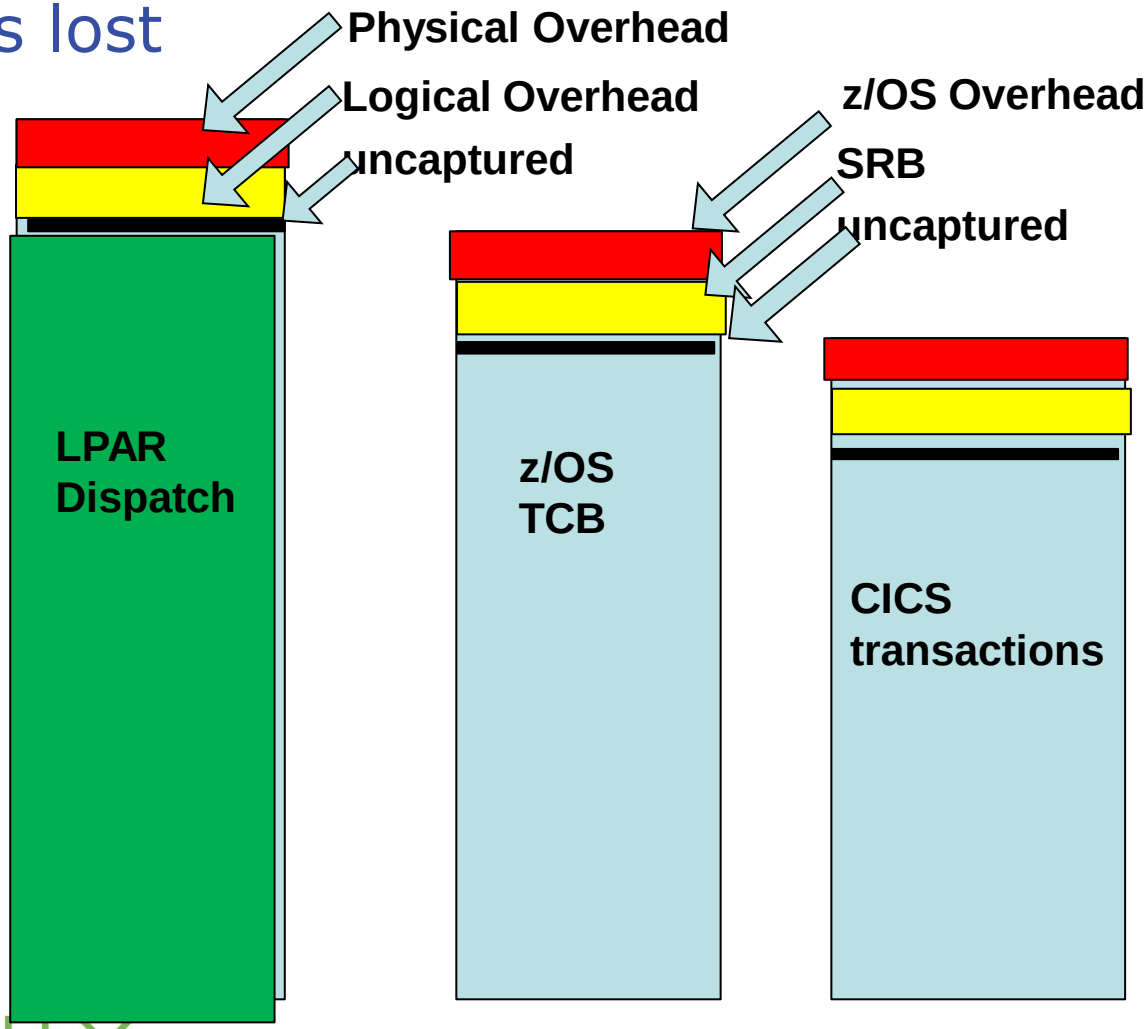

VELOCITY SOFTWARE

CICS Capture Ratio model

Question: how much CPU left for CICS transactions?

- SMF 70, LPAR Loss: 2%
- SMF 70, z/OS loss: 10%
- SMF 30 Logical Overhead/srb: 3%

What is left for transactions? 85%?

The major deficiency in this model is z/OS system time

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.5-1%	.5-1%	98%+	0%	100%
z/VM	1-2%	1-2%	95%+	0%	99%
z/VSE	<.1%	6-8%	92-94%	0%	99.9%
Linux	2%	7-10%	90%	< 1%	99% +
z/OS	????	3-13%	85-92%	8-15%?	85-92%
CICS		3%	96-97%	< 1%	99+%
DB2					

VELOCITY SOFTWARE

DB2 Capture Ratio model

SMF 30 – CPU by region

- CPU Data by region:

SMF 100 -

- QWSA – DB2 Statistics by region
- TCB,SRB,I/O CPU metrics

SMF 101

- CPU By DB2 Transaction
- QWACEJST – QWACBJST (total cpu per transaction)
  - (db2 time: qpacajst)
  - (package: qpactjst)
- Add up all transactions by transaction name, total

VELOCITY SOFTWARE

DB2 Capture Ratio model

SMF 30 – CPU by region

SYSID	<------JOB----------> Name	JobID	Step Nbr	<---CPU Percent Total	STD	SRB	
V24A	Totals		.	7.22	5.8	1.3	
	DSN1DBM1	STC04803	1	0.05	0.04	0.01	
	DSN1DIST	STC04804	1	0.02	0.02	0.00	
	DSN1MSTR	STC04801	1	0.64	0.61	0.01	
	IRLMPROC	STC04802	1	0.13	0	0.13	←SRB

SMF 100 (db2 statistics) – Match SMF 30 100%....

Date/ Time/ Sysid	SSID	DB2 Proc	<Percent> Intervl CPU
V24A	DSN1	DBM1	0.050
		DIST	0.023
		MSTR	0.642
		IRLM	0.132

VELOCITY SOFTWARE

DB2 Capture Ratio model

SMF 101 (db2 transactions - Capture low, needs work?)

<--------Transaction--------->			
Name	Count	Type	Total CPUPct
trantotal	2477	0	0.60 <--- vs .84 from smf100
POOLZDHD	753	4	0.16
POOLZDCD	747	4	0.16
POOLZDCU	741	4	0.23
POOLADCD	119	4	0.02
POOLADHD	117	4	0.02

SMF 100 (db2 statistics)

Date/ Time/ Sysid	SSID	DB2 Proc	<Percent> Intervl CPU
V24A	DSN1	DBM1	0.050
		DIST	0.023
		MSTR	0.642
		IRLM	0.132

VELOCITY SOFTWARE

Building the "Capture Model"

LPAR Layer (highest layer)

- Overhead "low"
- Capture Ratio 100%

	Mgmt	Logical	Work	Uncaptured	Capture Ratio
HMC/LPAR	.5-1%	.5-1%	98%+	0%	100%
z/VM	1-2%	1-2%	95%+	0%	99%
z/VSE	<.1%	6-8%	92-94%	0%	99.9%
Linux	2%	7-10%	90%	< 1%	99% +
z/OS	????	3-13%	85-92%	8-15%?	85-92%
CICS		3%	96-97%	< 1%	99+%
DB2		16%	70%+		70%

VELOCITY SOFTWARE

Recap and Thankyou

Capture Ratios validate the data for

- Capacity Planning – know consumption by app
- Chargeback – who consumed exactly what?
- Performance analysis – who is using cpu now?

Complementary data evaluation for anyone who sat thru 65 slides in 58 minutes

Thank you for your time!! See you in person next year

Questions and suggestions can be sent to 'barton@velocitysoftware.com'

VELOCITY SOFTWARE