Linux on z/VM
Chargeback and Accounting

Barton Robinson

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

Clifford Stoll's "The Cookoo Egg", Tracking a Spy Through the Maze of Computer Espionage

Why chargeback?

Accounting Issues

Accounting Models

Data Requirements

zVPS Accounting data

My first accounting program:

- written for IBM datacenter, 1976,
- Cobol, challenge: to correct CPU time for DOS/VS under vm/370

VELOCITY SOFTWARE

The Cuckoo's Egg

Clifford Stoll's "The Cuckoo's Egg"

1986, computer time at LBL cost $300/hour

Accounting had 75 cent discrepancy

The Process took almost a year:

- Set up logging
- Find the hacker's entry point (LBL, 1200 baud dial up modem)
- Access to a dozen networks, 1000's of computers
- Show the "defense contractors" that had been compromised

Result: Hackers (students in Germany), stealing computer time, sold information to KGB to support drugs, phone bills

When doing accounting, the numbers must add up!

VELOCITY SOFTWARE

IT as a business - Profit center vs Cost center

Cost center

- Corporate line item
- One price, percent of gross
- Easily outsourced to more efficient “profit organizations”
- Often inherent waste, decisions made based on religion
- Less regard to new initiatives (big data?)

Profit center

- “free” to corporation
- Funded by the users, they have choices
- More efficient because it has to be
- Platform decision results measurable by the users

Why chargeback?

Capital Expenditures “Large” vs “small”

- Capital Expenditures seen by executive committee
- Justification of $100K IFLs different than servers
- Distributed / Rack servers often under the “sign-off” limit

Justification for hardware

- Justify IFLs – knowing who is using them
- Knowing what application is paying for the new ones

In profit center model, IFLs are justified by user requirements and funding

VELOCITY SOFTWARE

Accounting Justification

Supporting Corporate Interests and service levels

- Understanding costs promotes competitiveness
- Understand service and service level requirements
- Encourages higher service levels at lower cost

Chargeback Encourages::

- Lower resource requirements
- Tuning applications
- Workload scheduling
- Better service levels if heavy CPU spikes moved off shift

Accurate accounting models takes religion out of platform decisions

- Bad decisions are not buried in a corporate line item

VELOCITY SOFTWARE

Accounting Justification

Charge structure resolves issues

A favorite “cron” job has a measureable cost

- Many sysprogs have their personally developed “cron job”
- It runs across many servers at “low cpu cost each”
- It may cost 3 IFLs in aggregate or:
- It may cost $300,000 per year or more to operate – is it worth it?

Putting a structured cost to operations changes perspectives

How much of the system resources are wasted because costs are unknown or not understood?

VELOCITY SOFTWARE

Cloud Challenges

Cloud:

- Cloud: Old concept, new name, same requirements
- Time Sharing existed in the 70's, same requirements
- Dynamic access to resource and service
- Often 100's or 1000's of servers (z)

Cloud as a business

- What service is being provided
- What to charge for a service?

Cloud REQUIRES accounting

VELOCITY SOFTWARE

Amazon charges

Amazon

- Charge by IP address
- Charge by bandwidth
- Storage (100tb -- $8400/month)
- DB: 1st 25 cpu hours per month free
- DB: row selects, adds, ….
- Email requests
- Data base loads
- 1TB of HPC: $28,000 / year

Everything is measured

VELOCITY SOFTWARE

Google

- **Google Cloud** Storage **pricing** is based on usage. Project storage usage and bandwidth usage are calculated in gigabytes (GB),Storage (100tb -- $8400/month)
- All machine types are charged a minimum of **10 minutes** .For example, if you run your instance for 2 minutes, you will be billed for 10 minutes of usage.
- After 10 minutes, instances are charged in **1 minute increments, rounded up to the nearest minute.**
- A 10TB persistent disk $400/month

VELOCITY SOFTWARE

VSI Colo Charges

Pricing model

- **10x10 cage**
- **30 Amp redundant circuit (15 amps used)**
- **A/C included in price of circuit**
- **20 Amp circuit**
- **16 IP addresses**
- **High bandwidth connection**
- **Max 1 gigabit / month, incremental charges after that**

Services include

- **UPS, Diesel generators (2 weeks of diesel on site)**
- **Security – badge in/badge out**
- **Internet**

VELOCITY SOFTWARE

Performance Management Requirements

Chargeback implies service levels as well:

Capacity Planning

- Methodology needed to project future requirements

Performance Analysis

- Tools/education needed to resolve current service issues

Operations

- Alerts for service issues
- If a process loops, who is responsible for detection and correction?

VELOCITY SOFTWARE

Define amount of money to be recovered

Charges for “z” must be “reasonable” as compared to “x”

Cost allocations must be correct

- Allocating data center costs these days to “Z” does not make sense
- Run as a business, departments don’t pay to support other departments

Honest allocations

- Electricity, A/C
- Floor space

System programmer who says: “I will quit if I have to support chargeback”

Challenges

Every data center has mixed technologies, other virtualization shared resource should be considered

If adding one more distributed server requires a new A/C Chiller, and one more chiller on the roof requires a new building, who pays the true costs? (True example, converted to “z” and saved a building)

If new data center is required to support more racks (or networking, or…) who pays?

- Incremental costs of the racks can be buried

VELOCITY SOFTWARE

Defining Costs

Hardware costs

- Processor, storage, network

Software “shared” licenses, by platform

- z/VM and Linux system software (and support)
- Linux application software
- Infrastructure (management, backup)
- Don’t charge HP Openview to the mainframe

Personnel Costs, by platform:

- Systems support (including office expenses)
- Operations

Network – based on bandwidth, user requirements

- Connection costs
- Personnel

VELOCITY SOFTWARE

Defining Costs

Environmentals, (BY PLATFORM)

- Floor space
- Power/Air Conditioning
- Power requirements should be documented for ALL hardware

Disaster Recovery (or lack of)

- Don’t pay for it if it won’t work
- Major strength of “z”, but has associated costs

Data backup

VELOCITY SOFTWARE

Accounting Model Considerations

Service level definitions

Design to meet Corporate Interests

- Move workloads to low resource periods when possible
- Minimize resource spikes during prime shift
- Convince users to tune their applications

Accounting model considerations

- Different support models (critical, production, development)
- Don’t charge for off shift resource consumption
- Charge for resources relevant to dollar purchases

Off shift work – when?

- Backups
- Database loads
- Server creates
- Batch updates

VELOCITY SOFTWARE

Influencing the Users

Resources that are the most expensive and most limiting:

- Processor / IFLs
- Real Storage

Chargeback should convince users to

- Move work off shift
- Tune applications
- Use real storage realistically (is 20 GB SGA really needed?)

- Utilize the I/O subsystem instead of caching the world

VELOCITY SOFTWARE

Planning for Chargeback

Accounting Methodology:

- Charge for resource consumption
- Flat rate
- Combination

Acknowledge resources have both fixed and variable costs

- Disk space is fixed or variable?
- CPU time is variable
- Network connections are fixed
- Network use is variable

VELOCITY SOFTWARE

Listserv comment, Feb 23, 2015

Hello! There are going to be a lot of ideas on how to do charging all the way from "we don't charge back" to "we charge by the micro-process".

Fixed-price charging is good for very well contained zLinux servers. Using SHARE ABSOLUTE and even CPU POOLING can help with this. CPU POOLING is really good if you are thinking of mixing those fixed-resource servers with high-performance or "premium services" servers.

VELOCITY SOFTWARE

Listserv comment, Feb 23, 2015

There are pros and cons to fixed price charging. A pro is it is MUCH easier than any other kind of charging (other than no-cost!).

A con is that some folks may think they are being charged too much if they really don't use it that much or that heavy, and some may not be thrilled being capped when they want to run something heavy for a short while.

VELOCITY SOFTWARE

Listserv comment, Feb 23, 2015

We use a variable rate based on consumption for most of our servers. We also have "lab" servers that are contained within SHARE ABS. We have processes to "cap" non-premium production servers when they get out of hand and the entire processor has been running in the high 90's for CPU.

We use zVPS to collect all the data for the processor/LPARs and the zLinux servers to evaluate what they are up to every minute. Using that data, we are able to charge-back to the business units based on what they consume by process/application.

VELOCITY SOFTWARE

Listserv comment, Feb 23, 2015

There is a small base charge for just having a server, then the use of resources adds to it. This was all built over the last 10 years and we are still tuning it! Figuring out what your business will accept for charging is the hard part. Between the knobs in z/VM and performance monitoring/collecting tools available, you can make it work the way you need it. -- ***James Vincent*** -- President, SHARE Inc

VELOCITY SOFTWARE

Planning for Chargeback

Computing Resources – traditional model

- by CPU second
- by resident storage (storage used)
- DASD I/O, disk space

Current model

- By CPU second
- By virtual machine size
- Disk space
- Network bandwidth required
- Network bandwidth utilized

VELOCITY SOFTWARE

Accounting Example

- Support types:
  - Premium: 7x24 technical support
  - Prime: Prime shift technical support
  - General: Technical support as available
- Charge for CPU consumption 7AM to 5PM
  - (No charge off shift)
- Charge for Average Storage residency
  - (Encourages smaller virtual machines)
- Charge for DISK space allocated

Model proven to minimize IFL and storage requirements

VELOCITY SOFTWARE

Accounting Model Failures

Accounting Model tried (large manufacturing company)

- Monthly flat fee per server

Flat fee encourages:

- Larger servers doing more work
- Wasting of resources

The need for additional IFLs with no additional revenue broke the model

- Fee charged needs to be the maximum to survive

VELOCITY SOFTWARE

Accounting Model Failures

Model based on Linux CPU data

- Prior to sles10, CPU data “VERY” bogus
- One test case showed Linux CPU data wrong by 100

- z/VM data is accurate to the microsecond -

Solution - prorate

VELOCITY SOFTWARE

Data Challenges – data objectives

- Data collection must not be prohibitive
- Spikes must be validated with historical data
- Capture Ratios (data model) must be validated
- Accurate accounting data source

Accounting Data Requirements

Resource utilization data requirement

- CPU by z/VM Virtual Machine (Linux server)
- CPU by Linux process
- CPU by application
- Real Storage residency

Historical data required

- Customer asks “why did I spike on October 20”?
- What processes, application was running?

Data Accuracy

- Data is accurate (Linux CPU measurements are not)
- Data must be validated
- Data collection must not be prohibitive

Capture Ratio (100%????)

- How much of resource consumed is measured?
- Validates the model

VELOCITY SOFTWARE

CPU Accounting

CPU consumption is key charge item

- 10-20GB ram required per IFL
- Software licenses
- IFL initial fee and maintenance

What CPU is charged?

- Infrastructure / support? (Overhead)
- System Overhead (by virtual machine)
- LPAR / CP Overhead?

Prorate function still needed?

- z/VM data very accurate
- Pre SLES10, Linux data under “VM” was bogus
- Linux data mostly valid (post sles10)
- Linux data does not include overhead
- Prorate used to include virtual machine system overhead

VELOCITY SOFTWARE

z/VM Data collection

Traditional data model (1989)

ZMON: Real time analysis

- **Uses Standard CP Monitor**

Real Time Analysis

ZMAP: Performance Reporting

Post Processing

Creates Long Term PDB

PDB or monwrite data input

PDB (Performance DataBase)

Complete data

By Minute, hour, day

Monthly/Yearly Archive

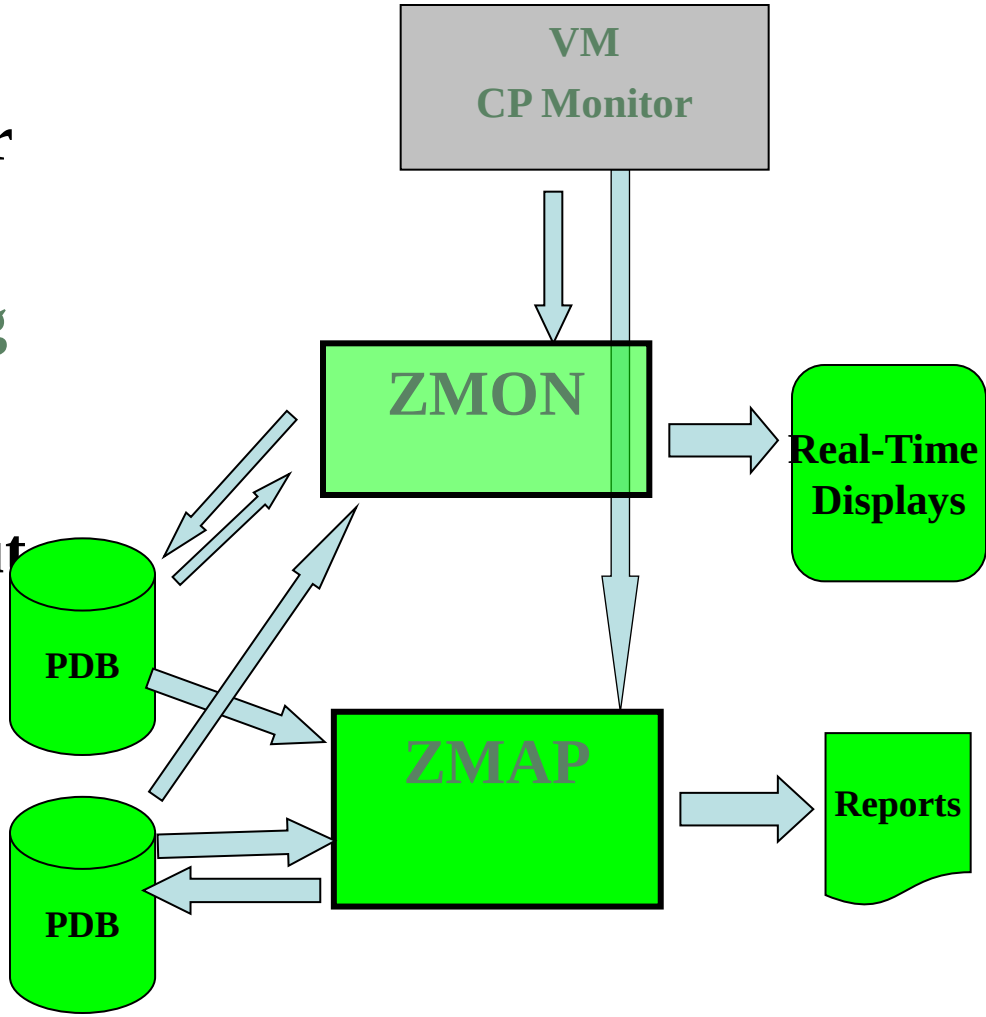

VELOCITY SOFTWARE

Linux and Network Data Acquisition

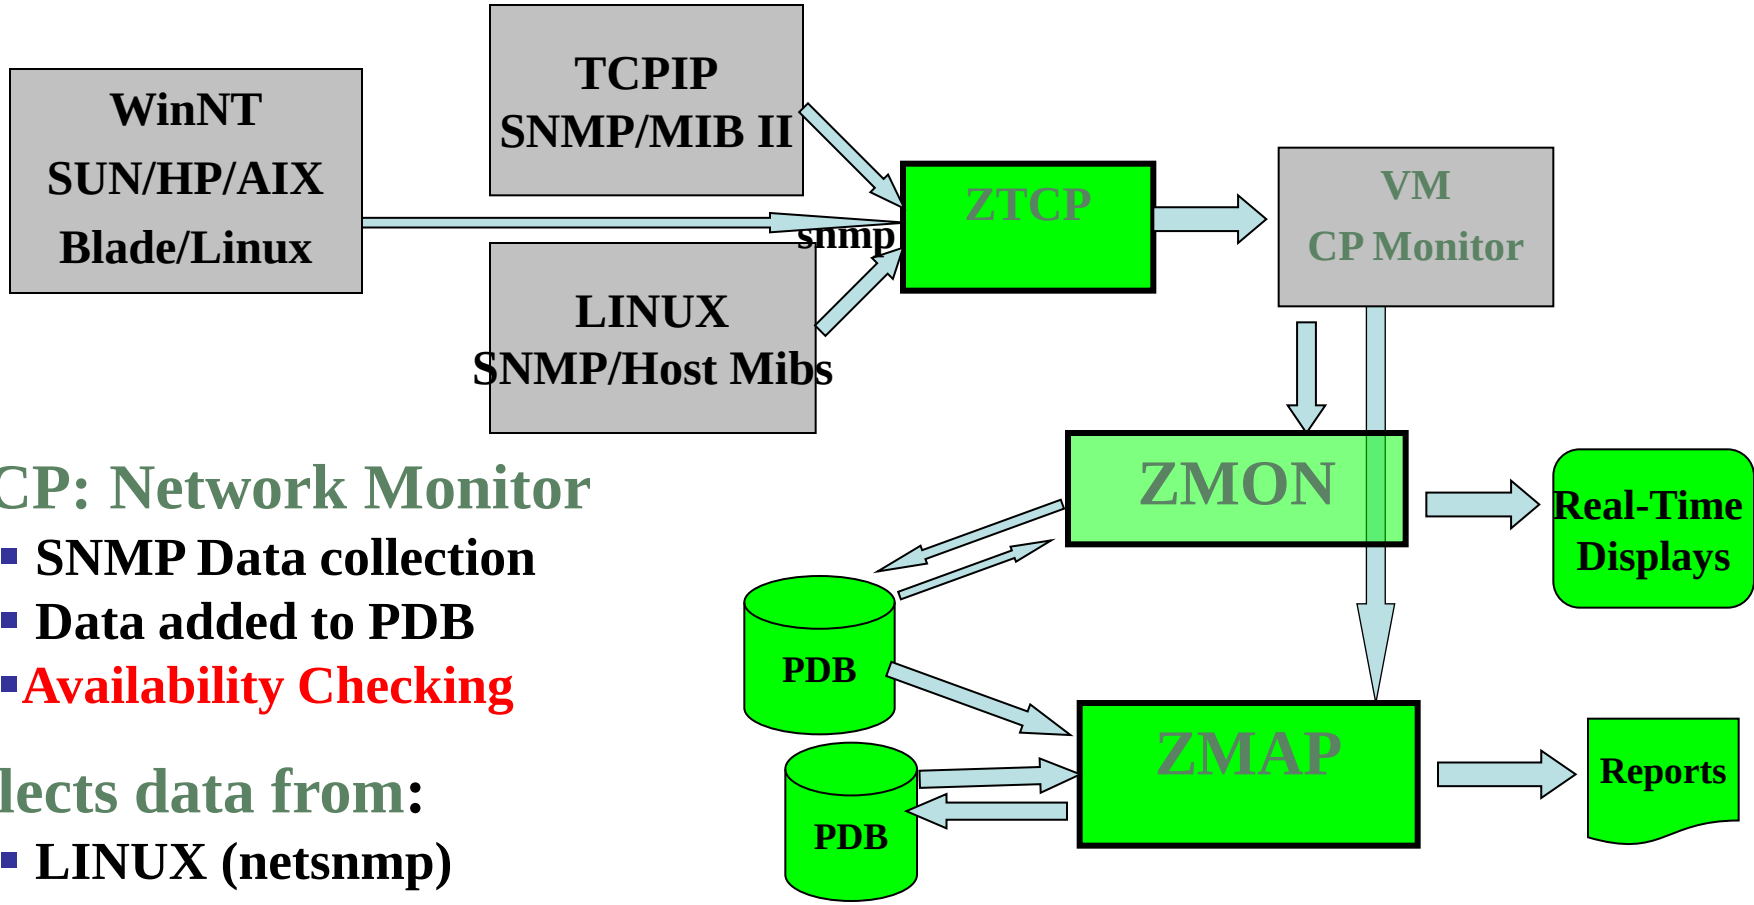

ZTCP: Network Monitor

- SNMP Data collection
- Data added to PDB
- Availability Checking

Collects data from:

- LINUX (netsnmp)
- NT/SUN/HP (native snmp)
- Printers/Routers….

VELOCITY SOFTWARE

Add Enterprise Support

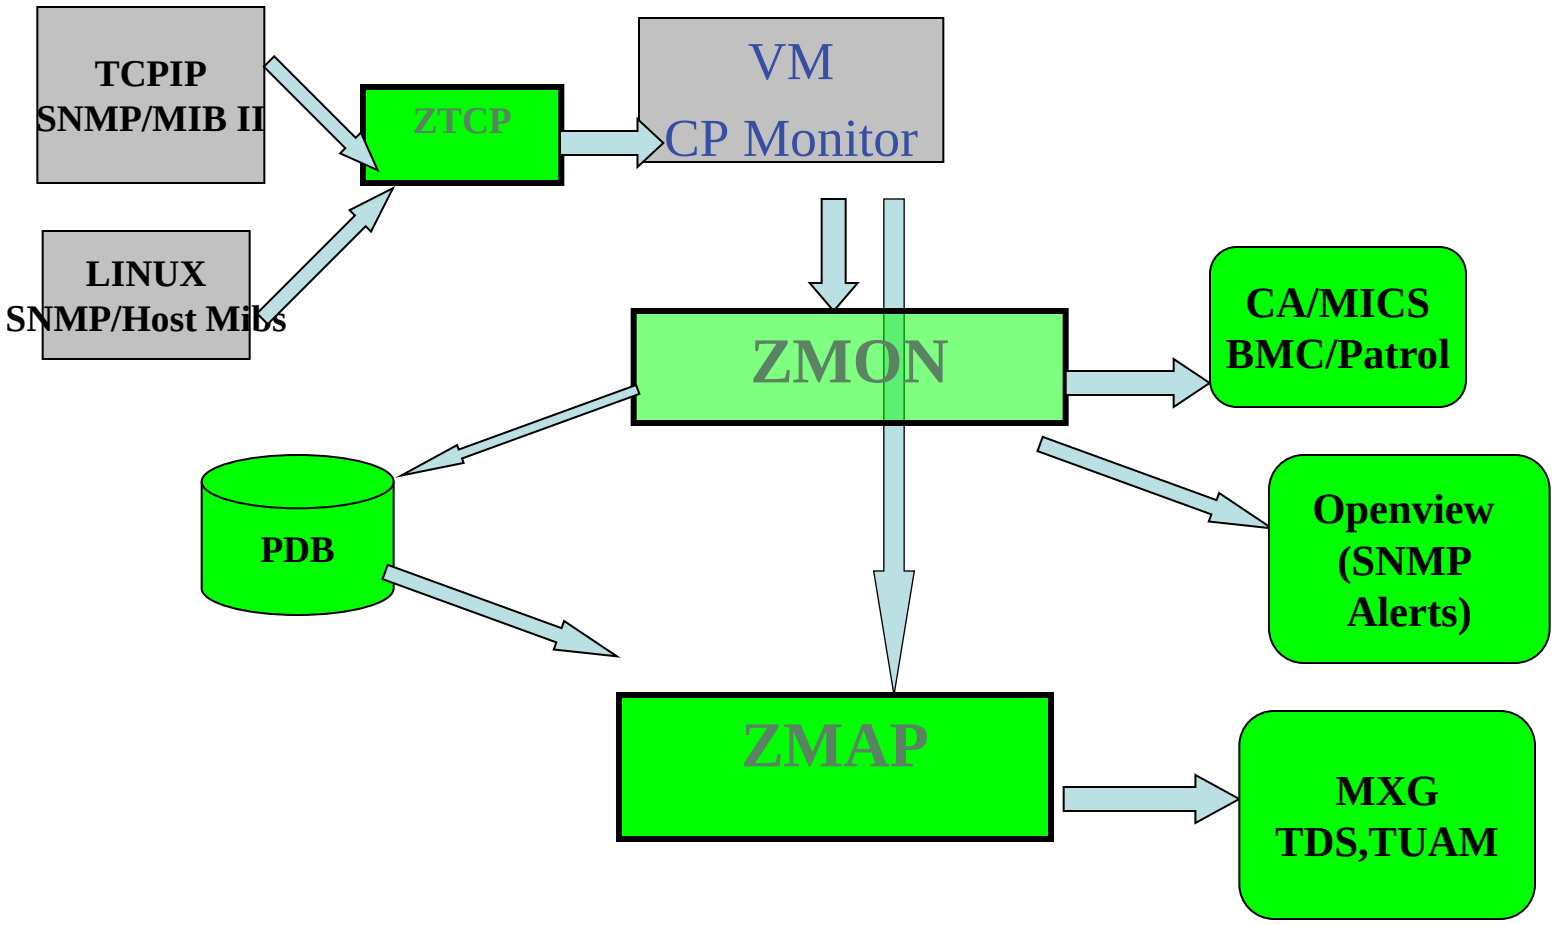

z/VM Accounting with ZVPS

Accounting using zVPS:

- CPU data from z/VM very accurate
- Linux process data accurate
- Linux application data accurate
- Historical data maintained
- Capture ratios 100%

Traditional z/VM accounting - ESAACCT

- by CPU second,
- by resident storage
- DASD I/O
- Reported on ESAACCT (accounting report)
- Extracted from performance database
- Users can see data real time

VELOCITY SOFTWARE

Traditional z/VM Accounting

Accounting Definitions:

- Service unit: Chargeable unit (compare z9 to z10)
- CPU Service Units: Service Units per CPU Second
- Storage Service Units: Service Units per page resident
- DASD I/O: Service Units per Virtual I/O

Real Storage Accounting (example)

- Storage costs $500 / GB per month
- Prime shift (20 days per month), storage costs $25 / GB
- Server averages 2GB resident per day, should be charged $50 day
- Choose STORAGE FACTOR

DASD I/O – More arbitrary

- Channels, Storage processors may be factors

VELOCITY SOFTWARE

Traditional z/VM Accounting

```
Report: ESAACCT     User Accounting Report                          Veloc
Monitor initialized: 10/19/08 at 14:00:00 on 2097 serial 274E2      First
-------------------------------------------------------------------------
Service Unit Virtual CPU Factor:    32989
Service Unit Overhead CPU Factor:   32989
Service Unit I/O Factor:            0.800
Service Unit Storage Factor:        0.500
Charge per Service Units:            0.00001
-------------------------------------------------------------------------
UserID   <--Users-->  <Service Units>  Charges  <---Resources Used---->
/Class   Logged Actv    /Sec   Total    Total     CPU   DASD I/O  Pages
-------- ------ ----  ------  ------   -------  ------  --------  -----
14:15:00     61 40.4   67058  60353K    603.52  1551.1    337088    18M
SYSTEM        .    .  715.67  644103      6.44   19.52
 ***Key User Analysis ***
RSCS          1  1.0    0.67  601.32      0.06    0.02         0     24
TCPIP         1  1.0   21.80   19619      1.96    0.58         0    810
VMSECURE      1  1.0    1.71  1535.0      0.15    0.04        26     65
 ***User Class Analysis***
*Servers     29 12.4  156.69  141023     14.10    3.91     12819   3719
*Linux       23 23.0   66161  59545K    595.45  1527.0    324243    18M
*Misc         4  2.0    0.50  448.13      0.04    0.00         0    667
 ***Top User Analysis***
LXSG0012      1  1.0   33662  30295K    302.94  903.65    228801   604K
LXSG0020      1  1.0  6365.4   5729K     57.28  146.40      4354  1792K
LXSG0011      1  1.0  5661.6   5095K     50.95  119.10      4641  2325K
LXSG0002      1  1.0  3778.2   3400K     34.00   72.64     30204  1960K
LXSG0003      1  1.0  3111.4   2800K     28.00   72.40      3625   818K
LXSG0017      1  1.0  1325.5   1193K     11.92   29.79      2039   417K
LXSG0016      1  1.0  1319.4   1187K     11.87   24.08      3079   781K
LXSG0008      1  0.9  1104.9  994451      9.94   15.27      1790   978K
LXSG0023      1  0.5  569.82  512841      5.12    9.26      1160   413K
LXSG0015      0  0.5  522.87  470584      4.76    8.75      1347   362K
```

VELOCITY SOFTWARE

Linux Accounting

Options:

- Charge per server
- Charge by application

Must have ability to show cost of “cron”

Requirements for application, infrastructure chargeback

- Must have full process table information every interval
- Must have process parent / child relationships
- Must have capture ratio exposed

zVPS Collects full process table every 60 seconds

- Data includes parent child relationship
- When process terminates, CPU is moved to parent “children buckets”
- An application is defined as either by process name (Oracle), or by parent
- Websphere identifies the “head of application”, that owns the worker processes

VELOCITY SOFTWARE

Linux Process Capture Ratio

High Linux CPU capture ratio

Report: ESALNXV LINUX Virtual Processor Analysis Report

Node/ Name	VM ServerID	<Linux Pct CPU> Total	Syst	User	<Process Data> Total	Syst	User	Capture Ratio	Prorate Factor
10:03:00									
NEALE1	LNEALE1	100.0	11.4	88.6	100.2	11.5	88.7	**1.002**	1.000

Report: ESALNXP LINUX HOST Process Statistics Report

node/ Name	<-Process Ident-> ID	PPID	GRP	Nice Valu	<------CPU Percents----> Tot	sys	user	syst	usrt
10:03:00									
NEALE1	0	0	0	0	100	0.43	3.35	11.0	85.4
kswapd0	100	1	1	0	0.12	0.12	0	0	0
snmpd	1013	1	1012	-10	**0.13**	0.03	0.10	0	0
sh	3653	3652	30124	0	52.7	0	0	9.37	43.3
gmake	9751	9750	30124	0	43.4	0.02	0.02	1.37	42.0
sh	10129	9751	30124	0	0.02	0.02	0	0	0
sh	10130	10129	30124	0	0.63	0.03	0.23	0.28	0.08
cc1	10307	10306	30124	0	3.12	0.18	2.93	0	0
rpmbuild	30124	16382	30124	0	0.07	0.03	0.03	0	0
sh	30125	30124	30124	0	0.02	0	0.02	0	0
gmake	30126	30125	30124	0	0.02	0	0.02	0	0

Report: ESALNXC LINUX Process Conf

Node/ Name	<-Process Ident-> ID	PPID	GRP	<-----Pr Path
NEALE1				
init	1	0	0	init [3]
migratio	2	1	0	migratio
ksoftirq	3	1	0	ksoftirq
events/0	4	1	0	events/0
khelper	5	4	0	khelper
kblockd/	6	4	0	kblockd/
cio	41	4	0	cio
cio_noti	42	4	0	cio_noti
kslowcrw	43	4	0	kslowcrw
appldata	96	4	0	appldata
aio/0	101	4	0	aio/0
pdflush	5266	4	0	pdflush
pdflush	26647	4	0	pdflush
kswapd0	100	1	1	kswapd0
kmcheck	158	1	1	kmcheck
syslogd	976	1	976	/sbin/sy
klogd	979	1	979	/sbin/kl
snmpd	1013	1	1012	snmpd
portmap	1030	1	1030	/sbin/po
rpciod	1034	1	1	rpciod
lockd	1035	1	1	lockd
sshd	1072	1	1072	/usr/sbi
sshd	16272	1072	16272	sshd: bu
sshd	16288	1072	16288	sshd: bu
sshd	16290	16288	16288	sshd: bu
bash	16291	16290	16291	bash
python	16312	16291	16291	python
do-bui	16313	16312	16291	/bin/sh
bb_do	16382	16313	16291	/usr/bin
rpmb	16415	16382	16415	rpmbuild
rpmb	30124	16382	30124	rpmbuild

VELOCITY SOFTWARE

Storage map - **CAPTURE RATIOs always critical for any instrumentation:**

- CP Fixed Storage
- CP Non Pageable
  - Free storage (only VMDBLKs)
  - Frame tables
- Dynamic Paging Area(DPA)
  - System Execution Space
  - User storage, MDC, Address Space, Vdisk **– This validates user resident storage**
  - Available List (greater/less than 2gb)

```
Report: ESASTR1      Main Storage Analysis              Velocity Software, Inc.       ESAMAP 3.6.0 05/15/06  Page  57
Monitor initialized: 06/06/05 at 08:42:16 on 2064 serial 11542     First record analyzed: 06/06/05 08:42:42
-------------------------------------------------------------------------------------------------------------------------
         Users <----------------------------Pages---------------------------------------------------------->
         Loggd System  Fixed Non-  Free Frame <Available> Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag    Capt-
Time        On Storage Store Pgble Stor Table <2gb  >2gb  ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98    Ratio
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- ----- ----  -----
08:45:42    22 7864304  2907  3816    5 61440  513K 6292K 33150 778370     8408   1090  6235     0  133K  333  0.995
*********************************************Summary*****************************************************
Average:    22 7864304  2907  3816    5 61440  513K 6292K 33150 778370     8408   1090  6235     0  133K  333  0.995
```

Database (Performance DataBase/PDB) Extraction

- Define headings as you like
- Define variables, summary data only

```
EXTRACT:            ; First extract system data

columns ='10 10 7 8 8 8'
TITLE = ' USERID    ACCTNBR      CPU   Pages     VIO'
TITLE = ' --------- -------- -------  ------  ------'

X = 'USERID'
Y = 'USRCON.VMDACTNO'
Y = 'USEACT.VMDTTIME' ; cpu utilization
Y = 'USEACT.VMDCTPVR / 256' ; MB Resident
Y = 'USEACT.VMDVDSCT' ; VIO

INTERVAL='SU'
CRITERIA='USRTYPE = USER'
CRITERIA='USRCON.CLASSID = Linux'
CRITERIA='STARTTIME >= 080000'
CRITERIA='STOPTIME <= 170000'
```

Database extract, traditional metrics

Date	Time	USERID	ACCTNBR	CPU	Pages	VIO
20100603	080000	REDHAT04	REDHAT	4	99	742
20100603	080000	REDHAT3	REDHAT	8	170	790
20100603	080000	REDHAT5	REDHAT	6	83	3138
20100603	080000	REDHAT5X	REDHAT	11	217	1
20100603	080000	SUSELNX1	SUSELNX	18	46	6985
20100603	080000	SUSELNX2	SUSELNX	194	15	1067
20100603	090000	REDHAT04	REDHAT	4	99	731
20100603	090000	REDHAT3	REDHAT	8	171	797
20100603	090000	REDHAT5	REDHAT	6	83	3211
20100603	090000	REDHAT5X	REDHAT	11	218	6
20100603	090000	SUSELNX1	SUSELNX	18	46	6977
20100603	090000	SUSELNX2	SUSELNX	195	17	1054
20100603	100000	REDHAT04	REDHAT	15	103	4618
20100603	100000	REDHAT3	REDHAT	8	173	792
20100603	100000	REDHAT5	REDHAT	6	84	3172
20100603	100000	REDHAT5X	REDHAT	11	220	0
20100603	100000	SUSELNX1	SUSELNX	33	46	7190
20100603	100000	SUSELNX2	SUSELNX	205	20	1288

VELOCITY SOFTWARE

Server Accounting with ZVPS

Linux accounting by “velocity mib”

- CPU by process, capture ratio 100%
  - Requires “prorate” technology
- CPU by Linux Application
- Reported on ESALNXA
- Extracted from performance database

Other virtual platforms need accounting too!

Microsoft, SUN, P-series Accounting by “host” process

- By CPU second, capture ratio 100% for static processes
- Capture ratio for dynamic processes low without Velocity mib
- Reported on ESAHST1
- Extracted from performance database

VELOCITY SOFTWARE

Linux accounting by “application”

```
EXTRACT:              ; First extract system data

columns ='10 10 10 8 8 8 8'
TITLE = ' STARTTIME  STOPTIME NODE      APPL        PID'
TITLEc= '     CPU'
TITLE = '---------- --------- --------- -------  ------'
TITLEc= ' -------'

INTERVAL='SU'
CRITERIA='USRTYPE=USER'
CRITERIA='USRCON.CLASSID=*Linux'
CRITERIA='STARTTIME >= 080000'
CRITERIA='STOPTIME <= 170000'
CRITERIA='VSINAP.TOTCPU > 0 '

X = 'STARTTIME'
Y = 'STOPTIME'
Y ='NODE'
Y ='VSINAP.NAME‘
Y ='VSINAP.ID  '
Y ='VSINAP.TOTCPU / 100'  ;Turn into CPU Seconds
```

By server, by application, CPU seconds consumed

Date	Time	NODE	APPL	PID	CPU
20090910	092200	LNXAS009	*Totals	0	34
20090910	092200	LNXAS009	agc2	29188	8
20090910	092200	LNXAS009	nodeage	318	7
20090910	092200	LNXAS009	automou	3028	0
20090910	092200	LNXAS009	cron	3313	1
20090910	092200	LNXAS009	kernel	1	10
20090910	092200	LNXAS009	multipa	24607	4
20090910	092200	LNXAS009	sendmai	3154	0
20090910	092200	LNXAS009	snmpd	30961	3
20090910	092200	LNXAS029	*Totals	0	170
20090910	092200	LNXAS029	.opende	3340	0
20090910	092200	LNXAS029	ace	28947	8
20090910	092200	LNXAS029	agenten	18597	7
20090910	092200	LNXAS029	auditau	26553	13
20090910	092200	LNXAS029	a2adata	25244	19
20090910	092200	LNXAS029	comm_um	15950	11
20090910	092200	LNXAS029	c2cb_mi	20312	3
20090910	092200	LNXAS029	ezacqui	4612	9
20090910	092200	LNXAS029	isodata	28948	7
20090910	092200	LNXAS029	nodeage	9756	15
20090910	092200	LNXAS029	webshee	22130	15

VELOCITY SOFTWARE

Websphere argument string used for accounting

wasadmin 27144 6846 0 Feb06 ? 00:43:13 /u01/was61/java/bin/java -Declipse.security -Dwas.status.socket=34229 -Dosgi.install.area=/u01/was61 -Dosgi.configuration.area=/u01/was61/profiles/appsrv/configuration -Dosgi.framework.extensions=com.ibm.cds -Xshareclasses:name=webspherev61_%g,groupAccess,nonFatal -Xscmx50M -Xbootclasspath/p:/u01/was61/java/jre/lib/ext/ibmorb.jar:/u01/was61/java/jre/lib/ext/ibmext.jar:/u01/J2EEProbe/DiagnosticsAgent/classes/IBM/1.5.0/instr.jre:/u01/J2EEProbe/DiagnosticsAgent/classes/boot -classpath /u01/was61/profiles/appsrv/properties:/u01/was61/properties:/u01/was61/lib/startup.jar:/u01/was61/lib/bootstrap.jar:/u01/was61/lib/j2ee.jar:/u01/was61/lib/lmproxy.jar:/u01/was61/lib/urlprotocols.jar:/u01/was61/deploytool/itp/batchboot.jar:/u01/was61/deploytool/itp/batch2.jar:/u01/was61/java/lib/tools.jar -Dibm.websphere.internalClassAccessMode=allow -verbose:gc -Xms1024m -Xmx1200m -Dws.ext.dirs=/u01/was61/java/lib:/u01/was61/profiles/appsrv/classes:/u01/was61/classes:/u01/was61/lib:/u01/was61/installedChannels:/u01/was61/lib/ext:/u01/was61/web/help:/u01/was61/deploytool/itp/plugins/com.ibm.etools.ejbdeploy/runtime -Dderby.system.home=/u01/was61/derby -Dcom.ibm.itp.location=/u01/was61/bin -Djava.util.logging.configureByServer=true -Duser.install.root=/u01/was61/profiles/appsrv -

Linux accounting by DISK Size, Space

```
EXTRACT:              ; First extract system data

columns ='10 8 8 10 10'
TITLE = ' NODE       disk     size(MB)   Used(MB)'
TITLE = ' ---------  -------  ---------  --------'

INTERVAL='SU'
CRITERIA='STARTTIME >= 080000'
CRITERIA='STOPTIME <= 170000'
CRITERIA='HSTMEM.ALLOCUN > 1024'
CRITERIA='HSTMEM.SIZE > 0'

X ='NODE'
Y ='HSTMEM.DESCR   '
Y ='HSTMEM.SIZE / 1024' ;Convert to MB
Y ='HSTMEM.USED / 1024' ;Convert to MB
```

Linux disk accounting metrics

Date	Time	NODE	disk	size(MB)	Used(MB)
20090910	092200	LNXAS009	/	272	106
20090910	092200	LNXAS009	/usr	1203	583
20090910	092200	LNXAS009	/opt	157	2
20090910	092200	LNXAS009	/webdat	4032	3131
20090910	092200	LNXAS009	/var	496	103
20090910	092200	LNXAS009	/home	248	67
20090910	092200	LNXAS009	/usr/op	248	71
20090910	092200	LNXAS009	/usr/X1	248	71
20090910	092200	LNXAS009	/opt/in	4032	3131
20090910	092200	LNXAS009	/local	248	71
20090910	092200	LNXAS009	/boot	51	18
20090910	092200	LNXAS009	/u01	12095	2467

Linux accounting by network activity by server

```
EXTRACT:              ; First extract system data

columns ='10 8 8 10 10'
TITLE = ' NODE      Descr    OutKB/Sec  InKb/Sec  Time(Secs)'
TITLE = ' --------- -------  ---------  --------  ----------'

INTERVAL='SU'
CRITERIA='STARTTIME >= 080000'
CRITERIA='STOPTIME <= 170000'
CRITERIA='IFTAB.type = 6'
CRITERIA='IFTAB.INOCTETS > 1000'

X ='NODE'
Y ='IFTAB.DESCR'
Y ='IFTAB.OUTOCTETS / 1024'
Y ='IFTAB.INOCTETS  / 1024'
Y ='IFTAB.SECONDS'
```

Linux Network Accounting

Date	Time	NODE	Descr	OutKB/Sec	InKb/Sec	Time(Secs)
20090910	092200	TCPIP	ETHERNE	30	82	3480
20090910	092200	LNXAS029	eth1	16	4	3480
20090910	092200	LNXASF25	eth1	16	5	3480
20090910	092200	LNXASA09	eth1	5	1	3480
20090910	092200	LNXAS019	eth1	23	3	3480
20090910	092200	LNXAS037	eth1	6	4	3480
20090910	092200	LNXAS021	eth1	7	2	3481
20090910	092200	LNXAS013	eth1	41	11	3481
20090910	092200	LNXAS031	eth1	25	18	3481
20090910	092200	LNXAS025	eth1	8	2	3480
20090910	092200	LNXAS002	eth1	13	2	3481
20090910	092200	LNXAS005	eth1	18	3	3480
20090910	092200	LNXAS011	eth1	8	2	3481

VELOCITY SOFTWARE

Accounting Summary with ZVPS

Many options, must focus on business drivers

- Determine costs
- Data provided by zVPS for resources utilized

VM Based Accounting

- CPU by virtual machine/Server, capture ratio 100%
- Storage residency
- Virtual I/O

Linux based accounting

- CPU by process, application, user
- DISK size and use by server
- Network by server

Distributed server data also collected for Accounting

VELOCITY SOFTWARE