Introduction to Performance Management

- Barton@VelocitySoftware.com
- HTTP://VelocitySoftware.com

"If you can't Measure it,

I am Just Not Interested ™"

Topics

Performance Management Infrastructure

- Performance Analysis
- Operational Alerts
- Capacity Planning
- Accounting/Charge back
- Data requirements

CP Monitor – where the data comes from

Importance of technology

- z/VM technology
- Linux Agent technology
- VSE

Infrastructure Requirements: Performance Analysis

Why Performance Analysis: Service Level Mgmt

- Diagnose problems real time
- Manage Shared resource environment
- Any application may impact other applications

Infrastructure Requirements:

- Analyze all z/VM Subsystems in detail, real time
- (DASD, Cache, Storage, Paging, Processor, Network)
- Analyze Linux
- (Applications, Processes, Processor, Storage, Swap)
- Historical view of same data important
- Why are things worse today than yesterday?
- Did adding new workload affect overall throughput?

Infrastructure Requirements: Capacity Planning

Why Capacity Planning: Future Service Levels

- How many more servers can you support with existing z15?
- What are the capacity requirements for an application?
- Avoid crises ***in advance***
- Consolidation Planning – Projecting requirements of the next 100 or 1000 servers

Infrastructure Requirements:

- Performance database (long term)
- z/VM **AND** Linux data
- Resource requirements by Server, Application, User
- z/VM and z/Linux data must be usable by existing planners
- **Interface to MICS, MXG, TUAM, TDS**

Infrastructure Requirements: Accounting and Chargeback

Why Chargeback?

- Distributed chargeback model is by server
- Shared chargeback model is by resource utilized
- Convincing customers to move applications to “z”
- Encourages efficient/effective resource use
- Carbon footprint by server becoming important
- Align IT to your business model

Infrastructure Requirements:

- Identify Resource by server
- Identify Resource by Linux application
- **High capture ratio**
- Every site does it differently, so flexible data is key

Infrastructure Requirements: Operational Alerts

Operational Requirements:

- Operations will manage 100's (1000's) of servers
- Requires active performance management
- Alerts for processes in loops, disks 90% full, missing processes
- One test server in a loop impacts all other servers
- Requires active performance management

Infrastructure Requirements:

- Fast problem detection
- Interface to SNMP management console (NETCOOL, HPOpenView)
- User tailored alerts
- Web based alerts

Data Requirement Summary

Performance data requirements:

- **Valid, correct – CPU data typically wrong or very wrong.**
- **SMT causes over reporting**
- z/VM and Linux data integrated?
- Helpful in solving problems?
- Validate benefits of tuning

Historical data requirements:

- Capacity Planning input
- Problem Analysis
- Linux
- z/VM

Accounting / Charge back:

- By server, by application, by process, by Linux userid

Manage Infrastructure cost:

- **Turning off “performance management” solves the performance problem?**

CP Monitor Overview

CP Monitor is standard API for z/VM Performance Data

- IUCV Function *MONITOR used by application to be alerted of data
- Data warranted by IBM

Data moved by CP to Monitor Data DCSS (MONDCSS)

- Specified interval (Default 1 minute)
- Specified sample rate (Default 2 seconds)
- Specified “domains”

Monitor Interval

- 60 seconds preferred for performance analysis and operational alerts
- 15 Minutes preferred for capacity planning and chargeback

Monitor DCSS (MONDCSS) areas and Data Types

- Sample data – collected at monitor interval
- Sample configuration data – collected at monitor “start”
- Event data – collected at event (user transaction, dasd seek, scheduler event)
- Event configuration – unused

Monitor Domains (11), Records (over 200)

CP Monitor Domains

11 domains (200+ records)

- 0: System (24 records)
- 1: Monitor (36 records)
- 2: Scheduler (14 records)
- 3: Storage (25 records)
- 4: User (13 records)
- 5: CPU (22 records)
- 6: Device (53 records)
- 7: Seek (1 record)
- 8: Virtual Network (4 records)
- 9: ISFC (4 records)
- 10: **Application data** - (100’s of records)
- 11: SSI (8 records)

CP Monitor DCSS

MONDCSS has 4 parts

- Sample configuration – created at monitor start (default 16mb)
- Sample data – populated at monitor interval
- Event configuration – unused
- Event data – populated at event – signaled at “block size”

“MONDCSS” Default location (7.3) 1GB
“MONDCSS” Default size (7.3) 96MB

Changes in 7.3 target incomplete data

- (zVPS increased MONDCSS size/location decades ago…)
- Issue is large DASD farms monitored exceed configuration and data areas

IBM’s “MONWRITE” reads from DCSS and produces “raw data”
zWRITE reads from DCSS and produces “history data”

CP Monitor Command

Monitor Enable all

Monitor EVENT disable scheduler ALL

- Eliminate high volume traffic with little purpose

Monitor EVENT enable scheduler user operator

- Really want the 2.8 record

Monitor EVENT disable seeks DEV rdev1-rdevn

- Useful in TPF environment

Monitor Sample config size 4096

- (Default)

Monitor Start Rate 1 sec Int 1 MIN

CP SET RESERVED DCSS MONDCSS 20000