

OpenVMS Performance Tools & Alpha Performance Update

Barry Kierstein

Software Engineer
HP OpenVMS Engineering



Overview for this update section



- Availability Manager
- DECamds
- OpenVMS Performance Data Collector
- HP OpenView

Availability Manager - Overview

- What is it?
 - Real-time performance and resource analysis tool
 - Highlights various performance problems and various resource contention problems
 - Collects data even on hung systems
 - Has ability to “fix” a number of these problems
 - Consists of two parts:
 - A GUI-based Data Analyzer
 - A OpenVMS driver-based Data Collector
 - A Java rewrite of the DECams tool
 - Ability to run the Data Analyzer on both OpenVMS and Windows systems
 - Faster program development

Availability Manager – Current status



- Current version is AM V2.3
 - DECams parity release
 - CPU Process State Summary display
 - Uses JRE 1.3.1 resulting in better performance
 - Decoding of various common lock resource names
 - Various minor feature enhancements and bug fixes
 - Data Analyzer runs on the following:
 - OpenVMS Alpha V7.2-2 and higher
 - Windows 2000 and Windows XP
 - Data Collector runs on OpenVMS V6.2 and higher, both VAX and Alpha

Availability Manager – Near future



- AM V2.3-1
 - Currently in field test
 - Support added OpenVMS Alpha V7.3-2
 - Will be released concurrent with OpenVMS Alpha V7.3-2 in the latter part of 2003

Availability Manager - Future

■ AM V2.4

- Display Extended File Cache (XFC) data
- Allow non-Administrator accounts on Windows to run the Data Analyzer
- Group and Node-level passwords
- Enable sorting of all tabular displays
- Show correct mount count for disk and nodes where the disk is mounted
- Fix various bugs when found

Availability Manager - Future

- AM V2.4+
 - Port to Itanium concurrent with OpenVMS Itanium port
 - One-screen single process view
 - Store and display site-specific information about a node
 - Better event threshold value settings
 - Data collection over IP
 - Possible handing events to OpenView for further processing
 - Various bug fixes and performance enhancements

DECams - Overview

- Predecessor to the Availability Manager
- Similar functionality to the Availability Manager
- Data Analyzer:
 - Runs on OpenVMS VAX and Alpha
 - DECWindows Motif-based GUI
- Same Data Collector as the Availability Manager

DECamds - Future

- DECamds V7.3-2
 - Currently in field test
 - Will release concurrent with OpenVMS Alpha V7.3-2
 - This will be the last release of DECamds

OpenVMS Performance Data Collector - Present



- TDC V1.5 - released in 2001 for limited distribution
 - V7.1+ for VAX & Alpha
 - 3 Formats – Binary, CSV and ASCII Text
 - Data available for clusters, CPU utilization, disk performance, system performance, system parameters & process information

OpenVMS Performance Data Collector - Future



- TDC V2 planned for next release of OpenVMS (V7.3-2)
 - Single data collector available with documented API
 - Includes network, XFC (hot files), lock statistics, CPU & RMS
 - Supports both “live” and “stored” data, can be extended by consumers and will track new OpenVMS enhancements
 - Ship concurrent with release of OpenVMS V7.3-2

HP OpenView Overview

- HP OpenView is a comprehensive and modular portfolio of software solutions for managing and optimizing business services over IT, voice and data infrastructures. It provides:
 - Better service with fewer resources
 - Maintains uptime
- Includes fault detection, performance, network, systems/servers, applications, web services and storage management across the enterprise
- Enterprise management for heterogeneous environments – not a replacement for OpenVMS system management specific products

HP OpenView Products

- Network Node Manager (NNM) provides a map of the network and a means to pinpoint problems and network bottlenecks quickly. Uses SNMP for data collection and events.
- OpenView Operations (OVO) Monitors, controls and reports on the health of the enterprise across boundaries through a "single pane of glass".
 - Automatically collects, correlates & responds to thousands of events from network devices, systems, databases and applications.
 - Native agents loaded onto systems for data collection, Event filtering & performing automatic action.
 - Smart Plug-Ins (SPIs) to manage applications and databases (Oracle, SAP, Exchange).

HP OpenView Products (cont'd)



- OpenView Storage Area Manager (OVSAM) centralizes and simplifies storage area management across distributed, multi-vendor storage, and efficiently manages availability, performance and growth.
- Storage Data Protector (OmniBack) Provides centralized and automated data protection and recovery.

Connectivity strategy for OpenView Operations



- Short term - partner with AppMind
 - AppMind™ OpenVMS System Management For HP OpenView Operations
 - Real-time monitoring of HP OpenVMS environment
 - Manage HP OpenVMS servers from an Enterprise Management System console
 - Correlate events from HP OpenVMS with other systems

Connectivity strategy for OpenView Operations(cont'd)



■ Long term solution

- Provide a “Native” agent for OpenVMS similar to agents available for HP-UX, Sun Solaris and Tru64 Unix today.
- VMS engineering is in the process of porting the existing HP-UX agent code to OpenVMS.
- Supports Alpha V7.3-1 and forward
- Functionality to include DCE based agents: control, message, logfile, action, monitor, distribution and message interceptor.

Connectivity strategy for OpenView Operations(cont'd)



- Long term solution (cont'd)
 - Monitor agent to include GETRMI/TDCV2 based data: CPU utilization, Memory Utilization, key process status, cluster member status, Buffered I/O counts, operations counts, Network utilization, Disk utilization, Thresholding, console messages, “hot” files as per XFC.
 - Future plans include:
 - Oracle SPI
 - Performance agent
 - Port to IPF

OpenVMS Performance Tools & Alpha Performance Update

Bart Z. Lederman

Software Engineer
HP OpenVMS Engineering



Overview for this update section



- MONITOR
- Performance API (\$GETRMI)
- Trends.

MONITOR

- Re-written in C.
- Some screen improvements, potential for more enhancements later.
- Testing on new platforms revealed performance problems requiring a change in the recorded data format.
- Time is needed to notify everyone who may be affected, so it won't be in the next release.
- There will probably be an update kit to allow testing in the future.
- Existing MONITOR data will be read by the new MONITOR, or there will be a conversion utility.

Performance API

- There is an interface that MONITOR uses to get data.
- Others would like the same data.
- The interface is not documented. It can (and probably will) change without warning.
- The interface would not scale well.
- Changing it would break MONITOR.
- The alternative is for users to write “Inner Mode” code.
 - Risky.
 - Difficult to test.
 - Changes with operating system upgrades.

- Create a new interface based on the old one, but documented.
- Incorporate input from customers and third-party vendors.
- Data provider only: you write your own applications.
- “Cloned” from an existing service, so people who used the old one should be able to switch easily.
- New items added as needed (and resources allow).

\$GETRMI

\$GETRMI Get Resource Monitor Information

Returns system performance information about the local system.

`SYSS$GETRMI [efn] ,[nullarg] ,[nullarg] , itmlst ,[iosb] ,[astadr]
,[astprm]`

\$GETRMI

RMI\$_FRLIST number of pages on the freelist.

RMI\$_MODLIST number of pages on the modified page list.

RMI\$_FAULTS number of pages faults per second

RMI\$_PREADS number of pages read.

RMI\$_PWRITES number of pages written.

RMI\$_PWRITIO physical page write I/O's.

RMI\$_PREADIO physical page read I/O's.

...

There are about 526 items in the next release, a significant increase.

Disk Access

Year	Drive	Average Device Access Time (mS)	I/O per Second
1956	RAMAC	1015.0	.98
1964	2314	112.5	8.9
1975	3350	36.7	27.3
1987	3380K	24.6	40.6
1996	3390-3	23.2	43.1
1998	Cheetah 18	18.0	55.5
1996	Elite 23	19.0	52.6
1997	Elite 47	19.0	52.6
2000		12.0	83.3

CPU Speed.

Processor	Instruction Cycle (nS)
PDP-1	5,000 / 10,000
PDP-8I	1,500 / 3,000
PDP-15	800 / 1,600
LSI-11	2,380
PDP-11/70	300
VAX-11/785	274
μVAX-II	430
VAX 3600	180
VAX 8550	140
VAX 4000-M90	30
VAX 9000	17
AS200 4/166	6
AS2100 5/250	4
AS1000A 56/333	3
GS160 6/833	1+
GS1280 7/1150	<1

What's getting faster?

- Drive Capacity increase: > **30,000**
- CPU Performance increase: \cong **4,000**
- (with word size, pipelining, etc. \cong **40,000 to 60,000**)
- Drive Performance increase: **9**
- CPU and Memory performance increase at a faster rate than I/O, and I/O probably won't catch up any time soon.
- I/O bottlenecks will continue to be a concern.

I/O is more than disks.

- Networks.
 - 100 times increase in speed since 1980s.
 - Next 5 to 10 times increase underway.
 - All electronic, therefore easier to increase in speed.
 - Multiple paths allow load sharing.
- Cluster Interconnects.
 - See Networking.
- Backplane and Memory Bandwidth?


Throw Hardware at the Problem?

μVAX II & RD53
↓
VAX 3600 & RA82
↓
AS4100 & HSZ40
↓
??????

Throw Hardware at the Problem?

- Larger systems are being developed.
- Moving applications requires review and / or re-design more often than in the past.
- Moving to large multi-processor systems can require significant design work.

It's not just the "big" systems.



OpenVMS is working on this.

- Shadowing.
- Fastpath.
- Lock Manager.
- Cluster Interconnect and communications.
- Storage Solutions.

Locks, SMP Sync.

- Locking schemes that work well on slower systems don't always scale.
- SMP traps that work well on slower system don't always scale.
- SMP schemes that work on 2 or 4 processors don't always scale.
- User-written schemes should be reviewed.

Code stops working?

- Faster CPUs sometimes expose ‘sloppy’ programming (not checking status, etc.).
- Multi-processors sometimes expose ‘sloppy’ programming (shared status fields / blocks).
 - You may have a multi-processor system and not know it.
- Assuming something will complete in a certain time (or in no less than a certain time) will come back to haunt you.
- Assuming the system operates serially is not safe.

Keep your tools up-to-date

- `CC /OPT=(TUNE=EV6)`
- But only if you have the current version of the compiler!
- Keeping compilers, etc., up-to-date will be even more important in the future.

Other Resources

- Freeware CD
- ENCOMPASS (CD, Tapes, Library, DECUServe)
- I've included my own tools on previous collections, and on the current collection.
- www.decus.org (www.encompass.org)
- OpenVMS home page.



HP WORLD 2003

Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

