



**halcyonsoftware**

The Experts in Multi-Platform Systems Management

**Best Practice**  
**AIX Monitoring Guidelines**

## Table of Contents

|                                     |    |
|-------------------------------------|----|
| AIX Monitoring Guidelines .....     | 1  |
| Introduction .....                  | 3  |
| AIX Error Report .....              | 4  |
| Volume Groups .....                 | 4  |
| Logical Volumes .....               | 4  |
| Filesystems.....                    | 5  |
| Disks .....                         | 5  |
| CPU, Memory and Paging Space.....   | 5  |
| Process Monitor .....               | 6  |
| Server Configuration .....          | 7  |
| Log Files.....                      | 8  |
| Network .....                       | 8  |
| Network File System .....           | 9  |
| Appendix – AIX Error Messages ..... | 10 |

## **Introduction**

This document aims to provide a high level overview of baseline monitoring for the IBM AIX platform. The vast majority of monitoring elements detailed in this document have been included as a Central Configuration Manager template and are now shipped as part of Network Server Suite.

## AIX Error Report

A list of all AIX error codes is attached (see Appendix) and takes the following format:

Field 1 – AIX error identifier (8 digit hexadecimal)

Field 2 – Error label

Field 3 – Error type (4 character) of values: ( INFO, PEND, PERF, PERM, TEMP, UNKN)

Field 4 – Error class (1 character) of values:

H - Hardware

S - Software

O - Error logger

U - Undetermined

Field 5 – Error description

Very few AIX errors would not need to be alerted to the local administrators so it is advised that all are alerted on. Your software should poll the error log and any new entries should be alerted as standard.

As of AIX 5 you can limit the number of duplicates reported within a certain time frame. Hardware failure errors, Logical Volume Manager errors are generally critical and should be flagged as such.

Although not directly attributed to the Error Report its good practice to run a dumpcheck (/usr/lib/ras/dumpcheck) once a day at a peak time (this is normally added as a cron job by default on installation of AIX). If the system dump device is not big enough for a potential machine failure an error will be generated in the Error Report.

## Volume Groups

Defined Volume Group statuses should always be active and there should never be any Stale Partitions.

Check that Quorum is set to off for any mirrored volume groups.

## Logical Volumes

There should never be any Stale Physical Partitions associated with Logical Volumes resigning on system.

The normal working status for Logical Volumes is open/syncd.

## Filesystems

Space checking should be against percentages configured by the end user for each filesystem as monitoring depends on the size and usage of the individual filesystem. For example, a 1 TB filesystem housing Oracle datafiles would have very different requirements from a 256Mb /var filesystem on the same server.

For normal filesystems space less than 20% free should be alerted upon.

Key Filesystems /, /tmp, /usr, /home and /var should always be resident.

The availability of sufficient Inodes can be critical and you should monitor for less than 10% being free.

Journalled Filesystems should have separate logging devices.

## Disks

Disk activity taken over a number of samples over 80% should be alerted upon.

Verify MultiPath I/O (MPIO) path failures – the normal status should read enabled.

Any disks that have a status of missing or removed would tend to indicate an issue and should be investigated (these will also be written to errpt).

A system could potentially be I/O bound if %iowait > 25% and / or %tm\_act > 70% and should be investigated.

## CPU, Memory and Paging Space

From AIX 5.3 the processor's simultaneous multi-threading mode should be checked and alerted upon if it is capable but not enabled. However in some rare situations, threading off will increase performance.

Historically the runqueue would be an element to keep a close check (alerting if the runqueue exceeded the number of CPU multiplied by 2) but with more powerful P-Series servers the value being monitored for should be increased to suit.

Although not an error as such it's worth checking lparstat on lpar'd machines for over utilization.

Check for sustained (multiple samples of) %usr + %sys >=95%. This would tend to indicate a very busy box that could possibly be CPU bound.

The percentage of CPU time spent waiting on disk I/O (wio time) should be measured in samples and be alerted on >25% as this indicates a potential I/O bound machine.

Paging Space Utilization should be tracked and a general rule being < 30% used means that there is probably too much page space while >70% means there is probably too little page space.

The paging of the Paging Space is absolutely critical and would indicate a shortfall of RAM. Check this by monitoring for Paging > 0.

Useful additional commands used for sampling software for memory footprints:

**svmon -G** provides a global report. You can see the size of memory, how much is in use and the amount that is free. It provides details about how it is being used and it also provides statistics on paging space. All numbers are reported as the number of frames. A frame is 4 KB in size.

**svmon -Pt 3** displays memory usage of the top 3 memory-using processes sorted in decreasing order of memory demand.

## Process Monitor

Mandatory processes that should be alerted if not running;

- syncd
- errdemon
- inetd
- biod
- cron
- qdaemon
- portmap
- writesrv

Optional processes (may not be configured to run);

- rpc.lockd
- rpc.statd
- snmpd
- aixmibd
- snmpmibd
- hostmibd
- sendmail
- sshd

## Server Configuration

### ***Configuration files – Check for modifications.***

- /etc/inetd.conf
- /etc/netsvc.conf
- /etc/resolv.conf
- /etc/security/login.cfg
- /etc/inittab
- /etc/aixmibd.conf
- /etc/hostmibd.conf
- /etc/snmpd.conf
- /var/spool/cron/crontabs/root
- /etc/sendmail.cf
- /etc/ssh/ssh\_config (optional)
- /etc/ssh/sshd\_config (optional)
- /etc/profile
- /etc/environment

### ***Security***

Check for following lines in /etc/inetd.conf – any services not being used in the environment should be commented out. Note, they are all available by default.

- ftp
- telnet
- shell
- login
- exec
- ntalk
- daytime
- time

Check users with pwck for inconsistencies in /etc/passwd & /etc/security/passwd

## Log Files

|               |                           |                                    |
|---------------|---------------------------|------------------------------------|
| Console log   | /var/adm/ras/conslog      | (alog -o -t console)               |
| Failed logins | /etc/security/failedlogin | (who -a /etc/security/failedlogin) |
| Error Report  | /var/adm/ras/errlog       | (last)                             |
| Cron log      | /var/adm/cron/log         |                                    |

Conslog, failedlogin and errlog are all binary based files and must be accessed via the methods above.

## Network

Network problems are by their nature very difficult to spot and normally require further investigation. The starting points are listed here and are not currently shipped as a template with the Halcyon AIX Server Manager.

### **netstat -m :**

Look for any failures on mbuf allocation. Indicates "thewall" parameter not set high enough.

### **netstat -p udp :**

Look for "packets dropped due to no socket" != 0 and "socket overflows" >0 indicates udp send socket parameter too low.

### **netstat -p tcpip :**

Compare the number of packets sent to the number of data packets retransmitted. If the number of packets retransmitted is over 10-15% of the total packets sent, TCP is timing out indicating that network traffic may be too high for acknowledgements (ACKs) to return before a timeout. A bottleneck on the receiving node or general network problems can also cause TCP retransmissions.

Compare the number of packets received with the number of completely duplicate packets. If TCP on a sending node times out before an ACK is received from the receiving node, it will retransmit the packet. Duplicate packets occur when the receiving node eventually receives all the retransmitted packets. If the number of duplicate packets exceeds 10-15%, the problem may again be too much network traffic or a bottleneck at the receiving node

### **netstat -in :**

If the Oerrs column is greater than 1% of Opkts, the send queue size for that interface may need to be increased. If Ierrs is greater than 1 % of Ipkts, then memory may be a problem. The transmit queue size can be changed via SMIT or the **chdev** command. The mtu size can be changed by the **ifconfig** or **chdev** commands or through SMIT.

### **netstat -v**

Look if there is any value in the S/W Transmit Queue Overflow field which would indicate a need for a larger transmit queue size.

### **netstat -D**

The **-D** option of **netstat** displays the number of packets received (Ipkts), transmitted (Opkts) and dropped (Idrops, Odrops) in the communications subsystem. The important information seen here are the dropped packets particularly with the device drivers (dd). If packets are being dropped at the device driver then you want to increase the queue size on the device driver.

## **Network File System**

This section of the baseline document covers the Network File System (NFS) is intended as general tips and has therefore not been configured in a standard template of the Halcyon AIX Server Manager.

### **nfsstat -s :**

If bad calls are greater than 0 the server is rejecting RPC requests. Whenever the **nfsd** daemon is scheduled to run but doesn't find a packet on the NFS server queue, the *nullrecv* field gets incremented by one. The server may be running an excessive number of **nfsd** daemons – this can be checked by using the command `lssrc`.

*badlen* refers to an empty or truncated RPC packet. The packet could have been damaged by a network problem.

*xdr call* refers to an XDR header that may have been damaged. This is rare, but can happen more often if the network is a WAN rather than a LAN

### **For nfs clients:**

*badcalls* indicates RPC failures due to timeouts (if a server does not respond within a timeout period) and interrupts (if a file system mount is interrupted with the `intr` option).

This differs from the *badcalls* as shown under the NFS statistics, which indicate authentication errors.

*retrans* indicates the number of retransmissions because no response was received for the server. If there is poor server response time, *retrans* will have a high number.

## Appendix – AIX Error Messages

Attached is a full list of all AIX error messages that could potentially appear in the Error Report.



aix-errors.txt