Hercules: Advanced Topics

Jay Maynard, Savail Consulting, Inc.
SHARE 108, Session 2865
16 February 2007, Tampa, Florida

Mailing address:
Jay Maynard
Savail Consulting, Inc.
1831 Oakwood Drive
Fairmont, Minnesota 56031-3225

(952) 960-3600 voice
(952) 960-1334 fax

jmaynard@conmicro.cx

http://www.hercules-390.org
http://www.savail.com
http://www.conmicro.cx

Agenda

- New features in 3.05
- Development status
- How to become a developer
- Known problems
- Networking setup
- Building Hercules from CVS
- Dynamic device and instruction libraries
- Shared DASD
- Future directions

New features in 3.05

- Major SCSI tape fixes
- Crypto instruction fixes
- Partial 3590 support
- Hercules Automatic Operator
- 360 interval timer fixes
  - Can now run ASP without problems
- Tapes can be mounted read-only
- TUN/TAP fixes for 2.6.18 and later Linux kernels
- x86_64 performance improvement
- Decimal Floating Point support
- S/370 I/O emulation fixes

The x86_64 improvement was a fix of a dumb mistake: it wasn't using byte swapping acceleration from x86. This resulted in about a 10% performance improvement.

The DFP support is a demonstration of the power of open source: IBM released Mike Cowlishaw's (inventor of REXX) DFP package under a license permitting free reuse, so we were able to add it to Hercules easily and quickly.

Development status

- 40 members of zhercules developers mailing list
- 3.05 in active development
  - Nearly a year since 3.04, time for a formal release
- As IBM adds features, Hercules developers start work
- No real roadmap for future plans
  - Is one really needed?

How to become a developer

- Contribute good code
- Ask a current developer
- Invited to join the zhercules mailing list
- Granted CVS repository commit access
- Don't act like a bull in a china shop

Known problems

- S/370 I/O problems
  - Only appear on MVT if DASD, tape devices on same channel, I/O attempted to both
  - Fix in testing
  - May require new parameter on device statement

Networking basics

- Hercules network adapter appears as another network interface on host system
- Hercules guest OS must have own IP address assigned
  - Can be either on same subnet as host system or different subnet
- Host system and Hercules know how to communicate
- Rest of LAN must be told how to route packets to Hercules
- Most network problems are in routing, not Hercules setup

Networking on Linux

- Either CTC or LCS connections supported
- In either case, host configuration issues are the same
- The key: setting up proxy ARP
  - Allows host Linux to respond to address queries from other machines on LAN
  - Set up via sysctl variable:
    net.ipv4.conf.all.proxy_arp = 1
- Enable packet forwarding
  - Set up via sysctl variable:
    net.ipv4.ip_forward = 1
- sysctl variables read at boot time from /etc/sysctl.conf

Networking on Linux, continued

- If Hercules IP on different subnet, must tell other hosts on LAN how to route to it
  - Need to add route definition on each host that will connect
  - If using a single gateway to Internet, may be able to simply put route on that gateway

Networking on Linux, continued

- Set permissions on tun/tap device
  - Many different mechanisms, depending on device manager
    - udev: /etc/udev/permissions.d/<file>
    - devfs: /etc/devfs.conf
    - Neither: just modify device file
  - Recommendation:
    - Make file group read/write (mode 0660)
    - Create new group
    - Assign file to new group (chgrp)
    - Add your userid to new group in /etc/group

Networking on Windows

- Must use version 3.x of WinPCAP driver
  - Any 3.x version now works with Fish's DLL version 3.1 or greater
  - 32-bit only
- Fish's driver does proxy ARP automatically
- Enable packet forwarding by registry change or RAS snap-in for Server 2K/2K3
- Disable checksum offloading in driver settings

Proxy ARP is only needed on Windows for CTCI devices; LCS devices just pass packets back and forth.

The details on how to set up IP forwarding are documented on Fish's CTCI-W32 Configuration page, at http://www.softdevlabs.com/ctci-w32 .

Some network cards have the ability to computer the checksum of Ethernet packets in hardware. This works fine, until you have to modify the packet's contents, as the TUN/TAP driver does. You must disable this feature on the network adapter used for Hercules communication, or things will not work.

Networking on Windows, continued

- If more than one physical network interface on host system, may use the wrong one
  - Connection order in advanced networking settings
  - First one will be used
  - If this is not correct, specify adapter's IP or MAC address on device definition statement

Networking on Windows, continued

- When you have problems, follow the diagnostic script before asking for help
  - Supplied diagnostic programs collect information that's needed to figure out problem
  - If the diagnostics won't work, neither will Hercules networking
  - Details are on the CTCI-W32 Problem Reporting page

The problem reporting page is at http://www.softdevlabs.com/Hercules/ctci-w32-problem-reporting.html .

Networking on Mac OS X

- Need Matthias Nissler's tunnel driver
  - Supplied in Hercules installation package
  - Version-specific for Panther (10.3) and Tiger (10.4)
- Must run Hercules as root
  - Can't set up permissions for TUN/TAP device
  - Use sudo
- May have to set up routing entry manually
- OS X does proxy ARP by default
- Panther version quite stable
- Current Tiger beta version of driver fixes hang at Hercules shutdown

The beta driver can be found at http://www-user.rhrk.uni-kl.de/~nissler/tuntap/ .

Building from CVS

- The absolute latest, up-to-the-minute code
- Not guaranteed to compile, much less run
- The only way to get bleeding-edge features
- Downloading:
  - Set CVSROOT environment variable
  - Log in to CVS
  - Hit enter when asked for the password
  - Check out Hercules
  - New subdirectory named hercules will be created
- Once downloaded, can be synchronized efficiently

The CVSROOT variable should be set to

```
:pserver:anonymous@cvs.hercules-390.org:/usr/cvs/hercules
```

To log in to CVS:

```
cvs login
```

To check out the Hercules package:

```
cvs co hercules
```

To update the Hercules source: from within the Hercules directory,

```
cvs update
```

Building from CVS, continued

- More tools will be needed than when building from source tarball
  - autoconf, automake - used to build the configure script
- cvslvlck utility provided to ensure necessary versions are present on your system
  - Supplied in util subdirectory of hercules
  - Requires perl to run
  - Reported failures are not fatal, but the first place to look if problems are encountered

Building from CVS, continued

- Extra step needed before usual build process
  - configure script not stored in CVS
  - Built with autoconf and friends
  - Run autogen.sh first
  - ./configure, make as usual

To run `autogen.sh`, do:

```
sh autogen.sh
```

Dynamic libraries

- Allows adding or replacing Hercules function without recompiling
- Libraries can be loaded or unloaded during execution
- New devices can be added to the configuration
- Instructions can be added or replaced
- As of 3.04, can also add CGI functions to Hercules web server
- Basic architecture is about all that can't be changed

Dynamic libraries, continued

- Configuration file statements
  - MODPATH path - Specifies search path for load modules
    - > Appended to default search path
    - > Default is OS-dependent and appropriate in all but exceptional cases
    - > This parameter should never be needed
    - > If misspecified, standard devices won't be loaded
  - LDMOD module list - Specifies modules to be loaded
    - > Adds modules to default list to be loaded at startup

Dynamic libraries, continued

- Operator commands
  - ldmod - Loads a module
  - rmmod - Unloads a module
  - lsmod - Lists loaded modules
  - lsdep - Lists module dependencies

Dynamic libraries - devices

- All devices are dynamically loaded
- Library provides entry points for:
  - Initialization
  - Cleanup
  - CCW processing, channel program suspend/resume
  - Read, write, query used, reserve/release
  - Immediate CCW code table
  - Hercules suspend/resume processing

Dynamic libraries - instructions

- Usually used to add instructions not in POO
- Used for crypto instructions to deal with export restrictions
  - if module not present, crypto feature is disabled
- One approach to providing TCP/IP facilities in 370 mode
  - being developed by team writing TCP/IP stack

Shared DASD

- True DASD sharing
  - RESERVE, RELEASE CCWs implemented
- One Hercules owns DASD images and acts as server
- Other Hercules images connect via IP
- Server Hercules is not required to have anything IPLed
- Depending on network, can actually run faster than DASD on the running image
- Can emulate PAV feature on a single host

Shared DASD - example

- Server system configuration file
  - In system parameters section: SHRDPORT port
  - If no SHRDPORT statement, device server is disabled
  - No special parameters on device statements
  - All DASD devices are shared

Shared DASD - example, continued

- Client system
  - Each device is specified individually as host:port:devnum
  - Example:
    0125 3390 hercserver:3390:125
    0126 3390 192.168.0.4
    > defaults are port 3390, device number same as the one being defined

Future directions

- Restructured I/O subsystem (again!)
  - Split control unit, device I/O functions
  - More code, but allows more correct execution
- Tape functions split into device emulation and media sections
  - Similar to current split in DASD code
  - Allows better independence of device emulation from vagaries of specific media
- Get mailing lists the heck off Yahoo! Groups
  - Google Groups 2 beta is coming along