

## **The Hercules S/370, ESA/390, and z/Architecture Systems Emulator**

Jay Maynard, Global MAINTeCH Corporation  
SHARE 97, 25 July 2001, Minneapolis, Minnesota

Mailing address:

Jay Maynard  
Global MAINTeCH Corporation  
7836 Second Avenue South, Suite 1  
Bloomington, MN 55420

(281) 477-8331 voice

(208) 248-3975 fax

[jmaynard@globalmt.com](mailto:jmaynard@globalmt.com) (work)

[jmaynard@conmicro.cx](mailto:jmaynard@conmicro.cx) (home)

<http://www.conmicro.cx>

## Agenda



- Introduction
- Capabilities
- Device emulations
- OS compatibility
- Performance
- Installation and configuration
- Operation



## Introduction



- What is Hercules?
  - Emulates S/370, ESA/390, or z/Architecture CPU and peripherals
  - Runs under Linux, Windows NT/2000, BSD
  - Hardware emulation only
  - Freely available
  - Open Source
  - Active user community



Hercules emulates the hardware of an IBM mainframe system. It doesn't provide any software services at all, any more than the bare iron does. You have to run some sort of OS on it.

One strength of Hercules is its user base. There are over 1300 members of the Hercules mailing list, and a large number of them are involved in running, fixing, and enhancing the package, and in getting OSes, old and new, running properly on it.

## Capabilities



- Today
  - Will run most software written for S/370, ESA/390, and z/Architecture
  - Most architectural features that make sense for a single system supported
  - Some bugs in the corners of the spec remain
- Goal
  - Complete compatibility
  - Limited only by documentation and IP restrictions



The biggest missing architecture features are the ones that aren't documented in the POO and the like in sufficient detail to implement. Some aren't likely to be implemented; others will be if they can be reverse-engineered. A complete list of features implemented and not is in the Hercules FAQ.

Hercules version 2 added z/Architecture support, as well as providing all three architectures in one executable program. It was the first OEM product to support the z/Architecture, doing so a mere three weeks after the POO was available.

## Device emulations



- DASD
- Tape
- Card reader/punch
- Printer
- CTCA
- Terminal



Other devices can be implemented, but some (like, say, the 3705) would require a large amount of work, and provide relatively little gain. Others, like the 3172 Ethernet adapter, would be quite a bit of help to users, but aren't documented well enough to emulate properly.

## Device emulations: DASD



- Emulated via image on disk
- CKD, FBA supported
- Most classical CKD devices from 2311 to 3390
- Several FBA devices
- Regular CKD and FBA files compatible with P/390
- CKD device files can be compressed, with improved performance



Regular CKD and FBA devices take up as much disk space as the actual device has; an emulated 3350 will take up 300 MB, for example.

Compressed CKD devices can take up much less, and the actual space used depends not only on how well the data compresses, but also how much space is actually in use. Empty tracks take up no space at all. The smaller space used also means less real I/O needs to be done.

CCKD is a relatively recent addition to Hercules; while it has been extensively tested and many bugs squashed, there may still be some lurking bugs. I recommend keeping good volume backups if you use it.

There are a few devices and models currently missing, mostly FBA. They can be easily added if needed, as long as the device geometry parameters are known.

Hercules 2.13 added support for the 9345 CKD device.

## Device emulations: tape



- SCSI-attached tape devices
  - 9-track, 3480, 4mm DAT, 8mm, QIC tested
- AWSTAPE files
- OMA tape files
  - Both same format as P/390
- Hercules Emulated Tape (HET)
  - Enhanced AWSTAPE, with compression
- Can emulate 3420 or 3480 drives



Other SCSI tapes should work, as long as they follow the SCSI tape specification.

SCSI tapes are only supported under Linux. The Cygwin library does not emulate Linux SCSI tape processing. Support under BSD is somewhat limited.

HET is a new format, upwardly compatible with AWSTAPE. It adds compression support, which can give a real space saving. Utilities are provided to manipulate HETs, including converting to and from AWSTAPE.

## Device emulations: card reader/punch



- Card reader
  - ASCII and EBCDIC/binary input files supported
  - Translation automatically enabled if needed
  - Can IPL binary decks
  - Emulates 1442, 2501, or 3505
- Card punch
  - Emulated via output file on disk
  - Can punch with or without translation to ASCII
  - Emulates 3525



Reader file records don't have to be 80 characters long; they'll be padded if necessary, and an option is provided to truncate long lines if present. Multiple files can be sent to the reader either as separate decks or concatenated into one.



## Device emulations: printer



- Emulated via output file on disk
- Fixed carriage control (3211 FCB support coming)
- Emulates 1403, 3203/3211



Printed output is currently sent to a file; work is in progress on allowing it to be spooled directly to a printer.

## Device emulations: CTCA



- Emulated via external program, Linux device, or TCP port
- Design allows flexibility in actual facility
- One external program currently available: vmnet
  - Uses SLIP support in Linux kernel to establish connection
  - OS/390 reported to work, others may
- TUN/TAP device support for IP connections
- CTCA-to-CTCA to another Hercules system
- Emulates 3088, more or less



TUN/TAP is a new facility provided in the 2.4 series of Linux kernels to allow software to tap directly into the network, below the IP stack. It replaces the ethertap interface in 2.2, and is much more robust and usable. Hercules 2.09 introduced TUN/TAP support.

The 3088 emulation is definitely incomplete, mainly due to a lack of documentation. (Heard that before?)

Neither TUN/TAP nor vmnet are available under Windows. CTCA-to-CTCA connections are, however.

## Device emulations: terminal



- 3270
  - Local, non-SNA controllers
  - Emulates 3274-41D
  - Supports capabilities of client program
  - Emulated via tn3270 session
- 1052/3215
  - Local console only
  - Emulates 1052 or 3215 console
  - Emulated via regular telnet session



The actual 3270 emulation provided is limited by the client's capabilities, not Hercules itself. The TN3270 protocol places almost all of the responsibility for interpreting the datastream on the client.

The most compatible clients for Hercules are x3270, on Linux, and Tom Brennan's Vista tn3270 on Windows.

Line-mode terminals, except for local processor consoles, aren't supported because it would require going through the hoops to emulate a 270x or 37x5, a major undertaking for a small gain.

## OS compatibility



- No formal IBM testing
- User reported successes
  - OS/390: from 1.2 through 2.10
  - VM: ESA 2.2, 2.4, 1.1.0 (370 Feature); SP: r5, r6
  - VSE: ESA 1.3.2, 2.2.0 through 2.4.0; AF 3.2
  - Linux/390
- Licensing issues abound
  - How about a personal use license, IBM?



Hercules has not been certified officially as a plug-compatible system, mainly because that takes real money. Because of that, we can't officially claim that newer stuff will run, nor can you officially license modern OSES for it.

That said, users have reported running these OSES with Hercules. There's a bug in the relative instructions in Hercules versions before 2.09 that causes 0C4s in OS/390 2.10 catalog processing; earlier versions of OS/390 don't have that problem.

Hercules is in use for development of Linux/390, since it allows developers to participate in work on that system without access to an actual mainframe. The developers of 64-bit Linux for the zSeries use it, as well.

IBM has never had much demand for a personal use license for its products before. Now that Hercules makes running a real OS in your home feasible, there's a need. Other manufacturers do it. Hopefully, IBM will too.

## OS compatibility, continued



- Public domain OSes
  - OS/360
  - MVS 3.8J
  - VM/370 r6



There have been several mods tapes located recently, but more would be helpful. The biggest missing piece is HASP. If anyone can provide an installable copy, it would really help.

MVS 3.8J is the complete release ordered from IBM in May 2000. (Yes, you could order 5752-VS2 until March 1, 2001.) Everything they shipped is there.

VM/370 r6 is the complete release, already genned and ready to load and run. It runs guest OSes as well under Hercules as it did on the real iron.

## Performance



- Depends on host system
- Will make effective use of multiple host CPUs
  - Emulate multiple CPUs
  - Overlap I/O and CPU activity
- Dual 1 GHz Pentium III can sustain about 4 MIPS with moderate I/O load
- Host RAID subsystem will dramatically improve I/O performance



The biggest boost in Hercules performance comes from using a multiprocessor host system. Hercules overlaps I/O and CPU by launching a new thread of execution for every I/O. These threads are dispatched on any available CPU in an SMP host.

The next biggest boost comes from speeding up disk I/O. As with real hosts, the system spends a lot of time waiting on I/O, especially when one emulated I/O can take several real ones. Hardware RAID can overlap those I/Os, and caching can save I/Os entirely - a real win, just as with a mainframe.

## Installation



- Windows or Linux?
  - Some features not supported under Windows
    - Most significant is CTCA networking
    - SCSI tape
  - Cygwin libraries needed for Windows
  - Most development done first on Linux
- What else is that computer used for?
- Consider VMware



VMware works well for Hercules, regardless of the host OS you run. It's a good alternative, if a little pricey at \$299.

## Installation



- Available in either source tarball or RPM file for Linux, self-extracting installer for Windows
- Part of the FreeBSD packages collection
- Basic installation
  - Build package from source (if desired)
  - Install package
  - Create configuration file
  - Create DASD image files
  - IPL system



Building from source requires specifying the host architecture, either i386, i586, i686, Alpha, or anything else. The choice of which is used sets compiler options needed to make Hercules run correctly. Note that the Intel selections are used for any Intel-compatible CPU, including AMD, Cyrix, and so on. “Anything else” means SPARC, PowerPC, and the like.

To compile on Linux: explode the source tarball with

```
tar xzvf hercules-version.tar.gz
```

and then

```
cd hercules-version  
make HOST_ARCH=architecture
```

To install the compiled version:

```
make install
```

To install a packaged RPM:

```
rpm -Uvh hercules-version.architecture.rpm
```

where *version* is the version of Hercules, and *architecture* is the host system’s CPU architecture.



## Configuration



- Text configuration file
  - Analogous to IOCDS
  - Specifies system options and devices
  - Selected at Hercules startup by command-line option



Hercules is started, for version 2, by:

```
hercules -f config-file
```

You can have multiple configurations available by having more than one configuration file, but a single copy of Hercules can only use one configuration file. You can dynamically add and delete devices during execution, but some system-wide configuration parameters can only be changed by stopping and restarting Hercules.

```
# Sample configuration file
CPUSERIAL 000611
CPUMODEL  3090
MAINSIZE  64
XPNDSIZE  0
CNLSLPORT 3270
NUMCPU    1
NUMVEC    1
LOADPARAM 0120....
OSTAILOR  OS/390
SYSEPOCH  1900
TZOFFSET  -0500
PANRATE   FAST
ARCHMODE  ESA/390
#
# Device definitions
000A  1442  adrdmprs.rdr
000C  3505  jcl.txt  ascii trunc
000D  3525  pch00d.txt  ascii
000E  1403  prt00e.txt
001F  3270
0120  3380  mvsv5r.120
0580  3420  ickdsf.ipl
```

## Configuration: system options



- CPUSERIAL, CPUMODEL
  - Set values returned by STIDP instruction
- MAINSIZE, XPNDSIZE
  - Allocate main and expanded storage
- CNSLPORT
  - Sets the TCP port terminal sessions connect to
- NUMCPU, NUMVEC
  - Number of CPUs, vector facilities online at IPL



The CPUSERIAL and CPUMODEL parameters don't change anything in the emulator's behavior beyond the STIDP results; in particular, no attempt is made to implement any model's specific behavior.

MAINSIZE and XPNDSIZE allocate that much memory at Hercules startup and hold it as long as Hercules is up. Don't allocate more than your machine's physical RAM and swap space, and if possible don't allocate more than physical RAM for best performance. If you have to page, let Linux or Windows do the paging so page I/Os don't have to be emulated.

The same TCP port is used for both 3215 and 3270 connections. Connecting with a regular telnet client will connect to a 3215 device, and connecting with a tn3270 client will connect to a 3270 device.

The total number of CPUs and vector facilities available in the configuration is set at compile time. NUMCPU and NUMVEC set the number online at IPL. Use the facilities of the OS (such as `CF CPU (x)`, `ON` in MVS) to bring others online as needed.

## Configuration: system options, continued



- LOADPARAM
  - Same as IPL parameter on ESA hosts
- OSTAILOR
  - Turns off reporting for normal program checks during operation
- SYSEPOCH
  - Sets the year for TOD clock value of zero
- TZOFFSET
  - Adjusts clock to local time if desired



Without OSTAILOR, for example, VM will flood the Hercules console with messages about privileged operations exceptions, and OS/390 and Linux will do the same for other program checks. These are perfectly normal during operation for those OSES, and reporting them just wastes time and CPU cycles. OSTAILOR turns them off.

SYSEPOCH is intended to make older, non-Y2K OSES work properly by fooling the TOD clock. For OS/360, SYSEPOCH 1988 is recommended to set the emulated TOD clock to 1973 for dates in 2001, since OS/360 thinks dates beyond 1999 are invalid, and the calendars match between 2001 and 1973. For MVS 3.8 and VM/370, SYSEPOCH 1928 is recommended for the same reason. Only 1900, 1928, 1960, 1988, and 1970 are valid in Hercules 2.13.

TZOFFSET is included in case you want to run your TOD clock on local time instead of GMT.

Both SYSEPOCH and TZOFFSET are present because the emulated TOD clock cannot be set directly; they're derived from the host system clock.

## Configuration: system options, continued



- PANRATE
  - Sets refresh rate for Hercules control panel
- TODDRAG
  - Slows rate of emulated clock
- ARCHMODE
  - Selects the architecture to be emulated



PANRATE FAST will make the control panel act more like a real mainframe by updating it every 50 milliseconds instead of every 500. This will cost some CPU speed, though. You can also specify a refresh interval in milliseconds directly.

TODDRAG slows down the emulated TOD clock by the specified factor. TODDRAG 2 will make the emulated TOD clock count one minute for every two minutes of real time. This helps with OS/390 performance on slower host systems. TODDRAG 1 specifies no slowdown.

ARCHMODE can be one of S/370, ESA/390, or ESAME. ESAME is the formal name for the z/Architecture level.

## Configuration: device entries



- Device entries follow system options
- One per device
- Specified as address, device type, device parms
- Parm specify filename and options



Devices can be specified in any order. I recommend putting card readers and tape drives first, then the IPL disk, then any other devices, as this makes controlling them easier in the graphical control panel.

## Building DASD images



- Two utilities: dasdinit and dasdload
- dasdinit makes empty volumes
  - Creates volume label
  - Initialize and load with normal IBM utilities
- dasdload builds volumes with data
  - Builds VTOC, EREP datasets, minimal OS CVOL
  - Creates empty datasets
  - Loads PDSes created with TSO XMIT
  - Optionally writes IPL text



When creating a volume with dasdinit, include alternate tracks for any volume to be formatted with ICKDSF or IBCDASDI. These programs will seek to the alternate cylinders even if they don't intend to write there, just to make sure they're accessible.

## Operation



- Two control panels: graphical and command line
- Most commands available in both
- Usual operator facilities available: IPL, start, stop, interrupt, restart
- Device controls: attach/detach, interrupt, initialize
- Debugging: breakpoint, single-step, trace, register and memory alter/display
- HMC console commands and messages



“Graphical” is a bit of a misnomer; it’s not really graphical, just laid out on one screen in an easy-to-use manner.

Device initialization is also used to mount tapes and card decks on emulated devices during operation.

When devices are added to the configuration via attach, or removed via detach, the OS is notified as well; OS/390 will respond with the IOS150I DEVICE *nnnn* NOW AVAILABLE FOR USE message, just as it does on the real iron.

## Information on the web



- Hercules home page: <http://www.conmicro.cx/hercules>
  - Installation and operation documentation
  - Downloads
- Hercules mailing list:  
<http://groups.yahoo.com/group/hercules-390>
- Hercules on Windows:  
<http://www.bsp-gmbh.com/hercules>

