



Sun HPC ClusterTools 7.1 Software Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 819-7478-12
October 2007, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Netra, Solaris, Sun Ray, docs.sun.com, Sun HPC ClusterTools, Sun Cluster, Sun Fire, UltraSPARC, Sun Performance Library and Sun are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Cette distribution peut comprendre des composants développés par des tierces parties.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Netra, Solaris, Sun Ray, docs.sun.com, Sun HPC ClusterTools, Sun Cluster, Sun Fire, UltraSPARC, Sun Performance Library et Sun sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

AMD, Opteron, le logo AMD, et le logo AMD Opteron sont des marques de fabrique ou des marques déposées de Advanced Micro Devices.

Les produits qui font l'objet de ce manuel d'entretien et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.



Contents

1. Sun HPC ClusterTools 7.1	
Software Release Notes	1
Major New Features	1
Product Migration	2
Related Software	3
Disabling Installation Notification	4
Outstanding CRs	4
The MPI Library is Not Thread-Safe (CR 6474910)	4
udapl BTL Fails in Heterogeneous Cluster (CR 6512878)	5
Problems With Heterogeneous Support (CR 6538714)	5
Using udapl BTL on Local Zones Fails for MPI Programs (CR 6480399)	5
udapl BTL in Open MPI: If Multiple HCAs Exist, User Should Be Able to Select One or More (CR 6532415)	6
udapl BTL in Open MPI should Detect That a udapl Connection is Not Accessible and Not Just Hang (CR 6497612)	6
MPI Is Not Handling Resource Exhaustion Gracefully (CR 6499679)	6
libmpi_cxx.so Is Incompatible With Applications Linked With stlport (CR 6532412)	7
OMPI Supposed to Bind Processes Even If CPU Is Down With --mca mpi_paffinity_alone Set to 1 (CR 6478417)	7
Sun Grid Engine qsub -notify (or SIGUSR1/2) Is Not Supported (CR 6535841)	8

Request Script Prevents <code>SUNWompiat</code> From Propagating to Non-global Zone During Zone Creation (CR 6539860)	8
<code>rsh</code> Could Run Out of Sockets When Launching Jobs on a Large Cluster (CR 6541735)	8
<code>udapl</code> BTL Use of Fragment Free Lists Can Potentially Starve a Peer Connection and Prevent Progress (CR 6542966)	9
Showing Message Queues Would cause SEGV on DDT (CR 6615467)	10
All-to-One Communication Patterns Can Cause SEGV on Larger Clusters (CR 6617724)	10
TotalView: MPI-2 Support Is Not Implemented (CR 6597772)	11
TotalView: Message Queue for Unexpected Messages is Not Implemented (CR 6597750)	11
Behavior of TotalView Message Queue Support with 32/64 SPARC/AMD User Executables (CR 6623686)	11
Solaris OS Issues	12
HCTS Network Test Failed Due to E1000G Issue (CR 6462893)	12

Sun HPC ClusterTools 7.1

Software Release Notes

This document describes late-breaking news about the Sun HPC ClusterTools™ 7.1 (ClusterTools 7.1) software. The information is organized into the following sections:

- [“Major New Features” on page 1](#)
- [“Product Migration” on page 2](#)
- [“Related Software” on page 3](#)
- [“Disabling Installation Notification” on page 4](#)
- [“Outstanding CRs” on page 4](#)
- [“Solaris OS Issues” on page 12](#)

Major New Features

The major new features of the ClusterTools 7.1 software include:

- Open source MPI (message-passing interface). Previous versions of Sun HPC ClusterTools were based on Sun™ MPI, Sun’s implementation of the MPI-1 and MPI-2 standards. Sun HPC ClusterTools 7.1 is based on Open MPI, an open-source MPI that is also based on the MPI-1 and MPI-2 standards. Open MPI has been extended to support Sun Grid Engine. Open MPI is delivered through Sun packages.
- uDAPL/Infiniband support. ClusterTools 7.1 supports a byte-transfer layer (BTL) capable of communicating over Infiniband networks using the uDAPL protocol.
- Improved memory pool usage, which allows the uDAPL BTL to use the same memory pool as other BTLs (including OpenIB). This improvement results in more maintainable uDAPL code, and it allows all BTLs to take advantage of improvements to the memory pool code in the future.
- Simplified system administration tasks. Open MPI needs no system-level daemons, so system administrators no longer need to monitor and maintain them.

- Support for the Allinea DDT debugger, as well as improved support for the TotalView debugger.
- Support for Intel processor-based systems running the Solaris OS as well as SPARC- and AMD-based systems.
- Improved I/O forwarding. This allows files larger than 16384 bytes to be redirected to stdin.
- Improved performance on large clusters (1000+ nodes).

Product Migration

This section lists some of the differences between Sun MPI and Open MPI. For more information about migrating MPI applications built using Sun HPC ClusterTools 6 software to Sun HPC ClusterTools 7.1, refer to the *Sun HPC ClusterTools 7.1 Software Migration Guide*.

For more information about the installation process, see the *Sun HPC ClusterTools 7.1 Installation Guide*.

Some components of the ClusterTools 6 product are not available in ClusterTools 7.1.

The following tools in Sun HPC ClusterTools 6 have no equivalents in Open MPI:

- `mpprof`
- `mpps`
- `mpinfo`
- `mpkill`

The *Sun HPC ClusterTools 7.1 Software Migration Guide* discusses some possible alternatives you can use in place of these tools.

These are some of the differences between Sun HPC ClusterTools 6/Sun MPI and Sun HPC ClusterTools 7.1/Open MPI:

- The `mprun` command is now `mpirun`.
- There is only a CLI-based installation for ClusterTools 7.1. GUI-based installation (using the `ctgui` command) is no longer available.
- There are no master or client nodes (daemons) in ClusterTools 7.1 clusters. For this reason, there is no longer a need to set up authentication during installation. Instead, `rsh/ssh` is now used to access remote nodes in order to launch remote processes on jobs across those remote nodes.
- At activation, a node establishes symbolic links to ClusterTools 7.1.

- In ClusterTools 6, the run-time environment was CRE. The run-time environment in ClusterTools 7.1 is now ORTE (Open Run-Time Environment), which is a component of Open MPI.
- ORTE can work with either the OpenPBS or PBS Professional resource management software, but not both. PBS is not a required application for ClusterTools 7.1.
- LSF (Load Sharing Facility) software is not currently supported in ClusterTools 7.1. You must use a different job launcher, such as Sun Grid Engine.
- Sun HPC ClusterTools 6 software included both non-thread-safe (`libmpi.so`) and thread-safe (`libmpi_mt.so`) MPI libraries. Sun HPC ClusterTools 7.1 includes only a non-thread-safe version of the library. Thread-safe libraries are not supplied in Sun HPC ClusterTools 7.1. Thread-safe libraries might be supplied in a future version of Sun HPC ClusterTools software.
- The UltraSPARC III processor is the minimum version of the SPARC processor that supports Sun HPC ClusterTools 7.1 software. Sun HPC ClusterTools 6 software was supported on UltraSPARC II processors.

Related Software

Sun HPC ClusterTools 7.1 software works with the following versions of related software:

- Solaris™ 10 03/05 OS, or any subsequent Solaris 10 OS release that supports Sun HPC ClusterTools 7.1 software.

Note – If you plan to use the uDAPL BTL (Byte Transfer Layer) with your applications, you must install Solaris 10 11/06 OS. Solaris 10 11/06 OS is the first version that supports this functionality. In addition, you must install patch 125792-01 (for SPARC-based systems) or patch 125793-01 (for AMD-based systems), plus any patches that those patches require.

- Sun™ Studio 10, 11, and 12 C, C++, and Fortran compilers
- Distributed resource management (DRM) frameworks:
 - Sun Grid Engine Version 6 (was Sun N1 Grid Engine)
 - OpenPBS Portable Batch System (PBS) 2.3.16, Altair PBS Professional 7.1 (formerly Veridian), or Torque 2.1.19
- Debugging software:
 - The TotalView 8.3 debugger from TotalView Technologies (formerly Etnus) supports debugging MPI applications on SPARC- and AMD-based systems running the Solaris OS. Compatibility limitations might exist. See the

TotalView Web site at www.totalviewtech.com for a list of compilers that TotalView supports. In addition, see the Open MPI FAQ at <http://www.open-mpi.org> for more information about using TotalView with Open MPI.

- The DDT 2.1.3 debugger from Allinea also supports debugging on SPARC- and AMD-based systems running the Solaris OS.

Disabling Installation Notification

To improve ClusterTools, Sun collects anonymous information about your cluster during installation. If you want to turn this feature off, use the `-w` option for `ctinstall`.

The communication between `ctinstall` and Sun works only if the Sun HPC ClusterTools software installation process completes successfully. It does not work if the installation fails for any reason.

Outstanding CRs

This section highlights some of the outstanding CRs (Change Requests) for the ClusterTools 7.1 software components. A CR might be a defect, or it might be an RFE (Request For Enhancement).

Note – The heading of each CR description includes the CR's Bugster number in parentheses.

The MPI Library is Not Thread-Safe (CR 6474910)

The Open MPI library does not currently support thread-safe operations. If your applications contain thread-safe operations, they might fail.

Workaround: None.

udapl BTL Fails in Heterogeneous Cluster (CR 6512878)

The udapl BTL fails when run in clusters of heterogeneous nodes.

Workaround: Use a different interconnect between heterogeneous nodes.

Problems With Heterogeneous Support (CR 6538714)

Occasionally, programs run on heterogeneous clusters will return truncated messages or incorrect information. Some examples of these issues include:

1. MPI_LONG_DOUBLE datatypes do not work properly
2. Some of the one-sided APIs do not work properly
3. Some derived datatypes will not work properly

Workaround: None.

Using udapl BTL on Local Zones Fails for MPI Programs (CR 6480399)

If you run an MPI program using the udapl BTL in a local (non-global) zone in the Solaris OS, your program might fail and display the following error message:

```
Process 0.1.3 is unable to reach 0.1.0 for MPI communication.  
If you specified the use of a BTL component, you may have  
forgotten a component (such as "self") in the list of  
usable components.
```

```
PML add procs failed  
--> Returned "Unreachable" (-12) instead of "Success" (0)
```

```
-----  
*** An error occurred in MPI_Init  
*** before MPI was initialized  
*** MPI_ERRORS_ARE_FATAL (goodbye)
```

Workarounds: Either run the udapl BTL in the Solaris global zone only, or use another interconnect (such as tcp) in the local zone.

udapl BTL in Open MPI: If Multiple HCAs Exist, User Should Be Able to Select One or More (CR 6532415)

If your cluster contains multiple Infiniband Host Channel Adapter (HCA) cards, you cannot select a particular HCA when running using the `udapl` BTL.

Workaround: None.

udapl BTL in Open MPI should Detect That a udapl Connection is Not Accessible and Not Just Hang (CR 6497612)

This condition happens when the `udapl` BTL is not available on one node in a cluster. The Infiniband adapter on the node could be unavailable or misconfigured, or there might not be an Infiniband adapter on the node.

When you run an Open MPI program using the `udapl` BTL under such conditions, the program might hang or fail, but no error message is displayed. When a similar operation fails under the `tcp` BTL, the failure results in an error message.

Workaround: Add the following MCA parameter to your command line to exclude the `udapl` BTL:

```
--mca btl ^udapl
```

For more information about MCA parameters and how to exclude functions at the command line, refer to the *Sun HPC ClusterTools 7.1 Software User's Guide*.

MPI Is Not Handling Resource Exhaustion Gracefully (CR 6499679)

If an MPI job exhausts the resources of the CPUs, the program can fail or show segmentation faults. This might happen when nodes are oversubscribed.

Workaround: Avoid oversubscribing the nodes.

libmpi_cxx.so Is Incompatible With Applications Linked With stlport (CR 6532412)

If you compile and link a C++ program using the `-library=stlport4` flag, the generated program will produce a segmentation fault.

Workaround: Do not link to `stlport4`.

OMPI Supposed to Bind Processes Even If CPU Is Down With `--mca mpi_paffinity_alone` Set to 1 (CR 6478417)

If one CPU in a cluster fails or goes offline while processor affinity is set, processes assigned to the failed CPU attempt to bind to that processor. When the attempt fails, processor affinity is turned off and the processes continue to run, even though they do not bind to the failed processor.

This issue produces the following error message:

```
% /opt/SUNWhpc/HPC7.1/bin/mpirun --mca mpi_paffinity_alone 1 -mca
btl self,sm -np 4 a.out

output:
Num_procs = 3;
Num_procs = 3;
Num_procs = 3;
Num_procs = 3;
-----
The MCA parameter "mpi_paffinity_alone" was set to a nonzero value,
but Open MPI was unable to bind MPI_COMM_WORLD rank 0 to a
processor.

Typical causes for this problem include:

- A node was oversubscribed (more processes than processors), in
  which case Open MPI will not bind any processes on that node
- A startup mechanism was used which did not tell Open MPI which
  processors to bind processes to
```

Sun Grid Engine `qsub -notify` (or `SIGUSR1/2`) Is Not Supported (CR 6535841)

If you start an MPI job using Sun Grid Engine and use the `qsub -notify` option to send a `SIGUSR1` (impending `SIGSTOP`) or `SIGUSR2` (impending `SIGKILL`) signal to the job, `mpirun` and `orted` receive the `SIGUSR` signal from Sun Grid Engine. They route the signal to the program and then exit before the actual `SIGSTOP` or `SIGKILL` is received. This might cause ORTE to hang instead of properly exiting.

Request Script Prevents `SUNWompiat` From Propagating to Non-global Zone During Zone Creation (CR 6539860)

When you set up non-global zones in the Solaris OS, the Solaris OS packages propagate from the global zone to the new zones.

However, if you installed Sun HPC ClusterTools software on the system before setting up the zones, `SUNWompiat` (the Open MPI installer package) does not get propagated to the new non-global zone. It causes the `Install_Uutilities` directory to not be available on non-global zones during new zone creation. This also means that the links to `/opt/SUNWhpc` do not get propagated to the local zone.

Workaround: There are two workarounds for this issue.

1. From the command line, use the full path to the Sun HPC ClusterTools executable you want to use. For example, you would type `/opt/SUNWhpc/HPC7.1/bin/mpirun` instead of `/opt/SUNWhpc/bin/mpirun`.
2. Reinstall Sun HPC ClusterTools 7.1 software in the non-global zone. This process allows you to activate Sun HPC ClusterTools 7.1 software (thus creating the links to the executables) on non-global zones.

`rsh` Could Run Out of Sockets When Launching Jobs on a Large Cluster (CR 6541735)

Note – This issue affects both `rsh` and the Sun Grid Engine program `qrsh`. `qrsh` uses `rsh` to launch jobs.

If you are using `rsh` or `qrsh` as the job launcher on a large cluster with hundreds of nodes, `rsh` might show the following error messages when launching jobs on the remote nodes:

```
rcmd: socket: Cannot assign requested address
rcmd: socket: Cannot assign requested address
rcmd: socket: Cannot assign requested address
[node0:00749] ERROR: A daemon on node m2187 failed to start as
expected.
[node0:00749] ERROR: There may be more information available from
[node0:00749] ERROR: the 'qstat -t' command on the Grid Engine
tasks.
[node0:00749] ERROR: If the problem persists, please restart the
[node0:00749] ERROR: Grid Engine PE job
```

This indicates that `rsh` is running out of sockets when launching the job from the head node.

Workarounds:

1. If you are using `rsh` as your job launcher, use `ssh` instead. Add the following to your command line:

```
-mca pls_rsh_agent ssh
```

2. If you are using Sun Grid Engine as your job launcher, you can modify the Sun Grid Engine configuration to allow Sun Grid Engine to use `ssh` instead of `rsh` to launch tasks on the remote nodes. The following web site describes how to perform this workaround:

http://gridengine.sunsource.net/howto/qrsh_qlogin_ssh.html

Note that this workaround does not properly track resource usage.

udapl BTL Use of Fragment Free Lists Can Potentially Starve a Peer Connection and Prevent Progress (CR 6542966)

When using a peer-to-peer connection with the `udapl` BTL (byte-transfer layer), the `udapl` BTL allocates a free list of fragments. This free list is used for send and receive operations between the peers. The free list does not have a specified maximum size, so a high amount of communication traffic at one peer might increase the size of the free list until it interferes with the ability of the other peers to communicate.

This issue might appear as a memory resource issue to an Open MPI application. This problem has only been observed on large jobs where the number of uDAPL connections exceeds the default value of `btl_udapl_max_eager_rdma_peers`.

Workaround: For example, if an Open MPI application running over uDAPL/IB (Infiniband) reports an out of memory error for `alloc` or for privileged memory, and if those two values have already been increased, the following might allow the program to run successfully.

1. At the command line, add the following MCA parameter to your `mpirun` command:

```
--mca btl_udapl_max_eager_rdma_peers x
```

where x is equal to the number of peer uDAPL connections that the Open MPI job will establish.

2. If the setting in [Step 1](#) does not fix the problem, then also set the following MCA parameter with the `mpirun` command at the command line:

```
--mca mpi_preconnect_all 1
```

Showing Message Queues Would cause SEGV on DDT (CR 6615467)

When using the Allinea DDT debugger message queue with Sun HPC ClusterTools 7.1 and Open MPI 1.2.4, DDT might cause a segmentation fault (SEGV) within the message queue DLL.

Workaround: This is a known issue in Open MPI 1.2.4. It will be fixed in a subsequent version of Open MPI.

All-to-One Communication Patterns Can Cause SEGV on Larger Clusters (CR 6617724)

Intermittent segmentation faults (SEGVs) can occur when using all-to-one communications (for example, when an `MPI_Put` operation sends the data from all ranks to rank 0, with rank 0 reporting on the origin of each part of the data). These faults occur in clusters with large numbers of processors.

Workaround: None.

TotalView: MPI-2 Support Is Not Implemented (CR 6597772)

The TotalView debugger might not be able to determine if an MPI_Comm_spawn operation has occurred, and might not be able to locate the new processes that the operation creates. This is because the current version of the Open MPI message dumping library (ompi/debuggers/ompi_dll.c) does not implement the functions and interfaces for the support of MPI 2 debugging and message dumping.

Workaround: None.

TotalView: Message Queue for Unexpected Messages is Not Implemented (CR 6597750)

The Open MPI DLL for the TotalView debugger does not support handling of unexpected messages. Only pending send and receive queues are supported.

Workaround: None.

Behavior of TotalView Message Queue Support with 32/64 SPARC/AMD User Executables (CR 6623686)

The TotalView debugger only provides a 32-bit debugger for Solaris SPARC and a 64-bit debugger for Solaris AMD. As a result, the following message might appear when TotalView attempts to load the message queue library.

```
The image claims that TotalView should use the dynamic library
'/opt/SUNWhpc/HPC7.1/lib/sparcv9/openmpi/libompitv.so'
for MPI debugging, but that file is not accessible.
This is probably an MPI installation problem.
TotalView will try to use the default debug library 'libtvmpeich.so'
```

This is a known issue with the TotalView debugger.

Workaround: Before you start TotalView, set the LD_LIBRARY_PATH or LD_LIBRARY_PATH_64 variable to the correct value (shown below) to ensure the correct message queue shared library is loaded.

- To show message queues on Solaris SPARC with a 32-bit user application, you do not need to set either variable.

- To show message queues on Solaris SPARC with a 64-bit user application, set `LD_LIBRARY_PATH` to `/opt/SUNWhpc/HPC7.1/lib/openmpi`
- To show message queues on Solaris AMD with a 32-bit user application, set `LD_LIBRARY_PATH_64` to `/opt/SUNWhpc/HPC7.1/lib/amd64/openmpi`
- To show message queues on Solaris AMD with a 64-bit user application you do not need to set either variable.

Note – Using the `msg_lib` parameter in the Totalview Parallel Configuration file to specify the location of the message queues library has no effect on the way Totalview finds the location of the library. Use the workaround described in this section until Totalview resolves this issue.

Solaris OS Issues

This section discusses issues that pertain to the Solaris OS. Although Sun HPC ClusterTools 7.1 supports the Solaris 10 3/05 release, updating to a later release (such as Solaris 10 11/06) can help you avoid these issues altogether.

HCTS Network Test Failed Due to E1000G Issue (CR 6462893)

When using the `tcp` BTL on the Sun Fire X4500 server with a high level of network traffic, running a program might result in data corruption.

Workaround: This issue has been fixed in the Solaris 10 11/06 release. It is strongly suggested that you upgrade your Solaris OS to Solaris 10 11/06 or a compatible release.

A workaround does exist for the Solaris 10 3/05 release, but you might experience performance degradation when using this workaround.

- **Use a text editor to add the following line to your `/etc/system` file:**

```
set ip:dohwcksum=0
```

This disables the hardware checksum and allows the program to complete successfully.