

AIX Questions



Compiled by Jeff Simon

The AIX Solution Provider Technical Support Group in Austin, Texas, supports software vendors who are developing or porting applications to AIX. This article is a compilation of questions that are frequently asked by vendors. The name of the responding Technical Support Group staff member appears after each response.

What is the limit for the number of files that can be opened by a single process?

The limit is 2,000 file descriptors per process on operating systems prior to BOS 4.3. It is not possible to exceed this limit. `OPEN_MAX` (`/usr/include/sys/limits.h`) is the maximum number of file descriptors that can be used to create an array of file descriptors in the process user block. Again, this array size cannot be altered.

BOS 4.3.1 increased this limit to 32,767.

—Jeff Simon

What is AutoFS?

AutoFS, which is being shipped with BOS 4.3.1, includes a port of the Sun[®] ONC+ implementation of automatic mounting and replaces the current automount facility. AutoFS has the following three components:

- ◆ The `automount` command
- ◆ AutoFS filesystem kernel extension
- ◆ Automount daemon

AutoFS allows map files and other changes to be implemented without stopping and restarting the automount daemon.

—Jeff Simon

What is CacheFS in BOS 4.3?

The Cache File System (CacheFS) is a general-purpose filesystem caching mechanism available in AIX[®] 4.3 that improves NFS server performance and scalability by reducing server and network load. Designed as a layered filesystem, CacheFS provides the ability to cache one filesystem on another.

In an NFS environment, CacheFS increases the client-per-server ratio, reduces server and network loads, and improves performance for clients on slow links, such as point-to-point protocol (PPP). Any filesystem that is primarily read is a likely candidate for caching. In AIX 4.3, CacheFS only supports the caching of NFS Version 2 and Version 3 filesystems.

—Jeff Simon



Jeff Simon

I have installed BOS 4.3. When I try to run `vmstat`, I get the following error:

```
0551-129 knlist failed
```

This problem usually can be resolved by ensuring that the link `ls -l /unix` is set up. The link should look like this:

```
unix -> /usr/lib/boot/unix_mp
```

Set up the proper link if needed:

```
ln -s /usr/lib/boot/unix_mp  
/unix
```

—Jeff Simon

What is the maximum number of continuation lines or size that the `xlf <4/5>` compiler can handle?

The maximum size of continuation lines for `xlf <4/5>` is 6700 bytes.

—Jeff Simon

Why are threads blocked inside the `clnt_create()` RPC call?

Generally, it is not considered safe to create threads in C++ code. For example, the use of I/O streams in the threaded C++ programs may cause problems. (`iostream.h` and `complex.h` are *not* threadsafe).

The `clnt_create()` remote procedure call (RPC) uses the `gethostbyname()` routine, which is not a thread-safe routine. This results in the `clnt_create()` call not being a true thread-safe function.

The `gethostbyname` returns a structure. However the `gethostbyname_r` function, which is thread safe, requires you to allocate and pass in the structure within your thread (plus a work structure). This makes the parameters different.

Routines that are not thread safe are sub-routines that use static or global variables—routines returning a pointer to a static structure.

RPC calls on AIX are not currently thread safe. When writing RPC calls, developers must put locks around the RPC calls to ensure that only one thread comes in at a time to prevent deadlocking.

—David McCloud

Why do I have problems mounting filesystems exported from Sun servers?

IBM supports Sun NFS Version 4.03 and newer versions. On earlier versions of `nfs`, all NFS calls had to use port numbers that were reserved (below 1024). That is not necessary with the new versions (Version 4.2.1 and greater) unless an `nfs` option `nfs_use_reserved_ports` is set. The problem may be that the other clients are expecting the RISC client to use a reserved port.

This problem has two possible solutions:

- ◆ Turn off port checking on the PC server. This may not be allowed because of security considerations, but it could be tried if the client allows you to unset port checking.
- ◆ Use the command `nfs -o nfs_use_reserved_ports=1` to tell the RISC client to use reserved ports on its calls.

If you are using BOS 4.2.1, you must use APAR# IX71891 for this setting to work.

The `noacl` option is another NFS mount option to consider to disable the ACL RPC checks. Disabling the ACL queries can be done by mounting with this option
(`mount -o intr,noacl
server:/nfs_server_dir /nfsdir`).

—David McCloud



David McCloud

How do you load a member of an archive using dlopen?

Figure 1 demonstrates how to load a member of an archive using dlopen.

```
compile: cc -g dltest.c -o dltest -ldl

/*-----dltest.c -----*/
#include <stdio.h>
#include <dlfcn.h>
#include <errno.h>
int main()

{
void *handle = dlopen("libXm.a(shr.o)",RTLD_MEMBER|RTLD_NOAUTODEFER);
/* you need to use the object module shr.o.
use "dump -Hv /usr/lib/libXm.a" to find which module to use*/

if (handle != 0)
{
void *addr = dlsym(handle,"XtToolkitInitialize");

if (addr == (void *)0)
{
printf("PASS: dlopen() : ok\tHandle = %p\n",handle);
if (errno == ENOEXEC)
{
printf("PASS: Find on XtToolkitInitialize ok, dlerror says:
[%s]\n",dlerror(
));
}
else
perror("PASS: Find on XtToolkitInitialize ok");
}

else
{
printf("PASS: loaded libXm.a and found libXt.a\n");
printf("INFO: %p=dlopen(), %p,=
dlsym(xx,XtToolkitInitialize)\n",handle,addr)
;
}
dlclose(handle);
}
else
{
if (errno == ENOEXEC) {
printf("Failed to load libXm.a - ok, dlerror says:
[%s]\n",dlerror());
}
else
perror("Failed to load libXm.a - ok");
}
}
}
```



Asthma Saudagar

Figure 1. Using dlopen to load a member of an archive

—Asthma Saudagar



What are the different logical volume types on AIX?

AIX has five logical volume types, shown in Figure 2.

Volume Type	Description
Filesystem	Holds filesystem data and metadata
Log	Holds journaled filesystem(JFS) metadata update log
Paging	Paging areas
Boot logical volume	Boot block and RAM filesystem code
Dump area	Holds panic dumps

Figure 2. AIX logical volume types

—Jeff Simon



My build line is very long and I get errors. Is there a workaround?

If the command line is too large for the linker to handle, it will give an error. To circumvent this, use the `-F` flag on the command line to insert a modified configuration file.

Copy the `/etc/xlC.cfg` to a local file `myconfig.cfg` and include the command-line options in the file under the appropriate stanza. InfoExplorer™ provides greater detail. Figure 3 shows an example.

```
xlC -Fmyconfig.cfg:stanza  
-Fconfig_file:stanza | -Fconfig_file | -:stanza
```

Figure 3. Using the `-F` flag to insert a modified configuration file

—Asthma Saudagar



Compiled by Jeff Simon, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Simon has worked in AIX technical support since 1990 and currently works as a technical lead and system administrator for a series of technical Web sites. He has a BS in Computer Science from Southwest Texas State University and certification in AIX support.