

# ***ptx/IBM-ESS V2.0.0*** ***Release Notes***

## **CAUTION**

*This release of ptx/IBM-ESS runs under DYNIX/ptx V4.5.1 **only**.  
For information on IBM ESS subsystem support under DYNIX/ptx V4.4.7,  
refer to the ptx/IBM-ESS V1.x.x release notes.*

This document contains information to support the installation and use of ptx/IBM-ESS V2.0.0, a layered software package that enables selected IBM NUMA-Q systems to communicate with IBM ESS 2105 storage subsystems. Topics covered in this document include:

- ptx/IBM-ESS overview, including ESS features supported in this release
- Requirements for use
- Caveats and limitations
- Known issues
- Related documentation
- Software installation
- Hardware integration overview
- Usage and administration notes

---

Copyright © 2000 by International Business Machines Corporation. All rights reserved. This document may not be copied or reproduced in any form without permission. Information in this document is subject to change without notice.  
Printed in the United States of America.

## Introduction

---

The ptx/IBM-ESS V2.0.0 software package enables communications between an IBM NUMA-Q host system running DYNIX/ptx V4.5.1 and an IBM Enterprise Storage Server (ESS) subsystem. The subsystem must be configured as a RAID-5 subsystem, and must be connected to the NUMA-Q system via Fibre Channel Bridge (FCB) SCSI bus.

Specific requirements for both the NUMA-Q system and the ESS are found under *Requirements for Use*, below. *Only connections between NUMA-Q systems and ESS subsystems that meet the conditions listed under **Requirements for Use** are supported in this release.*

The following are the main functionalities supported in this release of ptx/IBM-ESS:

- SCSI connection via Fibre Channel Bridge (FCB)
- Read/Write I/O, with block sizes from 512 bytes to 128KB
- Support for RAID-5 operation
- SCSI multiport target support
- Support for a maximum of 8 LUNS/target ID

The ptx/IBM-ESS software uses IBM NUMA-Q's **sd** device driver, and includes a software module, *sd\_ess*, that enables communications with the IBM 2105 disk devices.

No VTOCs are currently supplied; however, a DYNIX/ptx utility, **make\_vtoc**, is available as a patch (FP#249428) to the base operating system for creation of the appropriate VTOCs.

## Multiporting Overview

---

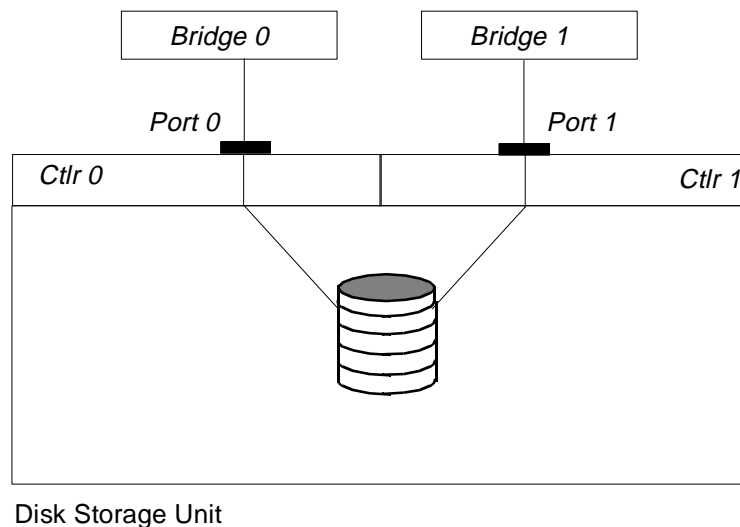
Certain disk storage units provide multiple physical connections (ports) to multiple I/O buses and then configure their logical units to allow access from those ports simultaneously. This feature is called *multiporting*.

Multiporting provides redundant access to individual logical units (LUNs), increasing their availability. If a particular port or I/O bus becomes unavailable, an unaffected port and/or I/O bus can be used to access the LUN. Locating multiple ports for a LUN on separate buses provides protection from both an I/O bus failure and the failure of a port within a storage unit. With multiporting:

- A port can have more than one LUN connected to it.
- A single LUN can be connected to more than one port. E.g., in Figure 1, the disk resource is connected to Port 0 and Port 1, and has a path through both bridges.

The operating system may also take advantage of multiporting to better distribute its I/O load. Doing so *may* reduce the latency of individual I/O requests, although this is not guaranteed.

---



---

**Figure 1.**  
**Sample Multiported Disk Storage Unit.**

## Requirements for Use

---

### NOTE

*For timely resolution of problems, the NUMA-Q/ESS solution should have the following capabilities:*

*Remote dial-in access to the ESS subsystem.*

*Remote dial-in access to the NUMA-Q system(s).*

*Call home SEQUENTLINK access to the NUMA-Q system(s).*

*For information about these capabilities, contact your IBM NUMA-Q Customer Support Center.*

### NUMA-Q System Requirements

To qualify for connection to an ESS subsystem via this release of ptx/IBM-ESS, the NUMA-Q system must meet the following requirements:

- Must be running DYNIX/ptx V4.5.1.
- Must be configured for FCB SCSI, with FCBRIDGE SW V1.5.3.
- Must have enough FCB SCSI channels available for dedicated connection of the ESS.
- If you intend to create custom VTOCs, the **make\_vtoc** utility must be installed. This utility is available as Fastpatch #249428.

### ESS Subsystem Requirements

To qualify for connection to a NUMA-Q system running this release of ptx/IBM-ESS, the ESS subsystem must meet the following requirements.

- Customers with ESS subsystems running V1.0 codebase *must upgrade to ESS V1.1 codebase* before installing or executing this release of ptx/IBM-ESS. When migrating an ESS subsystem to a new codebase, customers should contact their appropriate IBM NUMA-Q and IBM ESS support organizations.
- Must be configured for *differential* SCSI operation. NUMA-Q systems do not support single-ended SCSI connections.
- Must be configured as SCSI Bus Type Numaq\_Unix; see the discussion titled *Hardware Installation and Integration*, later in this document, for details.
- Must be configured as a RAID-5 subsystem. *JBOD operation is not supported.*
- Must be running a supported microcode release. For the currently supported microcode release(s), contact your IBM NUMA-Q Customer Support Center.

### NOTE

*There is no 2105 Host Attachment CD-ROM for attachment to IBM NUMA-Q hosts provided with the ESS subsystem. The ptx/IBM-ESS software supersedes this requirement.*

## ESS Specialist Requirements

To run ESS Specialist successfully, the console PC for ESS should meet the following requirement:

- Before running Specialist, you should check the *virtual memory total paging file size* of the console PC and make sure that this parameter is between 139MB and 189MB. The default value for the PC may be significantly smaller; if so, it should be increased. Virtual memory configurations smaller than this may cause the PC to run out of virtual memory.

## Cabling Requirements

- Connection between the NUMA-Q system and the Fibre Channel Bridge must be via multimode FC cable supplied by IBM NUMA-Q.

### NOTE

*For SCSI connections, total external cable length, FCB connector to ESS connector, cannot exceed 25m.*

- Each SCSI connection between the NUMA-Q system and the ESS requires one (1) NUMA-Q SCSI cable (NSBO-FCBXSS — 3m cable) and one (1) IBM ESS SCSI cable (features code 9701 — 10m cable or 9702 — 20m cable).

## ptx/IBM-ESS Software Requirements

- This release of ptx/IBM-ESS requires a software license key to execute. To obtain a key, contact your IBM NUMA-Q Customer Support Center and provide them with the *nodeid* of the target NUMA-Q system.

# Caveats and Limitations

---

The following caveats apply to the installation and use of this software package.

## General Caveats

- Only the NUMA-Q and ESS hardware and software configurations listed under *Requirements for Use*, earlier in this document, are supported.
- No data throughput speeds are expressed or implied. Overall performance, including throughput, is dependent on data/system load and other conditions that IBM NUMA-Q does not control.
- This product has not been tested in NUMACenter environments, and operation in NUMACenter environments is not supported in this release.
- ESS disks cannot be used as boot or dump devices.
- The standard DYNIX/ptx VTOC driver may work with ptx/IBM-ESS devices. However, stand-alone diagnostics and format utilities will not work. Attempting to use these programs will lead to unpredictable results.

## SCSI Connectivity Caveats

- The NUMA-Q systems support *only* differential SCSI connections.
- SCSI IDs 6 and 7 are reserved for FCB SCSI ports. Select a SCSI ID for the subsystem in the range 0—5 or 8—15. *Do not* configure the subsystem with SCSI ID 6 or 7.
- The FCB must be *dedicated* to ESS devices; there can be no other devices connected to the FCB.
- For adequate performance, IBM NUMA-Q recommends that no more than (4) four SCSI buses per FCB be connected. Two (2) SCSI buses per FCB is the configuration recommended for optimal performance.
- Multi-initiator SCSI is not supported. For path redundancy in NUMA-Q configurations, use multipointing via multipathing (refer to the discussion titled *Multipointing Overview*, earlier in this document).
- The NUMA-Q Fibre Channel Bridge (FCB) does not support Ultra-Wide SCSI protocol. However, it negotiates SCSI-2 protocol.
- DYNIX/ptx currently supports *not more than* 8 LUNs (LUNs 0—7) per SCSI target ID. Since target IDs 6 and 7 are reserved for the FC Bridge, each FC Bridge connection is limited to accessing no more than 112 **sd** devices.

## ESS Configuration and Support Caveats

- The IBM ESS subsystem must be running a microcode release supported by this release of ptx/IBM-ESS. Contact your IBM NUMA-Q Customer Support Center for the currently supported release(s) of ESS microcode.
- Concurrent microcode updates to the ESS are not certified on NUMA-Q. Scheduled downtime is required to perform microcode updates.

- This release supports only ESS subsystems configured as RAID-5 subsystems. JBOD configurations are not supported.
- Only ESS subsystems configured for SCSI operation are supported. For support of any ESCON connections, contact the appropriate server organization(s).
- This release makes no provision for upgrades from IBM ESS V1.0 to IBM ESS V1.1. For upgrade information, contact your IBM ESS support organization.
- This release supports up to eight (8) LUNs (numbered 0—7) per target ID. No more than 8 LUNs per target ID can be defined.
- Point-in-time copies via FlashCopy are not supported as part of this solution from IBM NUMA-Q. Contact the appropriate organization(s) regarding this functionality.
- Peer-to-peer Remote Copy (PPRC) is not supported as part of this solution from IBM NUMA-Q. Contact the appropriate organization(s) regarding this functionality.

## New or Changed in This Release

---

The following are new or changed in this release of ptx/IBM-ESS.

- This release runs *only* under DYNIX/ptx V4.5.1.
- Command queueing is *enabled* in this release, with a queue depth of 4.

## Problems Resolved in This Release

---

When running DYNIX/ptx V4.5.1 and IBM ESS codebase V1.1, the following items that were reported as problems in the ptx/IBM-ESS V1.0.0 release are no longer problems.

### **I/O on Alternate Path May Hang During Recovery of Original Path**

With the ESS V1.1. codebase, this problem no longer occurs.

### **Route Recovery May Result in Reinstate/Suspend Loop**

This problem has been fixed in the DYNIX/ptx V4.5.1 base operating system.

### **devctl -c May Not Bring Back All Devices**

With resolution of the **Route Recovery May Result . . .** problem, listed above, this is no longer a problem.



## Known Issues

---

The following are known problems or issues in this release of ptx/IBM-ESS.

### Hard Errors on Alternate Paths After Primary Path Failure

After I/Os that fail on the primary path are re-tried on the alternate path, hard errors on WRITE (Busy) are reported on the alternate path. Entries in **ktlog** will be similar to the following:

```
39008831 09:56:17 tolog/warn q1/e7/p2149 sd68(port 137): hard error on WRITE (Busy)
```

This problem is intermittent and has only been observed in scenarios in which either the SCSI bus is interrupted via a cable pull on the active channel or the SCSI bus fails under very specific conditions.

**Workaround.** None.

### I/O To a Path That Is Being Restored May Report Hard Errors

I/O to a path that is being restored may report WRITE (command aborted key/asc/ascq 0b/4e/00) and I/O timeout errors on its devices during the restoration of the path.

This occurs when active I/O is going down both paths and a **devctl -d <scsibus num>** with a follow-on **devctl -c <fchr num>** is executed on one cluster node. The other cluster node reports **ktlog** errors similar to the following:

```
13:54:32 warning:sd6(port 44): ioid 0x3C5AC300: soft error on WRITE (command aborted
key/asc/ascq 0b/4e/00)
13:54:32 warning:sd6(port 44): buffer 0x3C5AC300 being retried on route 0x3B4E2E40 -
sd6(port 172)
13:56:06 info:sd6(port 44): I/O timeout, hung bus?
13:56:06 warning:sd6(port 44): buffer 0x3C5BBF00 being retried on route 0x3B4E2E40 -
sd6(port 172)
13:56:06 warning:sd6(port 44): ioid 0x3C5B1AC0: soft error on WRITE (I/O timeout) route
0x3a021480
13:56:06 warning:sd6(port 44): buffer 0x3C5B1AC0 being retried on route 0x3B4E2E40 -
sd6(port 172)
13:56:06 info:sd6(port 44): resuming route
```

When this occurs, the failed I/Os are retried on the alternate path and routes are then resumed on the original path.

This problem has been detected in a 2-node cluster in which the nodes share the same SCSI paths.

The cause of this problem is under investigation.

## Related Documentation

---

The following documentation contains essential or useful additional information for installing and using this release of ptx/IBM-ESS:

- NUMA-Q system hardware and software documentation
- DYNIX/ptx V4.5 software installation documentation
- The following IBM Enterprise Storage Server documents:
  - *IBM Enterprise Storage Server* (SG24-5465-00)
  - *Introduction and Planning Guide, Models E10 and E20* (GC26-7294-01)
  - *Configuration Planner* (SC26-7353-00)
  - *Implementing the Enterprise Storage Server in Your Environment* (SG24-5420-00)
  - *User's Guide, 2105 Models E10 and E20* (SC26-7295-00)
  - *Host Systems Attachment Guide, 2105 Models E10 and E20* (SC26-7296-00)
  - *Quick Configuration Guide* (SC26-7354-00)
  - *Microcode Installation Instructions* for the correct 2105 microcode release
  - *Web Interface User's Guide* (SC26-7346-00)—for ESS Specialist software

# Software Installation

---

## Installation Notes

Installing ptx/IBM-ESS modifies the following DYNIX/ptx base operating system files to support IBM 2105 disk devices:

```
/usr/conf/uts/io/sd/sd_space.c  
/usr/conf/uts/symmetry/unix_std.a
```

Before modifying any files, the install script copies the following:

```
sd_space.c to /usr/lib/ibm_ess/sd_space.c
```

All object files used to modify *unix\_std.a* are saved under */usr/lib/ibm-ess/.save*.

## Installation Procedure

To install this software package, do the following:

1. Log in as **root**.
2. Make sure that the correct version of DYNIX/ptx is installed.
3. Install the ptx/IBM-ESS software. To do this:
  - a. Insert the ptx/IBM-ESS CD-ROM into the NUMA-Q system's CD drive.
  - b. Refer to the DYNIX/ptx software installation documentation and install the software. This should be a standard CD-ROM installation, with the following specifications:

```
Installation source type = CD-ROM  
Installation type = root  
Software Package Part Number = 1003-80183-00
```

There should be no file conflicts. If you abort the installation for any reason, then begin the installation procedure again, the menu system will ask if you want to continue the current installation; answer **yes**.

### CAUTION

*Do not recompile the kernel at this time. The recompile should take place **only** after the license key has been entered, as described below.*

4. You must enter a software license key for the ptx/IBM-ESS software to execute. If you do not have a key, contact your IBM NUMA-Q Customer Support Center and provide them with the *nodeid* of the NUMA-Q system. Once you have the key, enter it as follows:
  - a. Change to the directory */usr/conf/uts/io/sd*, and open the file *sd\_space.c* for editing.
  - b. In the file *sd\_space.c*, find the line that reads:

```
char *sd_ibm_ess_license_key = "No LicenseKey";
```

- c. Edit this line, replacing the string `No LicenseKey` with the ptx/IBM-ESS license key provided by the IBM NUMA-Q Customer Support Center. *Be sure that the key is enclosed in quotes and that the entire `No LicenseKey` string has been replaced.*
- d. Save and exit the file `sd_space.c`.
- e. Save a hard copy of the license key. If you should ever need to re-install the ptx/IBM-ESS software, you will need to re-enter the key.

5. Compile the new kernel.

If you are installing the ESS subsystem at this time, continue on to the discussion titled *Hardware Installation and Integration*, following. To reduce system downtime, we recommend that you do not reboot the system until the subsystem has been installed and connected.

# Hardware Installation and Integration

---

Once the ptx/IBM-ESS software has been installed, the following steps must be completed to integrate the ESS subsystem with the NUMA-Q system.

- Install and configure the subsystem.
- Connect the subsystem to the NUMA-Q host system.
- Verify the integration.

## Install and Configure the Subsystem

### CAUTIONS

*Only qualified IBM personnel should attempt to install and configure the subsystem.*

*In addition, please ensure that NUMA-Q and ESS installation personnel coordinate their integration efforts.*

Contact your ESS Customer Support Center for subsystem installation and configuration.

### CAUTION

*Refer to the **Requirements for Use** discussion, earlier in this document, for the PC console requirements to run ESS Specialist.*

### CAUTION

*Be sure to install any other required ESS software **before** installing the ESS operating system or LIC CD-ROMs. Refer to the appropriate microcode installation instructions.*

**NOTE:** *There is no 2105 Host Attachment CD-ROM for IBM NUMA-Q host types provided with the ESS software. The ptx/IBM-ESS layered software supersedes this requirement.*

The tasks necessary to configure the subsystem are discussed in the IBM document titled *Implementing the Enterprise Storage Server in Your Environment, Chapter 15. ESS configuration for fixed block storage.*

Configuring the subsystem requires the use of the IBM ESS Specialist. Information concerning the use of this software package is found in the *IBM Enterprise Storage Server Web Interface User's Guide*.

When configuring the subsystem, follow the procedures in the two documents referenced above. For connecting the ESS to NUMA-Q systems, the ESS should be configured as a SCSI Bus Type of Numaq\_Unix. To do this, refer to the Web Interface Users Guide and do the following:

### CAUTION

*The following procedure should only be done AFTER the LIC and PTF CD-ROMs have been loaded and are fully operational.*

1. Select the Numaq\_Unix host system type.
2. Apply the Numaq\_Unix bus configuration to each of the SCSI ports that connect to a NUMA-Q FCB.

## Connect the Subsystem to the NUMA-Q System

To connect the ESS subsystem to the NUMA-Q host system, do the following:


1. Warn all users that the NUMA-Q system will be coming down.
2. Perform an **init 0** on the NUMA-Q system.

### **NOTE**

*The ESS subsystem will take approximately 30 minutes to boot.*

3. Make sure the ESS subsystem is ready.
4. Refer to Figure 2 and make cable connections between the ESS subsystem and the NUMA-Q system's FCB switch.
5. Boot the NUMA-Q system.

**Example**  
**2105 E20 (SCSI)**  
**RAID 5 Configuration**



**NOTE:** ESCON cards may be replaced by SCSI cards. If this is done appropriate configuration changes will be required.

**Figure 2.**  
**2105 Configuration and Cabling.**

## Verify the Integration

To verify the integration:

1. After the NUMA-Q system has been booted with the new kernel (that includes ESS support), enter:

```
/etc/dumpconf -d | grep sd
```

and verify that the appropriate **sd** devices were found at system boot. You should see one entry for each data path to each IBM 2105 **sd** device.

2. You can then use **dumpconf** to determine which **sd** devices are IBM 2105 devices:

```
/etc/dumpconf | grep <scsibus #>
```

where **<scsibus #>** is the SCSI bus on which the **sd** devices reside. For additional detail on these devices, use the **diskid <device #>** command; refer to Figure 4 for an example of the output of **diskid**.

### NOTE

*In the following illustrations, firmware revision levels are for example purposes only. Firmware revision levels for your subsystem may be different. If you have any questions, contact your IBM NUMA-Q Customer Support Center for currently supported ESS microcode release(s).*

Figure 3 shows a typical ESS port ID mapping that identifies the association of ports to SCSI adapters, bays, and clusters. Figure 4 shows sample output of **diskid**, **infodev**, and **dumpconf** commands for 2105 devices.

### NOTE

***dumpconf** reports port information in hexadecimal.  
**ktlog** reports information in decimal.  
**diskid** reports the **cport ID** in decimal.*

3. After the devices have been verified, perform I/O to each device to verify its operation.
4. Complete the preparation of the subsystem(s) for use. This may include building VTOCs and filesystems or configuration using **ptx/SVM**. If you are using VTOCs, make sure that the required fastpatch has been installed, then refer to the **DYNIX/ptx make\_vtoc** on-line man page, along with the README file included with the fastpatch for usage information.
5. Notify users that the system is now available for use.

This completes the installation of **ptx/IBM-ESS** and the ESS subsystem.



Cluster #		1	1	1	1	1	1	1	1	Physical Labels
Bay #		1	1	1	1	2	2	2	2	
Adapter #		1	2	3	4	1	2	3	4	
Top (Port A)	HEX	0x01	0x05	0x09	0x0d	0x21	0x25	0x29	0x2d	
	DEC	1	5	9	13	33	37	41	45	Logical IDs
Bottom (Port B)	HEX	0x00	0x04	0x08	0x0c	0x20	0x24	0x28	0x2c	
	DEC	0	4	8	12	32	36	40	44	
Cluster #		2	2	2	2	2	2	2	2	
Bay #		3	3	3	3	4	4	4	4	
Adapter #		1	2	3	4	1	2	3	4	
Top (Port A)	HEX	0x81	0x85	0x89	0x8d	0xa1	0xa5	0xa9	0xad	
	DEC	129	133	137	141	161	165	169	173	
Bottom (Port B)	HEX	0x80	0x84	0x88	0x8c	0xa0	0xa4	0xa8	0xac	
	DEC	128	132	136	140	160	164	168	172	

**NOTES**

The Cluster #, Bay #, and Adapter # represent physical labels on the subsystem.  
 The Port ID is a logical assignment within each cluster. This value is shown in hex in **dumpconf**, and in decimal for both the **diskid cport** number and the **ktlog** port number.

**Figure 3.**  
**Port mappings for SCSI connection.**



# Usage and Administration Notes

---

## ESS Specialist

When the ESS Specialist GUI is set to *Graphical Display*, the display shows path connections as viewed by the device adapters. This display can be confusing: you may see two paths to a single host adapter and conclude that a loop on that adapter has more paths defined than are actually defined.

**Workaround.** At the top right of the GUI display, select *Tabular View* to see the path definitions for loops.

## Missing, Invalid or Expired License Key Will Cause Errors

If the ptx/IBM-ESS software license key is missing, invalid, or expired you will see warning messages in **ktlog**, similar to the following:

```
387a4251 12:34:25 warn q0/e0/p276 ibm_ess:947536465c:: Error <0xbad2>
IBM_ESS Driver disabled

387a4251 12:34:25 tolog/warn q0/e0/p276 mpt_set_device_num: +sd135:
set device num callout failed
```

If the license key is missing, you will see messages similar to the following on the console or in */usr/adm/messages*:

```
WARNING: ibm_ess:1b:: Error <0xbad1> IBM_ESS Driver disabled
WARNING: sd(unit 0 on +scsibus2): unable to get unit serial number (custom)
WARNING: sd(unit 1 on +scsibus2): unable to get unit serial number (custom)
...
```

In addition, on the NUMA-Q system console, you will see **devctl** failures as the kernel attempts to assign permanent names to the *ibm\_ess* devices. This will cause **dumpconf** to report devices with *temporary* device names of the form:

```
+sd<device_number>
```

If these symptoms occur, check that the software license key has been entered correctly and is still valid. If you do not have a license key or if it has expired, contact your IBM NUMA-Q Customer Support Center.

## Use Caution When Replacing VTOCs

If there is a VTOC on a disk, you can modify it or replace it (e.g., one built with **make\_vtoc**) without rebooting the system or using **devctl -M**. However, if you do this, the old "partition/slice" device nodes will still be present and *no new device nodes will be built*. Therefore, it is possible that some of the nodes required by the new VTOC will be missing. If this occurs, writes attempted to that partition will occur on the *host system's main hard drive* as a regular file write rather than to the drive in the subsystem.

If you attempt to replace VTOCs without rebooting or using **devctl**, you will need to check that all necessary device nodes are present; if they are not present, you will need to build them manually with **mknod(1M)**.

The safest method is to reboot or use **devctl** when changing VTOCs.

### **Using ptx/IBM-ESS with ptx/SVM**

If you are using ptx/IBM-ESS with ptx/SVM, you should *not* use **newfs** on volumes that are under SVM control. Instead, use **mkfs** to reduce the number of copies of superblocks, thereby increasing the usable disk space and improving performance. On RAID-5 partitions *not* under SVM control, you can use **newfs**, as usual.