

# 3PAR InForm<sup>®</sup> OS 2.3.1

## Command Line Interface Reference

3PAR, Inc.  
4209 Technology Drive  
Fremont, CA 94538 USA

Part No. 320-200166 Rev B  
March 2010

## Revision Notice

This is the second release of this manual. A complete revision history is provided at the end of this manual.

## Changes

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, 3PAR Inc. assumes no liability resulting from errors or omissions in this document or from the use of the information contained herein.

3PAR Inc. reserves the right to make changes in the product design without reservation and without notification to its users.

## Updates to the Documentation at 3PAR Central

Any updates to this document, or other 3PAR technical documents, can be found by logging into 3PAR Central's Document Control System from 3PAR's Support page, at: <http://support.3PAR.com>.

## 3PAR Technical Support and Services

Contact your local service provider for technical support and services at: <http://www.3PAR.com/services.html>.

## Sales and Ordering Information

For sales and ordering information contact:

3PAR, Inc.

4209 Technology Drive

Fremont, CA 94538 USA

Telephone: 510-413-5999

Fax: 510-413-5699

E-mail: [salesinfo@3PAR.com](mailto:salesinfo@3PAR.com)

## Reader Comments and Suggestions

Please E-mail your comments and suggestions about this document to [ReaderComments@3PAR.com](mailto:ReaderComments@3PAR.com).

## Copyright

Copyright © 2009-2010, 3PAR Inc. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of 3PAR Inc., 4245 Technology Drive, Fremont, CA 94538. By way of exception to the foregoing, the user may print one copy of electronic material for personal use only.

## NetBSD Notices

Certain sections of the InForm OS that handle crash dumps were derived from NetBSD under the BSD license. 3PAR, Inc. provides the following notices in accordance with the original license.

Copyright (c) 1996, 1997, 1998, 2001 The NetBSD Foundation, Inc. All rights reserved.

The code obtained from NetBSD is derived from software contributed to The NetBSD Foundation by Charles M. Hannum and by Jason R. Thorpe of the Numerical Aerospace Simulation Facility, NASA Ames Research Center.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:  
This product includes software developed by the NetBSD Foundation, Inc. and its contributors.
4. Neither the name of The NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

Copyright (c) 1998, 1999, 2000, 2001 Manuel Bouyer.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement:  
This product includes software developed by the University of California, Berkeley and its contributors.
- 4 Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

Copyright (c) 1996, 1998 Christopher G. Demetriou. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement:  
This product includes software developed by Christopher G. Demetriou for the NetBSD Project.
- 4 The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

Copyright (c) 1998 The NetBSD Foundation, Inc. All rights reserved.

The code obtained from NetBSD is derived from software contributed to The NetBSD Foundation by Charles M. Hannum, by Onno van der Linden and by Manuel Bouyer.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement:  
This product includes software developed by the NetBSD Foundation, Inc. and its contributors.
- 4 Neither the name of The NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

Copyright (c) 1991 The Regents of the University of California. All rights reserved.

The code obtained from NetBSD is derived from software contributed to Berkeley by William Jolitz.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement:  
This product includes software developed by the University of California, Berkeley and its contributors.
- 4 Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

The following applies to all of these notices:

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT OWNERS LISTED ABOVE AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Tcl

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, ActiveState Corporation and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses.

Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

GOVERNMENT USE: If you are acquiring this software on behalf of the U.S. government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as defined in Clause 252.227-7013 (c) (1) of DFARs. Notwithstanding the foregoing, the authors grant the U.S. Government and others acting in its behalf permission to use and distribute the software in accordance with the terms specified in this license.

### **GNU General Public License Materials**

The InForm OS uses the Linux kernel and lkcdutils crash dump utilities. The Linux kernel and lkcdutils crash dump utilities have been modified slightly by 3PAR Inc. and, as modified, are licensed under the GNU General Public License.

Copyright © 2002-2003 3PAR Inc.

A copy of the GNU General Public License is available on the CD-ROM provided by 3PAR and may additionally be obtained at <http://www.fsf.org/licenses/gpl.html>. As required by this license, for a period of three years after you receive the Linux kernel and lkcdutils crash dump utilities from 3PAR, a copy of the source code for such software, as modified, may be obtained from 3PAR at 3PAR's cost of providing you with such code.

The InForm OS uses a Linux gigabit adaptor base driver distributed by Intel under the GNU GPL. The driver has been modified slightly by 3PAR, Inc. and, as modified, is licensed under the GNU GPL.

Copyright © 2002, 2003, 3PAR, Inc.

A copy of the GNU GPL is available on the CD-ROM provided by 3PAR and may additionally be obtained at <http://www.fsf.org/licenses/gpl.html>. As required by this license, for a period of three years after you receive the Linux gigabit adapter base driver from 3PAR, a copy of the source code for such software, as modified, may be obtained from 3PAR at 3PAR's cost of providing you with such code.

The InForm OS contains hardware and firmware protocol definitions for the LSI Logic Fusion MPT architecture. These definitions are licensed under the GNU GPL.

Copyright © 2000-2002 LSI Logic Corporation.

A copy of the GNU GPL is available on the CD-ROM provided by 3PAR and may additionally be obtained at <http://www.fsf.org/licenses/gpl.html>. As required by this license, for a period of three years after you receive these definitions from 3PAR a copy of the source code may be obtained from 3PAR at 3PAR's cost of providing you with such code.

### **GNU Lesser General Public License Materials**

The InForm OS uses the following unmodified GNU LGPL libraries: glibc (Copyright © 1991-2001 Free Software Foundation, Inc), libgmp (Copyright © 1991, 1993-2002 Free Software Foundation, Inc), libncurses (Copyright © 1998 Free Software Foundation, Inc), libpopt (Copyright © Red Hat Software), and libstdc++ (Copyright © 1986-2000 Free Software Foundation, Inc). These libraries are licensed under the GNU Lesser General Public License.

A copy of the GNU Lesser General Public License is available on the CD-ROM provided by 3PAR and may additionally be obtained at <http://www.fsf.org/licenses/lgpl.html>. A copy of the source code for such software may be obtained from 3PAR or from <http://www.debian.org>

### **OpenSSL License Materials**

The InForm OS uses the unmodified libssl OpenSSL library. This library is licensed under dual licenses, the OpenSSL License and the SSLeay License.

Copyright (c) 1998-1999 The OpenSSL Project. All rights reserved.

Redistribution and use of the libssl OpenSSL library in source and binary forms, with or without modification, is permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:  
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used endorse or promote products derived from this software without prior written permission. For written permission, please contact [openssl-core@openssl.org](mailto:openssl-core@openssl.org).
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:  
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

This product includes cryptographic software written by Eric Young ([eyay@cryptsoft.com](mailto:eyay@cryptsoft.com)). This product includes software written by Tim Hudson ([tjh@cryptsoft.com](mailto:tjh@cryptsoft.com)).

#### **Original SSLeay License**

Copyright (C) 1995-1998 Eric Young ([eyay@cryptsoft.com](mailto:eyay@cryptsoft.com))  
All rights reserved.

This package is an SSL implementation written  
by Eric Young ([eyay@cryptsoft.com](mailto:eyay@cryptsoft.com)).  
The implementation was written so as to conform with Netscape's SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson ([tjh@cryptsoft.com](mailto:tjh@cryptsoft.com)). Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed.

If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used.

This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use of the libssl OpenSSL library in source and binary forms, with or without modification, is permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:  
"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)." The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:  
"This product includes software written by Tim Hudson (tjh@cryptsoft.com)."

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e., this code cannot simply be copied and put under another distribution license [including the GNU Public License.]

### **Other Open Source Materials**

The InForm OS uses the unmodified zlib library.

Copyright © 1995-1998 Jean-loup Gailly and Mark Adler.

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

**Trademarks**

3PAR, InServ, InForm, InSpire and Serving Information are registered trademarks of 3PAR Inc.

Intel and Pentium are registered trademarks of Intel Corporation.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, and Windows NT, Exchange Server, and SQL Server are either registered trademarks or trademarks of Microsoft Corporation.

Linux is a registered trademark of Linus Torvalds.

Redhat is a registered trademark of Red Hat, Inc.

SuSE is a registered trademark of Novell, Inc.

Oracle is a registered trademark of Oracle Corporation.

Sun, Solaris, and Java are trademarks of Sun Microsystems, Inc.

UNIX is a registered trademark of The Open Group.

All other trademarks and registered trademarks are owned by their respective owners.





# Table of Contents

---

## **1 Introduction**

1.1 Audience	<b>1.2</b>
1.2 User Interfaces	<b>1.2</b>
1.3 Supported Operating Systems	<b>1.2</b>
1.4 Related Documentation	<b>1.3</b>
1.5 Organization	<b>1.4</b>
1.6 Typographical Conventions	<b>1.6</b>
1.7 Advisories	<b>1.6</b>

## **2 CLI Command Syntax and Conventions**

2.1 Syntax and Conventions	<b>2.2</b>
2.2 Syntax Rules	<b>2.3</b>
2.3 Glob-Style Pattern	<b>2.4</b>
2.4 Typical Command Layout in this Book	<b>2.5</b>
2.5 Global Options and Environment Variables	<b>2.6</b>
2.6 Exit Status	<b>2.6</b>

## **3 Commands Quick Reference**

3.1 Overview	<b>3.2</b>
3.2 Disk Enclosure Management Commands	<b>3.2</b>

3.2.1	Drive Cage Management	<b>3.2</b>
3.2.2	Physical Disk Management	<b>3.3</b>
3.3	Domain Management Commands	<b>3.3</b>
3.4	Health and Alert Management Commands	<b>3.5</b>
3.4.1	Alerts	<b>3.5</b>
3.4.2	Events	<b>3.5</b>
3.4.3	System Health	<b>3.5</b>
3.5	Help and Utility Commands	<b>3.6</b>
3.6	Task Management Commands	<b>3.7</b>
3.7	LDAP Management Commands	<b>3.7</b>
3.8	Licensing Management Commands	<b>3.8</b>
3.9	Node Subsystem Management Commands	<b>3.8</b>
3.9.1	Firmware Versions	<b>3.8</b>
3.9.2	Node Date Information	<b>3.9</b>
3.9.3	Node Properties	<b>3.9</b>
3.9.4	Node EEPROM Log	<b>3.9</b>
3.9.5	Array and Node Information	<b>3.10</b>
3.9.6	Network Interface Configuration	<b>3.10</b>
3.9.7	Port Information	<b>3.11</b>
3.9.8	Battery Management	<b>3.12</b>
3.9.9	System Manager	<b>3.13</b>
3.10	Performance Management Commands	<b>3.13</b>
3.10.1	Chunklet Statistics	<b>3.13</b>
3.10.2	Data Cache Memory Statistics	<b>3.14</b>
3.10.3	Node CPU Statistics	<b>3.14</b>
3.10.4	Logical Disk Statistics	<b>3.14</b>
3.10.5	Link Statistics	<b>3.15</b>
3.10.6	Physical Disk Statistics	<b>3.15</b>
3.10.7	Port Statistics	<b>3.15</b>
3.10.8	System Tuner	<b>3.15</b>
3.10.9	Dynamic Optimization	<b>3.16</b>

3.10.10	Virtual LUN (Export) Statistics	<b>3.16</b>
3.10.11	Virtual Volume Statistics	<b>3.17</b>
3.11	Preserved Data Commands	<b>3.17</b>
3.12	Replication Commands	<b>3.18</b>
3.12.1	Physical Copy	<b>3.18</b>
3.12.2	Remote Copy	<b>3.18</b>
3.12.3	Virtual Copy	<b>3.21</b>
3.13	Service Commands	<b>3.21</b>
3.13.1	Disk Enclosure	<b>3.21</b>
3.13.2	General System Maintenance	<b>3.22</b>
3.13.3	System Upgrade	<b>3.22</b>
3.14	SNMP Agent Commands	<b>3.23</b>
3.15	CIM Server Commands	<b>3.24</b>
3.16	Sparing Commands	<b>3.24</b>
3.17	SSH Access Commands	<b>3.25</b>
3.18	Task Schedule Commands	<b>3.26</b>
3.19	User Management Commands	<b>3.26</b>
3.20	Volume Management Commands	<b>3.27</b>
3.20.1	Common Provisioning Group Management	<b>3.27</b>
3.20.2	Host Management	<b>3.28</b>
3.20.3	Logical Disk Management	<b>3.29</b>
3.20.4	Space and Storage Management	<b>3.30</b>
3.20.5	Template Management	<b>3.31</b>
3.20.6	Virtual Volume Management	<b>3.32</b>
3.20.7	Virtual LUN (Export) Management	<b>3.34</b>

## **4 Add Command**

addsnmpmgr	<b>4.2</b>
------------	------------

## **5 Admit Commands**

admithw	<b>5.2</b>
---------	------------

admitpd	5.4
admitrcopylink	5.6
admitrcopytarget	5.8
admitrcopyvv	5.10
<b>6 Cancel Command</b>	
canceltask	6.2
<b>7 Check Commands</b>	
checkhealth	7.2
checkId	7.4
checkpassword	7.6
checkpd	7.8
checkport	7.11
checkvv	7.13
<b>8 CLI Command</b>	
cli	8.2
<b>9 Compact Commands</b>	
compactcp	9.2
compactId	9.4
<b>10 Control Commands</b>	
controlscsiport	10.2
controlmag	10.5
controlpd	10.7
controlport	10.9
<b>11 Create Commands</b>	
createald	11.3
createaldvv	11.9

createavv	<b>11.18</b>
createcpg	<b>11.22</b>
createdomain	<b>11.30</b>
createdomainset	<b>11.32</b>
creategroupsv	<b>11.34</b>
creategroupvvcopy	<b>11.37</b>
createhost	<b>11.40</b>
createhostset	<b>11.43</b>
createld	<b>11.45</b>
creatercopygroup	<b>11.48</b>
creatercopytarget	<b>11.50</b>
createsched	<b>11.52</b>
createspare	<b>11.56</b>
createsv	<b>11.58</b>
createtemplate	<b>11.61</b>
createtpvv	<b>11.71</b>
createuser	<b>11.77</b>
createvlun	<b>11.79</b>
createvv	<b>11.83</b>
createvvcopy	<b>11.90</b>
createvvset	<b>11.94</b>

## **12 Dismiss Commands**

dismisspd	<b>12.2</b>
dismissrcopylink	<b>12.3</b>
dismissrcopytarget	<b>12.5</b>
dismissrcopyvv	<b>12.6</b>

## 13 Free Command

freespace	13.2
-----------	------

## 14 Grow Commands

growaldv	14.2
growavv	14.10
growtpvv	14.12
growvv	14.14

## 15 Hist Commands

histch	15.2
histld	15.6
histpd	15.10
histport	15.18
histvlun	15.23
histvv	15.28

## 16 Locate Commands

locatecage	16.2
locatesys	16.4

## 17 Move Commands

movech	17.2
movechtospare	17.6
movepdtospare	17.9
moverelocpd	17.12
movetodomain	17.18

## 18 Promote Commands

promotesv	18.2
promotevvcopy	18.4

## 19 Remove Commands

removealert	19.3
removecpq	19.5
removedomain	19.7
removedomainset	19.8
removeeventlog	19.10
removehost	19.12
removehostset	19.14
removeld	19.16
removercopygroup	19.18
removercopytarget	19.20
removesched	19.22
removesnmpmgr	19.24
removesnmppw	19.26
removespare	19.28
removesshkey	19.30
removetask	19.31
removetemplate	19.33
removeuser	19.35
removeuserconn	19.37
removevlun	19.39
removevv	19.43
removevvset	19.46

## 20 Service Commands

servicecage	20.2
servicehost	20.5
servicemag	20.8

## 21 Set Commands

setalert	21.3
setauthparam	21.5
setbattery	21.12
setcage	21.14
setcim	21.16
setclienv	21.18
setcpg	21.20
setdate	21.27
setdomain	21.31
setdomainset	21.33
sethost	21.35
sethostset	21.39
setlicense	21.40
setnet	21.42
setnode	21.47
setntp	21.49
setpassword	21.51
setpd	21.53
setrcopygroup	21.55
setrcopytarget	21.62
setsched	21.66
setsnmppw	21.68
setsshkey	21.70
setstatch	21.72
setstatpdch	21.74
setsys	21.75



setsysmgr	21.79
settemplate	21.82
setuser	21.84
setuseracl	21.86
setvv	21.88
setvvset	21.94

## 22 Show Commands

showalert	22.4
showauthparam	22.6
showbattery	22.8
showblock	22.14
showcage	22.17
showcim	22.23
showclienv	22.25
showcpg	22.26
showdate	22.33
showdomain	22.34
showdomainset	22.36
showeeprom	22.38
showeventlog	22.41
showfirmwaredb	22.45
showhost	22.47
showhostset	22.52
showinventory	22.54
showiscsisession	22.56
showld	22.58
showldch	22.66

showldmap	<b>22.71</b>
showlicense	<b>22.73</b>
shownet	<b>22.75</b>
shownode	<b>22.77</b>
shownodeenv	<b>22.87</b>
showpatch	<b>22.89</b>
showpd	<b>22.91</b>
showpdata	<b>22.107</b>
showpdch	<b>22.108</b>
showpdvv	<b>22.115</b>
showport	<b>22.119</b>
showportarp	<b>22.130</b>
showportdev	<b>22.132</b>
showportisns	<b>22.135</b>
showportlesb	<b>22.137</b>
showrcopy	<b>22.143</b>
showrctransport	<b>22.148</b>
showrsv	<b>22.151</b>
showsched	<b>22.153</b>
showsnmpmgr	<b>22.155</b>
showsnmppw	<b>22.157</b>
showspace	<b>22.159</b>
showspare [-used]	<b>22.164</b>
showsshkey	<b>22.167</b>
showsys	<b>22.169</b>
showsysmgr	<b>22.174</b>
showtarget	<b>22.177</b>

showtask	22.178
showtemplate	22.182
showtoc	22.183
showtocgen	22.185
showuser	22.186
showuseracl	22.188
showuserconn	22.189
showversion	22.191
showvln	22.193
showvv	22.199
showvvmap	22.215
showvvpd	22.217
showvvset	22.222

## 23 Shutdown Commands

shutdownnode	23.2
shutdownsys	23.4

## 24 Start Commands

startcim	24.2
startld	24.3
startcopy	24.4
startcopygroup	24.5
starttask	24.7
startvv	24.8

## 25 Stat Commands

statch	25.2
statcmp	25.7
statcpu	25.10

	statiscsi	25.12
	statiscsisession	25.16
	statld	25.18
	statlink	25.23
	statpd	25.25
	statport	25.33
	statrcopy	25.39
	statvlun	25.42
	statvv	25.48
<b>26</b>	<b>Stop Commands</b>	
	stopcim	26.2
	stoprcopy	26.4
	stoprcopygroup	26.6
<b>27</b>	<b>Sync Command</b>	
	syncrcopy	27.2
<b>28</b>	<b>Tune Commands</b>	
	tunealdvv	28.2
	tunepd	28.10
	tunetpvv	28.15
	tunevv	28.18
<b>29</b>	<b>Update Commands</b>	
	updatesnapspace	29.2
	updatevv	29.4
<b>30</b>	<b>Upgrade Commands</b>	
	upgradecage	30.2
	upgradePd	30.4

## **31 Wait Command**

waittask

**31.2**

## **IX Index**

## **RH Revision History**



# 1

# Introduction

---

## In this chapter

This chapter includes the following sections:

1.1 Audience	<b>1.2</b>
1.2 User Interfaces	<b>1.2</b>
1.3 Supported Operating Systems	<b>1.2</b>
1.4 Related Documentation	<b>1.3</b>
1.5 Organization	<b>1.4</b>
1.6 Typographical Conventions	<b>1.6</b>
1.7 Advisories	<b>1.6</b>

This reference describes the 3PAR® InForm® Command Line Interface (CLI) commands that are used to administer and maintain the 3PAR InServ® Storage Server.

---

## 1.1 Audience

This reference is for system and storage administrators who monitor and direct system configurations and resource allocation for 3PAR InServ storage systems.

---

## 1.2 User Interfaces

Two user interfaces are offered as part of the 3PAR InForm® operating system: the 3PAR InForm Command Line Interface (CLI) and the 3PAR InForm Management Console graphical user interface. This manual discusses the InForm CLI. For information about the InForm Management Console, refer to the *3PAR InForm OS Management Console Online Help*.

---

## 1.3 Supported Operating Systems

For a list of operating systems that support the InForm CLI, refer to the *3PAR InForm OS Configuration Matrix*.



## 1.4 Related Documentation

The following documents also provide information related to InServ Storage Servers and the InForm Operating System:

For Information About...	Read the...
Using the InForm Command Line Interface (CLI) to configure and administer InServ Storage Servers	<i>3PAR InForm OS CLI Administrator's Guide</i>
Using the InForm Management Console (IMC) graphical user interface to configure and administer InServ Storage Servers	<i>3PAR InForm OS Management Console Online Help</i>
Storage server hardware configurations, component numbering and layout, and system cabling	<i>3PAR InServ E-Class/F-Class Storage Server and Third-Party Rack Physical Planning Manual</i>
	<i>3PAR InServ S-Class/IT-Class Storage Server Physical Planning Manual</i>
Identifying storage server components and detailed alert information	<i>3PAR InForm OS Messages and Operator's Guide</i>
Using 3PAR Remote Copy	<i>3PAR Remote Copy User's Guide</i>
Using 3PAR CIM API	<i>3PAR CIM API Programming Reference</i>

---

## 1.5 Organization

This guide is organized as follows:

- [Chapter 1, \*Introduction\*](#) (this chapter), provides an overview of this reference, including information on audience, related documentation, and typographical conventions.
- [Chapter 2, \*CLI Command Syntax and Conventions\*](#), describes the standard syntax and conventions used by the InForm CLI.
- [Chapter 3, \*Commands Quick Reference\*](#), provides a list of the commands included in this reference, ordered by functionality.
- [Chapter 4, \*Add Command\*](#), provides the command used to add an SNMP manager.
- [Chapter 5, \*Admit Commands\*](#), describes the commands used to create and admit physical disks and virtual volumes into the system.
- [Chapter 6, \*Cancel Command\*](#), describes how to cancel a running task.
- [Chapter 7, \*Check Commands\*](#), presents the commands used to check the integrity of storage server resources.
- [Chapter 8, \*CLI Command\*](#), provides the general CLI command used to enter the interactive CLI shell.
- [Chapter 9, \*Compact Commands\*](#) presents the commands used to consolidate disk space.
- [Chapter 10, \*Control Commands\*](#), provides the commands used to control storage server components.
- [Chapter 11, \*Create Commands\*](#), describes the commands used to create new logical resources within the system (or the current service group) such as logical disks, hosts, and virtual volumes.
- [Chapter 12, \*Dismiss Commands\*](#), presents the commands used to remove physical disks from the system.
- [Chapter 13, \*Free Command\*](#), describes the command used to free snapshot administration and snapshot data spaces from a virtual volume.
- [Chapter 14, \*Grow Commands\*](#), describes the commands used to enlarge administration space.

- [Chapter 15, \*Hist Commands\*](#), includes the commands used to monitor existing service group resources.
- [Chapter 16, \*Locate Commands\*](#), provides the commands used to identify physical system resources.
- [Chapter 17, \*Move Commands\*](#), lists the commands used to relocate logical entities to various resources in the service group.
- [Chapter 18, \*Promote Commands\*](#), describes the commands used to copy snapshots to base volumes.
- [Chapter 19, \*Remove Commands\*](#), describes the commands used to remove logical resources (such as logical disks, hosts, and virtual volumes) from within the system or current service group.
- [Chapter 20, \*Service Commands\*](#), provides the commands used when replacing drive magazines and FCAL cards in the storage server.
- [Chapter 21, \*Set Commands\*](#), provides the commands used to set specific system parameters.
- [Chapter 22, \*Show Commands\*](#), presents commands used to display information and status for storage server hardware components.
- [Chapter 23, \*Shutdown Commands\*](#), covers the commands used to shut down entire clusters and nodes.
- [Chapter 24, \*Start Commands\*](#), provides the commands used to start 3PAR Remote Copy and the CIM service.
- [Chapter 25, \*Stat Commands\*](#), lists the commands used to display operational statistics for storage server hardware.
- [Chapter 26, \*Stop Commands\*](#), provides the commands used to stop 3PAR Remote Copy and the CIM service.
- [Chapter 27, \*Sync Command\*](#), covers the command used to synchronize Remote Copy volume groups.
- [Chapter 28, \*Tune Commands\*](#), provides the commands used to detect and rebalance physical disks with high service times.
- [Chapter 29, \*Update Commands\*](#), presents the commands used to update snapshot virtual volumes.

- [Chapter 30, Upgrade Commands](#), indicates the commands used to provide firmware upgrades to system components.
- [Chapter 31, Wait Command](#), provides information about the command for pausing a job.

This guide also contains an index and a revision history for your reference.

---

## 1.6 Typographical Conventions

The following typographical conventions are used in this guide:

Typeface	Meaning	Example
<b>ABCDabcd</b>	Used for dialog box elements such as titles and button labels.	Enter your system name in the <b>Value</b> box and click <b>OK</b> .
ABCDabcd	Used for system output and text you are to enter.	Enter <code>cli</code> at the Windows command prompt.

---

## 1.7 Advisories

To facilitate use of the InForm CLI, observe the notes and cautions used throughout this reference.



**NOTE:** Notes are reminders, tips, or suggestions that supplement the procedures included in this reference.



**CAUTION:** Cautions alert you to actions that can cause damage to equipment, software, or data.



**WARNING:** Warnings alert you to actions that can cause injury to people or irreversible damage to data or the operating system.

# 2

## CLI Command Syntax and Conventions

---

In this chapter

2.1 Syntax and Conventions	<b>2.2</b>
2.2 Syntax Rules	<b>2.3</b>
2.3 Glob-Style Pattern	<b>2.4</b>
2.4 Typical Command Layout in this Book	<b>2.5</b>
2.5 Global Options and Environment Variables	<b>2.6</b>
2.6 Exit Status	<b>2.6</b>

This chapter describes the command syntax for the CLI commands listed in this reference. General control commands that do not follow the syntax rules are also listed. In addition, glob-style patterns, as used in the CLI, are discussed.

## 2.1 Syntax and Conventions

Most CLI commands use the following syntax. Commands that do not use this syntax are listed in *General Control and Help Commands* on page 3.5.

```
cmd subcmd [options [arg...]] <spec>... [<pattern>...]
```

[Table 2-1](#) Lists all syntax elements and provides their meanings:

**Table 2-1.** CLI Command Syntax

Element	Meaning
cmd	Specifies an operation to be executed, such as <code>create</code> , <code>move</code> , or <code>show</code> .
subcmd	Specifies a subcommand. Subcommands specify actions for commands to be executed.
options	Indicates an optional command line element such as <code>histch -rw</code> .
arg	Indicates a specific variable of an option or subcommand. The argument is often used to identify a specific node, volume, or disk.
spec	Indicates a specifier used with a required command line element such as the command or option.
	Specifies that only one of the command specifiers or options separated by this character can be specified at a time.
{ }	Indicates grouped elements. Do not type the braces; type only the information inside the braces.
[ ]	Indicates optional elements. Do not type the brackets; type only the information inside the brackets.
< >	Indicates user-supplied input.
...	Indicates that a specifier or an option can be used more than once in a command.

## 2.2 Syntax Rules

The command syntax described in *Syntax and Conventions* on page 2.2 obeys the following rules:

- All command line elements are in lowercase letters, except where indicated in this reference.
- Subcommands immediately follow the commands they modify.
- Options, as indicated in this guide, are indicated by one or more letters, are preceded by a hyphen, and are enclosed in brackets (for example: `removealert [-a]`).
- Options follow subcommands on the command line and precede any specifier.
- An argument must directly follow the option or subcommand it is modifying and is required on the command line, unless otherwise specified (for example: `removealert -i <alert_ID>`).
- Multiple options and arguments on a command line are separated with a space.
- Specifiers follow options.
- User supplied input is identified by angled brackets (< >).
- Unless noted otherwise, valid character formats include alphanumeric characters, periods, dashes, and underscores. In general, the following length limits are enforced by the InForm CLI:
  - ◆ virtual volume name ≤ 31 characters
  - ◆ Thinly Provisioned Virtual Volume (TPVV) name ≤ 31 characters
  - ◆ virtual copy names ≤ 31 characters
  - ◆ logical disk name ≤ 31 characters
  - ◆ host name ≤ 31 characters
  - ◆ Common Provisioning Group (CPG) name ≤ 31 characters
  - ◆ template name ≤ 31 characters
  - ◆ domain name ≤ 31 characters
  - ◆ snapshot name ≤ 31 characters

- ◆ user name ≤ 31 characters
- ◆ system name ≤ 31 characters

---

## 2.3 Glob-Style Pattern

Several of the commands in this reference allow a pattern to be specified to refine the output of the commands. The patterns are used to match against a string. Strings are typically names such as host or virtual volume names. For example, in the `showhost` `host_name|pattern...` command, a pattern can be specified to refine the command output for a string of host names matching the specified pattern.

The `pattern` specifier, as used in the CLI commands, is specified in the form of a glob-style pattern. Glob-style matching in the CLI is implemented by Tcl. A glob-style pattern consists of the symbols in the following table.

Symbol	Explanation of Action
*	Matches any sequence of characters in a string, including a null string.
?	Matches any single character in a string.
[chars]	Matches any character in the set given by chars. A range of chars can be specified using a dash (-). For example, [a-z] represents all the ASCII characters from a through z.
\x	Matches the single character x.



The following actions are practical examples of the glob-style pattern in use:

Example	Explanation of Action
*	Shows all names.
[a]*	Matches all names beginning with the letter a.
[a-z]x*	Matches any character a-z in the first character position and "x" in the second position, followed by any character "*".



**NOTE:** Brackets ([ ]) are significant in Tcl and must be escaped using a backslash (\) or enclosed in braces ({ }). Other characters such as star (\*) are significant in most shells and must be escaped or quoted if running CLI commands from the shell.

## 2.4 Typical Command Layout in this Book

Typical CLI reference pages are formatted similarly as the examples that follow:

COMMAND

This section includes the name of the CLI command.

DESCRIPTION

This section describes the use or purpose of the command.

AUTHORITY

This section defines the user access required within the InServ Storage Server to use the command.

SUBCOMMANDS

This section indicates any necessary subcommands required to complete the use of the command.

**SYNTAX**

See [Syntax and Conventions](#) on page 2.2.

**OPTIONS**

This section lists the available options you can use in conjunction with the command for varied results. Some options and their specifiers are required, as indicated in the syntax, whereas some are provided to increase the level and functionality of your output.

**SPECIFIERS**

Specifies a suggested input required by the user.

**RESTRICTIONS**

This section includes any restrictions that must be followed in order to achieve maximum results.

**EXAMPLES**

This section lists sample output with results similar to what you should expect when running the command and its available options.

**NOTES**

Any pertinent information about the command that might help increase understanding and reliability is often provided in the Notes section.

---

## 2.5 Global Options and Environment Variables

Several options and environment variables are available at the global level. For complete information about these, refer to the *InForm OS CLI Administrator's Manual*.

---

## 2.6 Exit Status

Except where noted, the following codes are returned indicating success or failure for each individual command:

- 0 indicates that the command was successful.
- 1 indicates that the command failed.

# 3

## Commands Quick Reference

---

In this section

3.1 Overview	<b>3.2</b>
3.2 Disk Enclosure Management Commands	<b>3.2</b>
3.3 Domain Management Commands	<b>3.3</b>
3.4 Health and Alert Management Commands	<b>3.5</b>
3.5 Help and Utility Commands	<b>3.6</b>
3.6 Task Management Commands	<b>3.7</b>
3.7 LDAP Management Commands	<b>3.7</b>
3.8 Licensing Management Commands	<b>3.8</b>
3.9 Node Subsystem Management Commands	<b>3.8</b>
3.10 Performance Management Commands	<b>3.13</b>
3.11 Preserved Data Commands	<b>3.17</b>
3.12 Replication Commands	<b>3.18</b>
3.13 Service Commands	<b>3.21</b>
3.14 SNMP Agent Commands	<b>3.23</b>

3.15	CIM Server Commands	3.24
3.16	Sparing Commands	3.24
3.17	SSH Access Commands	3.25
3.18	Task Schedule Commands	3.26
3.19	User Management Commands	3.26
3.20	Volume Management Commands	3.27

---

## 3.1 Overview

The InForm CLI provides a host of commands allowing you to administer your InServ Storage Server. To facilitate your navigation through this manual, this section provides an overview of the commands grouped by functionality.

---

## 3.2 Disk Enclosure Management Commands

### 3.2.1 Drive Cage Management

Command	Description	Authority	For Details See
locatecage	Locates a particular drive cage.	Service, Service*	<a href="#">locatecage</a> on page 16.2
setcage	Sets parameters for a drive cage.	Super, Service*	<a href="#">setcage</a> on page 21.14
showcage	Displays drive cage information.	Super, Service, Edit, Browse	<a href="#">showcage</a> on page 22.17

## 3.2.2 Physical Disk Management

Command	Description	Authority	For Details See
<code>admitpd</code>	Admits one or all physical disks to enable their use.	Super, Service*	<a href="#">admitpd</a> on page 5.4
<code>checkpd</code>	Executes surface scans on physical disks.	Super, Service†	<a href="#">checkpd</a> on page 7.8
<code>controlpd</code>	Spins physical disks up or down.	Super, Service*	<a href="#">controlpd</a> on page 10.7
<code>dismisspd</code>	Dismisses one or more physical disks from use.	Super, Service*	<a href="#">dismisspd</a> on page 12.2
<code>setpd</code>	Marks physical disks as allocatable for logical disks.	Super, Service*	<a href="#">setpd</a> on page 21.53
<code>showpd</code>	Displays physical disks in the system.	Super, Edit, Service, Browse	<a href="#">showpd</a> on page 22.91

\* You need access to all domains in order to run this command.

† You need access to all domains in order to run this command.

## 3.3 Domain Management Commands

Command	Description	Authority	For Details See
<code>createdomain</code>	Shows a list of domains on the system.	Super*	<a href="#">createdomain</a> on page 11.30
<code>createdomainset</code>	Defines a new set of domains and provides the option of assigning one or more domains to that set.	Super, Edit	<a href="#">createdomainset</a> on page 11.32

Command	Description	Authority	For Details See
<code>movetodomain</code>	Moves objects from one domain to another.	Super*	<a href="#">movetodomain</a> on page 17.18
<code>removedomain</code>	Removes an existing domain from the system.	Super*	<a href="#">removedomain</a> on page 19.7
<code>removedomainset</code>	Removes a domain set or removes domains from an existing set.	Super, Edit	<a href="#">removedomainset</a> on page 19.8
<code>setdomain</code>	Sets the parameters and modifies the properties of a domain.	Super*	<a href="#">setdomain</a> on page 21.31
<code>setdomainset</code>	Sets the parameters and modifies the properties of a domain set.	Super, Edit	<a href="#">setdomainset</a> on page 21.33
<code>showdomain</code>	Displays the list of domains on a system.	Super, Service, Edit, Browse	<a href="#">showdomain</a> on page 22.34
<code>showdomainset</code>	Displays the domain sets defined on the InServ and their members.	Super, Service, Edit, Browse	<a href="#">showdomainset</a> on page 22.36

\* You need access to all domains in order to run this command.

## 3.4 Health and Alert Management Commands

### 3.4.1 Alerts

Command	Description	Authority	For Details See
removealert	Removes one or more alerts.	Super, Service *	<a href="#">removealert</a> on page 19.3
setalert	Sets the status of system alerts.	Super, Service*	<a href="#">setalert</a> on page 21.3
showalert	Displays system alerts.	Super, Service, Edit, Browse	<a href="#">showalert</a> on page 22.4

\* You need access to all domains in order to run this command.

### 3.4.2 Events

Command	Description	Authority	For Details See
removeeventlog	Removes event logs.	Super*	<a href="#">removeeventlog</a> on page 19.10
showeventlog	Displays event logs.	Super, Service, Edit, Browse	<a href="#">showeventlog</a> on page 22.41

\* You need access to all domains in order to run this command.

### 3.4.3 System Health

Command	Description	Authority	For Details See
checkhealth	Displays the status of the system hardware and software components.	Super, Service	<a href="#">checkhealth</a> on page 7.2

## 3.5 Help and Utility Commands

Command	Description	Authority	For Details See
cli	Provides a means to set up your CLI session or to enter directly into a CLI shell.	Super, Edit, Browse, Service	<a href="#">CLI Command</a> on page 8.1
clihelp	Lists all commands or details for a specified command.	Super, Edit, Browse, Service	<i>General Control and Help Commands</i> on page 3.5
cmore	Pages the output of commands.	Super, Edit, Browse, Service	<i>General Control and Help Commands</i> on page 3.5
help	Lists all commands or details for a specified command.	Super, Edit, Browse, Service	<i>General Control and Help Commands</i> on page 3.5
setclienv	Sets the CLI environment parameters.	Super, Service, Edit, Browse	<a href="#">setclienv</a> on page 21.18
showclienv	Displays the CLI environment parameters.	Super, Service, Edit, Browse	<a href="#">showclienv</a> on page 22.25



## 3.6 Task Management Commands

Command	Description	Authority	For Details See
canceltask	Cancels one or more tasks.	Super, Edit	<a href="#">canceltask</a> on page 6.2
removetask	Removes information about one or more tasks and their details.	Super, Edit	<a href="#">removetask</a> on page 19.31
showtask	Displays information about tasks.	Super, Service, Edit, Browse	<a href="#">showtask</a> on page 22.180
starttask	Executes commands with long running times.	Super, Service	<a href="#">starttask</a> on page 24.7
waittask	Asks the CLI to wait for a task to complete before proceeding.	Super, Service, Edit, Browse	<a href="#">waittask</a> on page 31.2

## 3.7 LDAP Management Commands

Command	Description	Authority	For Details See
setauthparam	Sets the authentication parameters.	Super	<a href="#">setauthparam</a> on page 21.5

Command	Description	Authority	For Details See
<code>showauthparam</code>	Shows authentication parameters and integrates the authentication and authorization features using LDAP.	Super	<a href="#">showauthparam</a> on page 22.6
<code>checkpassword</code>	Supports authentication and authorization using LDAP.	Super*	<a href="#">checkpassword</a> on page 7.6

\* You need access to all domains in order to run this command.

## 3.8 Licensing Management Commands

Command	Description	Authority	For Details See
<code>setlicense</code>	Sets the license key.	Super, Service*	<a href="#">setlicense</a> on page 21.40
<code>showlicense</code>	Displays the installed license info or key.	Super, Service, Edit, Browse	<a href="#">showlicense</a> on page 22.73

\* You need access to all domains in order to run this command.

## 3.9 Node Subsystem Management Commands

### 3.9.1 Firmware Versions

Command	Description	Authority	For Details See
<code>showfirmwaredb</code>	Displays a current database of firmware levels.	Super, Service, Edit, Browse	<a href="#">showfirmwaredb</a> on page 22.45

## 3.9.2 Node Date Information

Command	Description	Authority	For Details See
setdate	Sets the system time and date on all nodes.	Super, Service*	<a href="#">setdate</a> on page 21.27
showdate	Displays the date and time on all system nodes.	Super, Service, Edit, Browse	<a href="#">showdate</a> on page 22.33

\* You need access to all domains in order to run this command.

## 3.9.3 Node Properties

Command	Description	Authority	For Details See
setnode	Sets the properties of the node components such as the serial number of the power supply.	Super, Service*	<a href="#">setnode</a> on page 21.47
shownode	Displays an overview of the node specific properties.	Super, Service, Edit, Browse	<a href="#">shownode</a> on page 22.77
shownodeenv	Displays the node's environmental status.	Super, Service, Edit, Browse	<a href="#">shownodeenv</a> on page 22.87

\* You need access to all domains in order to run this command.

## 3.9.4 Node EEPROM Log

Command	Description	Authority	For Details See
showeeprom	Displays node EEPROM information.	Super, Service, Edit, Browse	<a href="#">showeeprom</a> on page 22.38

### 3.9.5 Array and Node Information

Command	Description	Authority	For Details See
<code>locatesys</code>	Locates a system by blinking its LEDs.	Super, Service*	<a href="#">locatesys</a> on page 16.4
<code>setsys</code>	Enables you to set system-wide parameters such as the raw space alert.	Super, Service*	<a href="#">setsys</a> on page 21.75
<code>showsys</code>	Displays the InServ system properties, including system name, model, serial number, and system capacity.	Super, Service, Edit, Browse	<a href="#">showsys</a> on page 22.171

\* You need access to all domains in order to run this command.

### 3.9.6 Network Interface Configuration

Command	Description	Authority	For Details See
<code>setnet</code>	Sets the administration network interface configuration.	Super, Service*	<a href="#">setnet</a> on page 21.42
<code>setntp</code>	Sets the NTP server to which the InServ Storage Server synchronizes.	Super, Service*	<a href="#">setntp</a> on page 21.49
<code>shownet</code>	Displays the network configuration and status.	Super, Service, Edit, Browse	<a href="#">shownet</a> on page 22.75

\* You need access to all domains in order to run this command.

### 3.9.7 Port Information

Command	Description	Authority	For Details See
checkport	Performs a loopback test on Fibre Channel ports.	Super, Service*	<a href="#">checkport</a> on page 7.11
controlport	Controls Fibre Channel or Remote Copy ports.	Super, Service*	<a href="#">controlport</a> on page 10.9
controliscsiport	Used to set up the parameters and characteristics of an iSCSI port.	Super, Service*	<a href="#">controliscsiport</a> on page 10.2
showiscsisession	Shows the iSCSI active sessions per port.	Super, Service, Edit, Browse	<a href="#">showiscsisession</a> on page 22.56
showport	Displays system port information.	Super, Service, Edit, Browse	<a href="#">showport</a> on page 22.121
showportarp	Shows the ARP table for iSCSI ports in the system.	Super, Service, Edit, Browse	<a href="#">showportarp</a> on page 22.132
showportdev	Displays detailed information about devices on a Fibre Channel port.	Super, Service, Edit, Browse	<a href="#">showportdev</a> on page 22.134
showportisns	Show iSNS host information for iSCSI ports in the system.	Super, Service, Edit, Browse	<a href="#">showportisns</a> on page 22.137

Command	Description	Authority	For Details See
showportlesb	Displays Link Error Status Block information about devices on a Fibre Channel port.	Super, Service, Edit, Browse	<a href="#">showportlesb</a> on page 22.139
showtarget	Displays unrecognized targets.	Super, Service, Edit, Browse	<a href="#">showtarget</a> on page 22.179
statiscsi	Displays the iSCSI statistics.	Super, Service, Edit, Browse	<a href="#">statiscsi</a> on page 25.12
statiscsisession	Displays the iSCSI session statistics.	Super, Service, Edit, Browse	<a href="#">statiscsisession</a> on page 25.16

\* You need access to all domains in order to run this command.

### 3.9.8 Battery Management

Command	Description	Authority	For Details See
setbattery	Sets battery properties.	Super, Service *	<a href="#">setbattery</a> on page 21.12
showbattery	Displays battery status information.	Super, Service, Edit, Browse	<a href="#">showbattery</a> on page 22.8

\* You need access to all domains in order to run this command.

## 3.9.9 System Manager

Command	Description	Authority	For Details See
<code>setsysmgr</code>	Sets the system manager startup state.	Super*	<a href="#">setsysmgr</a> on page 21.79
<code>showsysmgr</code>	Displays the system manager startup state.	Super, Service, Edit, Browse	<a href="#">showsysmgr</a> on page 22.176
<code>showtoc</code>	Displays the system table of contents summary.	Super, Service, Edit, Browse	<a href="#">showtoc</a> on page 22.185
<code>showtocgen</code>	Displays the system table of contents generation number.	Super, Service, Edit, Browse	<a href="#">showtocgen</a> on page 22.187

\* You need access to all domains in order to run this command.

## 3.10 Performance Management Commands

### 3.10.1 Chunklet Statistics

Command	Description	Authority	For Details See
<code>histch</code>	Displays histogram data for individual chunklets.	Super, Service, Edit, Browse	<a href="#">histch</a> on page 15.2
<code>setstatch</code>	Sets statistics collection mode on chunklets.	Super, Edit	<a href="#">setstatch</a> on page 21.72

Command	Description	Authority	For Details See
setstatpdch	Sets statistics collection mode on physical disk chunklets.	Super, Edit	<a href="#">setstatpdch</a> on page 21.74
statch	Displays statistics for individual chunklets.	Super, Service, Edit, Browse	<a href="#">statch</a> on page 25.2

### 3.10.2 Data Cache Memory Statistics

Command	Description	Authority	For Details See
statcmp	Displays statistics for cache memory pages.	Super, Service, Edit, Browse	<a href="#">statcmp</a> on page 25.7

### 3.10.3 Node CPU Statistics

Command	Description	Authority	For Details See
statcpu	Displays statistics for CPU use.	Super, Service, Edit, Browse	<a href="#">statcpu</a> on page 25.10

### 3.10.4 Logical Disk Statistics

Command	Description	Authority	For Details See
histld	Displays histogram data for logical disks.	Super, Service, Edit, Browse	<a href="#">histld</a> on page 15.6
statld	Displays statistics for logical disks.	Super, Service, Edit, Browse	<a href="#">statld</a> on page 25.18



### 3.10.5 Link Statistics

Command	Description	Authority	For Details See
statlink	Displays statistics for links.	Super, Service, Edit, Browse	<a href="#">statlink</a> on page 25.23

### 3.10.6 Physical Disk Statistics

Command	Description	Authority	For Details See
histpd	Displays histogram data for physical disks.	Super, Service, Edit, Browse	<a href="#">histpd</a> on page 15.10
statpd	Displays statistics for physical disks.	Super, Service, Edit, Browse	<a href="#">statpd</a> on page 25.25

### 3.10.7 Port Statistics

Command	Description	Authority	For Details See
histport	Displays histogram data for Fibre Channel ports.	Super, Service, Edit, Browse	<a href="#">histport</a> on page 15.18
statport	Displays statistics for Fibre Channel ports.	Super, Service, Edit, Browse	<a href="#">statport</a> on page 25.33

### 3.10.8 System Tuner

Command	Description	Authority	For Details See
tunepd	Displays physical disks with high service times and optionally performs load balancing.	Super, Edit	<a href="#">tunepd</a> on page 28.10

### 3.10.9 Dynamic Optimization

Command	Description	Authority	For Details See
tunealdevv	Allows the RAID and Availability characteristics of an existing Thin Provisioned Virtual Volume to be dynamically modified. See the InForm OS Administrator's Guide for a complete discussion of 3PAR System Tuner and the use of the tunealdevv command.	Super, Edit*	<a href="#">tunealdevv</a> on page 28.2
tunetpvv	Changes the layout of a Thinly Provisioned Virtual Volume (TPVV).	Super, Edit	<a href="#">tunetpvv</a> on page 28.15
tunevv	Changes the layout of a virtual volume.	Super, Edit	<a href="#">tunevv</a> on page 28.18

\* You need access to all domains in order to run this command.

### 3.10.10 Virtual LUN (Export) Statistics

Command	Description	Authority	For Details See
histvln	Displays histogram data for VLUNs.	Super, Service, Edit, Browse	<a href="#">histvln</a> on page 15.23
statvln	Displays statistics for VLUNs.	Super, Service, Edit, Browse	<a href="#">statvln</a> on page 25.42

### 3.10.11 Virtual Volume Statistics

Command	Description	Authority	For Details See
histvv	Displays histogram data for virtual volumes.	Super, Service, Edit, Browse	<a href="#">histvv</a> on page 15.28
statvv	Displays statistics for virtual volumes.	Super, Service, Edit, Browse	<a href="#">statvv</a> on page 25.48

## 3.11 Preserved Data Commands

Command	Description	Authority	For Details See
showpdata	Displays preserved data status.	Super, Service, Edit, Browse	<a href="#">showpdata</a> on page 22.108

## 3.12 Replication Commands

### 3.12.1 Physical Copy

Command	Description	Authority	For Details See
<code>creategroupvvcopy</code>	Creates consistent group physical copies of a list of virtual volumes.	Super, Edit	<a href="#">creategroupvvcopy</a> on page 11.37
<code>createvvcopy</code>	Copies a virtual volume.	Super, Edit	<a href="#">createvvcopy</a> on page 11.90
<code>promotevvcopy</code>	Promotes a physical copy back to a base volume.	Super, Edit	<a href="#">promotevvcopy</a> on page 18.4

### 3.12.2 Remote Copy



**NOTE:** Functionality of 3PAR Remote Copy requires a 3PAR Remote Copy license. See the 3PAR InForm OS Concepts Guide, Chapter 2 for additional information.

Command	Description	Authority	For Details See
<code>admitrcopylink</code>	Admits a network link for Remote Copy use.	Super, Edit	<a href="#">admitrcopylink</a> on page 5.6
<code>admitrcopytarget</code>	Adds a target to a Remote Copy volume group	Super, Edit	<a href="#">admitrcopytarget</a> on page 5.8
<code>admitrcopyvv</code>	Admits a virtual volume to a Remote Copy volume group.	Super, Edit	<a href="#">admitrcopyvv</a> on page 5.10

Command	Description	Authority	For Details See
<code>creatercopygroup</code>	Creates a group for Remote Copy.	Super, Edit	<a href="#">creatercopygroup</a> on page 11.48
<code>creatercopytarget</code>	Creates a target for Remote Copy.	Super, Edit*	<a href="#">creatercopytarget</a> on page 11.50
<code>dismissrcopylink</code>	Dismisses a network link from Remote Copy use.	Super, Edit*	<a href="#">dismissrcopylink</a> on page 12.3
<code>dismissrcopytarget</code>	Dismisses a Remote Copy target from a Remote Copy volume group.	Super, Edit	<a href="#">dismissrcopytarget</a> on page 12.5
<code>dismissrcopyvv</code>	Dismisses a virtual volume from a Remote Copy volume group.	Super, Edit	<a href="#">dismissrcopyvv</a> on page 12.6
<code>removercopygroup</code>	Removes a group used for Remote Copy.	Super, Edit	<a href="#">removercopygroup</a> on page 19.18
<code>removercopytarget</code>	Removes a target used for Remote Copy.	Super, Edit*	<a href="#">removercopytarget</a> on page 19.20
<code>setrcopygroup</code>	Sets the volume group's policy for dealing with I/O failure and error handling, or switches the direction of a volume group.	Super, Edit	<a href="#">setrcopygroup</a> on page 21.55

Command	Description	Authority	For Details See
setrcopytarget	Sets the Remote Copy target state.	Super, Edit*	<a href="#">setrcopytarget</a> on page 21.62
showrcopy	Displays the details of a Remote Copy configuration.	Super, Service, Edit, Browse	<a href="#">showrcopy</a> on page 22.145
showrctransport	Shows status and information about end-to-end transport for Remote Copy in the system.	Super, Service, Edit, Browse	<a href="#">showrctransport</a> on page 22.150
startrcopy	Starts a Remote Copy subsystem.	Super, Edit*	<a href="#">startrcopy</a> on page 24.4
startrcopygroup	Starts a Remote Copy volume group.	Super, Edit	<a href="#">startrcopygroup</a> on page 24.5
statrcopy	Displays Remote Copy statistics.	Super, Service, Edit, Browse	<a href="#">statrcopy</a> on page 25.39
stoprcopy	Stops a Remote Copy subsystem.	Super, Edit*	<a href="#">stoprcopy</a> on page 26.4
stoprcopygroup	Stops a Remote Copy volume group.	Super, Edit	<a href="#">stoprcopygroup</a> on page 26.6
syncrcopy	Synchronizes Remote Copy volume groups.	Super, Edit	<a href="#">syncrcopy</a> on page 27.2

## 3.12.3 Virtual Copy

Command	Description	Authority	For Details See
<code>createsv</code>	Creates snapshot volumes.	Super, Edit	<a href="#">createsv</a> on page 11.58
<code>creategroupsv</code>	Creates consistent group snapshots of a list of virtual volumes.	Super, Edit	<a href="#">creategroupsv</a> on page 11.34
<code>promotesv</code>	Copies the differences of a virtual copy back to its base volume.	Super, Edit	<a href="#">promotesv</a> on page 18.2
<code>updatevv</code>	Updates a snapshot virtual volume with a new snapshot.	Super, Edit	<a href="#">updatevv</a> on page 29.4

## 3.13 Service Commands

### 3.13.1 Disk Enclosure

Command	Description	Authority	For Details See
<code>admithw</code>	Admits new hardware into the system.	Super, Service <sup>*</sup>	<a href="#">admithw</a> on page 5.2
<code>controlmag</code>	Takes drives or magazines on or off loop.	Super, Service <sup>†</sup>	<a href="#">controlmag</a> on page 10.5
<code>servicecage</code>	Prepares a drive cage for service.	Super, Service <sup>*</sup>	<a href="#">servicecage</a> on page 20.2
<code>servicehost</code>	Prepares a port for host attachment.	Super, Service <sup>*</sup>	<a href="#">servicehost</a> on page 20.5

Command	Description	Authority	For Details See
<code>servicemag</code>	Prepares a drive magazine for service.	Super, Service*	<a href="#">servicemag</a> on page 20.8
<code>upgradecage</code>	Upgrades drive cage firmware.	Super, Service*	<a href="#">upgradecage</a> on page 30.2
<code>upgradedpd</code>	Upgrades disk firmware.	Super, Service*	<a href="#">upgradedpd</a> on page 30.4

\* You need access to all domains in order to run this command.

† You need access to all domains in order to run this command.

### 3.13.2 General System Maintenance

Command	Description	Authority	For Details See
<code>shutdownnode</code>	Shuts down an individual system node.	Super, Service*	<a href="#">shutdownnode</a> on page 23.2
<code>shutdownsys</code>	Shuts down the entire system.	Super, Service*	<a href="#">shutdownsys</a> on page 23.4

### 3.13.3 System Upgrade

Command	Description	Authority	For Details See
<code>showpatch</code>	Displays patches applied to a system.	Super, Service, Edit, Browse	<a href="#">showpatch</a> on page 22.89
<code>showversion</code>	Displays software versions.	Super, Service, Edit, Browse	<a href="#">showversion</a> on page 22.193



## 3.14 SNMP Agent Commands

Command	Description	Authority	For Details See
<code>addsnmpmgr</code>	Adds an SNMP manager to receive trap notifications.	Super*	<a href="#">addsnmpmgr</a> on page 4.2
<code>removesnmpmgr</code>	Removes an SNMP trap manager.	Super*	<a href="#">removesnmpmgr</a> on page 19.24
<code>removesnmppw</code>	Removes an SNMP password.	Super*	<a href="#">removesnmppw</a> on page 19.26
<code>setsnmppw</code>	Allows users to update SNMP passwords.	Super*	<a href="#">setsnmppw</a> on page 21.68
<code>showsnmpmgr</code>	Displays SNMP trap managers.	Super, Service, Edit, Browse	<a href="#">showsnmpmgr</a> on page 22.157
<code>showsnmppw</code>	Displays SNMP access passwords.	Super, Edit, Browse	<a href="#">showsnmppw</a> on page 22.159

\* You need access to all domains in order to run this command.

## 3.15 CIM Server Commands

Command	Description	Authority	For Details See
setcim	Sets the properties of the CIM server, including options to enable or disable the SLP, HTTP and HTTPS ports for the CIM server.	Super, Service*	<a href="#">setcim</a> on page 21.16
showcim	Displays the CIM server setting information and status.	Super, Service, Edit, Browse	<a href="#">showcim</a> on page 22.23
startcim	Starts the CIM server to service CIM requests.	Super, Service*	<a href="#">startcim</a> on page 24.2
stopcim	Stops the CIM server from servicing CIM requests.	Super, Service*	<a href="#">stopcim</a> on page 26.2

\* You need access to all domains in order to run this command.

## 3.16 Sparing Commands

Command	Description	Authority	For Details See
createspare	Creates spare chunklets.	Super, Service*	<a href="#">createspare</a> on page 11.56
movech	Moves specified chunklets.	Super, Service, Edit*	<a href="#">movech</a> on page 17.2
movechtospare	Moves specified chunklets to spare.	Super, Service, Edit*	<a href="#">movetodomain</a> on page 17.18

Command	Description	Authority	For Details See
<code>movepdtospare</code>	Moves specified physical disks to spare.	Super, Service, Edit*	<a href="#">movepdtospare</a> on page 17.9
<code>moverelocpd</code>	Moves chunklets relocated from a physical disk to another physical disk.	Super, Service, Edit*	<a href="#">moverelocpd</a> on page 17.12
<code>removespare</code>	Removes spare chunklets.	Edit, Service*	<a href="#">removespare</a> on page 19.28
<code>showspare</code>	Displays information about spare and relocated chunklets.	Super, Service, Edit, Browse	<a href="#">showspare [-used]</a> on page 22.166

\* You need access to all domains in order to run this command.

## 3.17 SSH Access Commands

Command	Description	Authority	For Details See
<code>setsshkey</code>	Sets the SSH public key for users enabling login without a password.	Super, Service, Edit, Browse	<a href="#">setsshkey</a> on page 21.70
<code>showsshkey</code>	Displays all SSH public keys that have been set with <code>setsshkey</code> .	Super, Service, Edit, Browse	<a href="#">showsshkey</a> on page 22.169
<code>removesshkey</code>	Removes a user's SSH public key.	Super, Service, Edit, Browse	<a href="#">removesshkey</a> on page 19.30

## 3.18 Task Schedule Commands

Command	Description	Authority	For Details See
<code>createsched</code>	Allows users to schedule tasks that are periodically run by the scheduler.	Super, Service	<a href="#">createsched</a> on page 11.52
<code>removesched</code>	Removes a scheduled task from the system.	Super, Service	<a href="#">removesched</a> on page 19.22
<code>setsched</code>	Allows users to suspend, pause, change the schedule, change the parameters, and change the name of currently scheduled tasks.	Super, Service	<a href="#">setsched</a> on page 21.66
<code>showsched</code>	Displays the state of tasks currently scheduled on the system.	Super, Service	<a href="#">showsched</a> on page 22.155

## 3.19 User Management Commands

Command	Description	Authority	For Details See
<code>createuser</code>	Creates user accounts.	Super*	<a href="#">createuser</a> on page 11.77
<code>removeuser</code>	Removes user accounts.	Super*	<a href="#">removeuser</a> on page 19.35
<code>removeuserconn</code>	Removes user connections.	Super*	<a href="#">removeuserconn</a> on page 19.37

Command	Description	Authority	For Details See
setpassword	Changes your password.	Super, Edit, Service Browse <sup>†</sup>	<a href="#">setpassword</a> on page 21.51
setuser	Sets your user properties.	Super, Edit, Browse	<a href="#">setuser</a> on page 21.84
setuseracl	Sets your Access Control List (ACL).	Super, Edit	<a href="#">setuseracl</a> on page 21.86
showuser	Displays user accounts.	Super, Service, Edit, Browse	<a href="#">showuser</a> on page 22.188
showuseracl	Displays your access control list (ACL).	Super, Service, Edit, Browse	<a href="#">showuseracl</a> on page 22.190
showuserconn	Displays user connections.	Super, Service, Edit, Browse	<a href="#">showuserconn</a> on page 22.191

\* You need access to all domains in order to run this command.

† Only the Super User can edit other user's passwords.

## 3.20 Volume Management Commands

### 3.20.1 Common Provisioning Group Management

Command	Description	Authority	For Details See
compactcpg	Consolidates logical disk space in a CPG into as few logical disks as possible, allowing unused logical disks to be removed.	Super, Edit	<a href="#">compactcpg</a> on page 9.2
createcpg	Creates a Common Provisioning Group (CPG).	Super, Edit*	<a href="#">createcpg</a> on page 11.22

Command	Description	Authority	For Details See
<code>removecpg</code>	Removes CPGs.	Super, Edit*	<a href="#">removecpg</a> on page 19.5
<code>setcpg</code>	Changes the properties CPGs.	Super, Edit*	<a href="#">setcpg</a> on page 21.20
<code>showcpg</code>	Displays CPGs.	Super, Service, Edit, Browse	<a href="#">showcpg</a> on page 22.26

\* You need access to all domains in order to run this command.

## 3.20.2 Host Management

Command	Description	Authority	For Details See
<code>createhost</code>	Creates host and host path definitions.	Super, Edit*	<a href="#">createhost</a> on page 11.40
<code>createhostset</code>	Creates a new set of hosts and provides the option of assigning one or more existing hosts to that set.	Super, Edit	<a href="#">createhostset</a> on page 11.43
<code>removehost</code>	Removes host definitions from the system.	Super, Edit*	<a href="#">removehost</a> on page 19.12
<code>removehostset</code>	Removes a host set or removes hosts from an existing set.	Super, Edit	<a href="#">removehostset</a> on page 19.14
<code>showhost</code>	Displays defined hosts in the system.	Super, Service, Edit, Browse	<a href="#">showhost</a> on page 22.47
<code>showhostset</code>	Displays the host sets defined on the InServ and their members.	Super, Service, Edit, Browse	<a href="#">showhostset</a> on page 22.52

Command	Description	Authority	For Details See
<code>sethost</code>	Sets properties on existing system hosts, including options to annotate a host with descriptor information such as physical location, IP address, operating system, model, and so on.	Super, Edit*	<a href="#">sethost</a> on page 21.35
<code>sethostset</code>	Sets the parameters and modifies the properties of a host set.	Super, Edit	<a href="#">sethostset</a> on page 21.39

\* You need access to all domains in order to run this command.

### 3.20.3 Logical Disk Management

Command	Description	Authority	For Details See
<code>checkld</code>	Performs validity checks of data on logical disks.	Super, Service*	<a href="#">checkld</a> on page 7.4
<code>compactld</code>	Consolidates space on the logical disks.	Super, Edit	<a href="#">compactld</a> on page 9.4
<code>createald</code>	Automatically creates logical disks.	Super, Edit*	<a href="#">createald</a> on page 11.3
<code>createld</code>	Creates logical disks.	Super, Edit*	<a href="#">createld</a> on page 11.45
<code>removeld</code>	Removes logical disks.	Super, Service, Edit	<a href="#">removeld</a> on page 19.16

Command	Description	Authority	For Details See
<code>showld</code>	Displays logical disks.	Super, Service, Edit, Browse	<a href="#">showld</a> on page 22.58
<code>startld</code>	Starts logical disks.	Super, Service*	<a href="#">startld</a> on page 24.3

\* You need access to all domains in order to run this command.

### 3.20.4 Space and Storage Management

Command	Description	Authority	For Details See
<code>showblock</code>	Displays block mapping information for virtual volumes, logical disks, and physical disks.	Super, Service, Edit, Browse	<a href="#">showblock</a> on page 22.14
<code>showldch</code>	Displays logical disk to physical disk chunklet mapping.	Super, Service, Edit, Browse	<a href="#">showldch</a> on page 22.66
<code>showldmap</code>	Displays logical disk to virtual volume mapping.	Super, Service, Edit, Browse	<a href="#">showldmap</a> on page 22.71
<code>showpdch</code>	Displays the status of selected chunklets of physical disks.	Super, Service, Edit, Browse	<a href="#">showpdch</a> on page 22.109
<code>showpdvv</code>	Displays physical disk to virtual volume mapping.	Super, Service, Edit, Browse	<a href="#">showpdvv</a> on page 22.116
<code>showspace</code>	Displays estimated free space.	Super, Service, Edit, Browse	<a href="#">showspace</a> on page 22.161



Command	Description	Authority	For Details See
<code>showvmap</code>	Displays virtual volume to logical disk mapping.	Super, Service, Edit, Browse	<a href="#">showvmap</a> on page 22.217
<code>showvvpd</code>	Displays virtual volume distribution across physical disks.	Super, Service, Edit, Browse	<a href="#">showvvpd</a> on page 22.219

### 3.20.5 Template Management

Command	Description	Authority	For Details See
<code>createtemplate</code>	Creates templates for the creation of logical disks, virtual volumes, thinly provisioned virtual volumes, and common provisioning groups.	Super <sup>*</sup>	<a href="#">createtemplate</a> on page 11.61
<code>removetemplate</code>	Removes one or more templates.	Super <sup>*</sup>	<a href="#">removetemplate</a> on page 19.33
<code>settemplate</code>	Modifies template properties.	Super <sup>*</sup>	<a href="#">settemplate</a> on page 21.82
<code>showtemplate</code>	Displays existing templates.	Super, Service, Edit, Browse	<a href="#">showtemplate</a> on page 22.184

<sup>\*</sup> You need access to all domains in order to run this command.

### 3.20.6 Virtual Volume Management

Command	Description	Authority	For Details See
checkvv	Performs validity checks of virtual volume administrative information.	Super, Service*	<a href="#">checkvv</a> on page 7.13
createaldvv	Automatically creates virtual volumes and their underlying logical disks.	Super, Edit*	<a href="#">createaldvv</a> on page 11.9
createavv	Automatically creates virtual volumes.	Super, Edit*	<a href="#">createavv</a> on page 11.18
createtpvv	Creates a TPVV.	Super, Edit	<a href="#">createtpvv</a> on page 11.71
createvv	Creates a virtual volume from logical disks.	Super, Edit*	<a href="#">createvv</a> on page 11.83
createvvset	Defines a new set of virtual volumes provides the option of assigning one or more existing virtual volumes to that set.	Super, Edit	<a href="#">createvvset</a> on page 11.94
freespace	Frees SA and SD spaces from a virtual volume if they are not in use.	Super, Edit	<a href="#">freespace</a> on page 13.2

Command	Description	Authority	For Details See
growaldivv	Automatically increases the size of a virtual volume and its underlying logical disks.	Super, Edit*	<a href="#">growaldivv</a> on page 14.2
growavv	Automatically increases the size of a virtual volume.	Super, Edit*	<a href="#">growavv</a> on page 14.10
growtpvv	Enlarges a thin provisioning virtual volume.	Super, Edit	<a href="#">growtpvv</a> on page 14.12
growvv	Increases the size of a virtual volume by adding logical disks.	Super, Edit*	<a href="#">growvv</a> on page 14.14
removevv	Removes virtual volumes or logical disks from common provisioning groups.	Super, Edit	<a href="#">removevv</a> on page 19.43
removevvset	Removes a virtual volume set or virtual volumes from an existing set.	Super, Edit	<a href="#">removevvset</a> on page 19.46
setvv	Modifies properties associated with a virtual volume.	Super, Edit	<a href="#">setvv</a> on page 21.88
setvvset	Sets the parameters and modifies the properties of a virtual volume set.	Super, Edit	<a href="#">setvvset</a> on page 21.94

Command	Description	Authority	For Details See
<code>showrsv</code>	Displays information about reservation and registration of VLUNs connected on a Fibre Channel port.	Super, Service, Edit, Browse	<a href="#">showrsv</a> on page 22.153
<code>showvv</code>	Displays virtual volumes in the system.	Super, Service, Edit, Browse	<a href="#">showvv</a> on page 22.201
<code>showvvset</code>	Displays the virtual volume sets defined on the InServ and their members.	Super, Service, Edit, Browse	<a href="#">showvvset</a> on page 22.224
<code>startvv</code>	Starts virtual volumes.	Super, Service*	<a href="#">startvv</a> on page 24.8
<code>updatesnapspace</code>	Starts a task to update the actual snapshot space used by a virtual volume.	Super, Edit	<a href="#">updatesnapspace</a> on page 29.2

\* You need access to all domains in order to run this command.

### 3.20.7 Virtual LUN (Export) Management

Command	Description	Authority	For Details See
<code>createvln</code>	Creates a virtual volume as a SCSI LUN.	Super, Edit	<a href="#">createvln</a> on page 11.79
<code>removevln</code>	Removes VLUNs.	Super, Edit	<a href="#">removevln</a> on page 19.39
<code>showvln</code>	Displays VLUNs in the system.	Super, Service, Edit, Browse	<a href="#">showvln</a> on page 22.195

# 4

## Add Command

---

In this chapter

`addsnmpmgr`

**4.2**

---

**COMMAND**

`addsnmpmgr`

**DESCRIPTION**

The `addsnmpmgr` command adds an SNMP manager to receive alert (traps) notifications.

**SYNTAX**

`addsnmpmgr [options <arg>] <manager_IP>`

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-p <port_number>`

Specifies the port number where the SNMP manager receives traps. This option is used if the port number differs from the default of 162.

`-pw <password>`

Specifies the SNMP manager's access community string (password), using up to 32 alphanumeric characters. If unspecified, the agents send traps without an access password.

`-r <retry>`

Specifies the number of times to send a trap (`<retry>`) using an integer from 1 through 15 if the SNMP manager is not available. If not specified, the number of times a trap is sent defaults to 2.

`-t <timeout>`

Specifies the number of seconds to wait before sending a trap (`timeout`) using an integer from 1 through 300. If not specified, the time defaults to 200 seconds.

## SPECIFIERS

<manager\_IP>

Specifies the IP address of the host where the manager runs. It must be a valid IPv4 or IPv6 address. The IPv6 address is in hexadecimal, is case insensitive, and is separated by colons.

For example:

```
5de:2008:0:0abcd:0:0:161a
```

In addition, a double colon (: :) can be used once in an address to replace multiple fields of zeros. For example:

```
5de:2008:0:0abcd::161a.
```

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the addition of SNMP manager IPv4 address 123.45.67.89 with the assigned password of alpha1:

```
cli% addsnmpmgr -pw alpha1 123.45.67.89
```

The following example displays the addition of SNMP manager IPv6 address 5def:2008:abcd::161a with port number 9162:

```
cli% addsnmpmgr -p 9162 5def:2008:abcd::161a
```

## NOTES

- The InServ Storage Server does not support any form of name resolution. You must specify these IP addresses directly.
- Issue the `showsnmpmgr` command to display the list of registered SNMP managers.
- Issue the `setsnmpw` command to change the SNMP passwords.
- Issue the `removesnmpw` command to remove SNMP passwords.
- Issue the `removesnmpmgr` command to remove SNMP managers. See [removesnmpmgr](#) on page 19.24 for additional information.





# 5

## Admit Commands

---

In this chapter

<code>admithw</code>	<b>5.2</b>
<code>admitpd</code>	<b>5.4</b>
<code>admitrcopylink</code>	<b>5.6</b>
<code>admitrcopytarget</code>	<b>5.8</b>
<code>admitrcopyvv</code>	<b>5.10</b>

---

**COMMAND**

`admithw`

**DESCRIPTION**

The `admithw` command admits new hardware into the system.

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

`admithw [options]`

**OPTIONS**

`-checkonly`

Only performs passive checks; does not make any changes.

`-f`

If errors are encountered, the `admithw` command ignores them and continues. The messages remain displayed.

**SPECIFIERS**

None.

**RESTRICTIONS**

Access to all domains is required to run this command.

## EXAMPLES

The following example displays a passive check of the system:

```
cli% admithw -checkonly
Checking nodes...

Checking volumes...

Checking system LDs...

Checking ports...

Checking state of disks...

Checking cabling...

Check complete.
```

## NOTES

- The `admithw` command handles any nodes, disks, or cages that have been added into the system.
- In addition to verifying that all expected hardware is present, the `admithw` command handles all checks, including valid states, cabling, and firmware revisions.
- The `admithw` command also handles creating system logical disks while adding and rebalancing spare chunklets.
- Spares are allocated according to the algorithm specified by the Sparing Algorithm system parameter.

---

**COMMAND**`admitpd`**DESCRIPTION**

The `admitpd` command creates and admits physical disk definitions to enable the use of those disks.

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

```
admitpd [option] [<WWN>...]
```

**OPTIONS**

`-nold`

Do not use the physical disk (as identified by the WWN specifier) for logical disk allocation. See [Notes](#) on page 5.5 for more information about the `-nold` option.

**SPECIFIERS**

[<WWN>...]

Indicates the World-Wide Name (WWN) of the physical disk to be admitted. If WWNs are specified, only the specified physical disk(s) are admitted. Otherwise, all available physical disks are admitted.

**RESTRICTIONS**

- Access to all domains is required to run this command.
- If multiple WWNs are specified and not all can be admitted, the `admitpd` command fails.

**EXAMPLES**

The following example admits physical disks in a 20-disk system:

```
cli% admitpd
20 disks admitted
```

## NOTES

- Physical disks cannot be used by the InForm operating system for storage until they are admitted into the system.
- Specify the `-nold` option when adding a physical disk to replace a failed disk, whose chunklets were moved to spare space. Specifying `-nold` prevents the allocation of the newly added physical disk and allows chunklets to be moved back to the new disk. After chunklets have been moved back to the new disk, the administrator can allow logical disks to be allocated again by issuing the `setpd` command.
- Verify the admittance of physical disk definitions by issuing the `showpd` command. See [showpd](#) on page 22.91.
- If no WWN is specified or if all the specified WWNs are admitted, the command succeeds. If all the specified WWNs could not be admitted, the command fails.

---

## COMMAND

`admitrcopylink`

## DESCRIPTION

The `admitrcopylink` command adds one or more links (connections) to a Remote Copy target system.

## SYNTAX

- For Remote Copy over IP (RCIP), the syntax for the `admitrcopylink` command is as follows:

```
admitrcopylink <target_name> <N:S:P:IP_address>...
```

- For Remote Copy over Fibre Channel (RCFC), the syntax for the `admitrcopylink` command is as follows:

```
admitrcopylink <target_name> <N:S:P:WWN>...
```

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

None.

## SPECIFIERS

`<target_name>`

The target name, as specified with the `creatercopytarget` command (see [creatercopytarget](#) on page 11.50).

`<N:S:P:IP_address>...`

Specifies the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.

`<N:S:P:WWN>...`

Specifies the node, slot, and port of the Fibre Channel port on the local system and a World Wide Name (WWN) address on the target system.

## RESTRICTIONS

- Functionality of this command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- Access to all domains is required to run this command.

## EXAMPLES

The following example adds a link on `System2` (targetname), node `1`. The IP address `193.1.2.11` specifies the address on the target system:

```
cli% admitrcopylink System2 1:2:1:193.1.2.11
```

## NOTES

- See the *Remote Copy User's Guide* for more examples.
- This command concludes by returning a list of one or more links to be admitted.
- IP targets are made up of pairs composed of the node containing the Ethernet port on the local systems and an IP address on the target system.
- FC targets are made up of sets with the node, slot, and port of the fibre channel port on the local system and a WWN on the target system.

---

## COMMAND

`admitrcopytarget`

## DESCRIPTION

The `admitrcopytarget` command adds a target to a Remote Copy volume group.

## SYNTAX

`admitrcopytarget <target_name> <mode> <group_name>`

`[ <pri_VV_name>:<sec_VV_name> ] ...`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

None.

## SPECIFIERS

`<target_name>`

Specifies the name of the target that was previously created with the `creatercopytarget` command.

`<mode>`

Specifies the mode of the target as either synchronous (`sync`) or asynchronous periodic (`periodic`).

`<group_name>`

Specifies the name of the existing Remote Copy volume group created with the `creatercopygroup` command to which the target will be added.

`[ <pri_VV_name>:<sec_VV_name> ] ...`

Specifies the names of the primary and secondary volume groups to which the specified target is added. This specifier is not required.



## RESTRICTIONS

Functionality of this command requires the 3PAR Remote Copy license. Contact your local service provider for further information.

## EXAMPLES

In the following example, the `admitrcopytarget` command adds the target `target1` in synchronous mode to volume group `Group1`.

```
cli% admitrcopytarget target1 sync Group1
```

## NOTES

A primary to secondary volume mapping must be provided for each volume currently in the group.

---

**COMMAND**

admitrcopyvv

**DESCRIPTION**

The `admitrcopyvv` command adds an existing virtual volume to an existing Remote Copy volume group.

**SYNTAX**

```
admitrcopyvv <VV_name> <group_name> <target_name>:<sec_VV_name>
               [<target_name>:<sec_VV_name> ...]
```

**AUTHORITY**

Super, Edit

**OPTIONS**

None.

**SPECIFIERS**

<VV\_name>

Specifies the name of the existing virtual volume to be admitted to an existing Remote Copy volume group that was created with the `creatercopygroup` command (see [page 11.48](#)).

<group\_name>

Specifies the name of the existing Remote Copy volume group created with the `creatercopygroup` command (see [page 11.48](#)), to which the volume will be added.

<target\_name>:<sec\_VV\_name>

The target name associated with this group, as set with the `creatercopygroup` command (see [page 11.48](#)). The target is created with the `creatercopytarget` command (see [page 11.50](#)). <sec\_VV\_name> specifies the name of the secondary volume on the target system. One <target\_name>:<sec\_VV\_name> must be specified for each target of the group.

**RESTRICTIONS**

Functionality of this command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

## EXAMPLES

In the following example, the `admitrcopyvv` command adds the volume `vv1` to the primary volume group `Group1`. At the same time, it adds the volume `vv1.r` on the target system `InServ1_in` to the corresponding secondary volume group that was previously created when the `creatercopygroup` command was issued:

```
cli% admitrcopyvv vv1 Group1 InServ1_in:vv1.r
```

## NOTES

- A secondary volume mapping must be provided for each target in the group.
- The virtual volume and the remote copy group must be in the same domain or both in no domain.



# 6

## Cancel Command

---

In this chapter

`canceltask`

**6.2**

---

**COMMAND**

`canceltask`

**DESCRIPTION**

The `canceltask` command cancels one or more tasks.

**SYNTAX**

`canceltask [-f] all | <task_ID>...`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-f`

Forces the command. The command completes the process without prompting for confirmation.

**SPECIFIERS**

`all`

Cancels all active tasks. If not specified, a task ID(s) must be specified.

`<task_ID>...`

Cancels only tasks identified by their task IDs. The `<task_ID>` must be an unsigned integer from 1 through 9999. If not specified, all tasks are cancelled.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example shows how to cancel a task using the task ID:

```
cli% canceltask 1
Are you sure you want to cancel task 1?
select q=quit y=yes n=no: y
```

## NOTES

- See the *3PAR InForm OS CLI Administrator's Manual* for additional information and examples regarding task management and task management commands.
- The `canceltask` command can return before a cancellation is completed. As a result, resources reserved for a task might not be immediately available. This can prevent actions like restarting the canceled task. Use the `waittask` command to ensure orderly completion of the cancellation before taking other actions. See [waittask](#) on page 31.2 for details.
- The `tunevv restart` commands enables you to resume a canceled `tunevv` task. See [tunevv](#) on page 28.18 for more information.





# 7

## Check Commands

---

In this chapter

checkhealth	<b>7.2</b>
checkId	<b>7.4</b>
checkpassword	<b>7.6</b>
checkpd	<b>7.8</b>
checkport	<b>7.11</b>
checkvv	<b>7.13</b>

---

## COMMAND

checkhealth

## DESCRIPTION

The `checkhealth` command checks the status of system hardware and software components, and reports any issues.



**NOTE:** Refer to the *3PAR InForm Messages and Operator's Guide* and *3PAR InForm OS CLI Administrator's Guide* for alert messages and administering your InServ Storage Server.

## SYNTAX

`checkhealth [<options> | <component>]`

## AUTHORITY

Super, Service

## OPTIONS

`-list`

Will list all components which `checkhealth` can be run on.

`-quiet`

Will not display which component is currently being checked.

`-detail`

Will display detailed information regarding the status of the system.

## SPECIFIERS

`<component>`

Indicates the component to check. Use `-list` option to get the list of components.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays as sample output of the `checkhealth` command:

```
cli% checkhealth
Checking alert
Checking cage
Checking date
Checking ld
Checking license
Checking network
Checking node
Checking pd
Checking pdch
Checking port
Checking rc
Checking snmp
Checking task
Checking vlun
Checking vv
Component -----Description----- Qty
Alert      New alerts                      6
Cage       Cages not on current firmware      2
Date       Date is not the same on all nodes  1
LD         LDs not mapped to a volume         3
License    Golden License.                   1
PD         PDs that are degraded or failed    2
pdch       LDs with chunklets on a remote disk 5
vlun       Hosts not connected to a port      3
vlun       Paths not reported by host agent   2
```

## NOTES

None.

---

## COMMAND

checkld

## DESCRIPTION

The `checkld` command executes consistency checks of data on logical disks in the event of an uncontrolled system shutdown and optionally repairs inconsistent logical disks.

## SYNTAX

`checkld [options] <LD_name>...`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-y|-n`

Specifies that if errors are found they are either modified so they are valid (`-y`) or left unmodified (`-n`). If not specified, errors are left unmodified (`-n`).

## SPECIFIERS

`<LD_name>...`

Requests that the integrity of a specified logical disk is checked. This specifier can be repeated to execute validity checks on multiple logical disks.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The `checkld` command should not be issued for logical disks that have been started because it can return incorrect mismatches.

## EXAMPLES

The following example displays a validity check of logical disk `vv63.admn.0`:

```
cli% checkld -n vv63.adm.0
Performing a consistency check only (no repairs)
Working on ld vv63.adm.0 (955)
compare:/dev/tpddev/pd/33 0x2180000 with:/dev/tpddev/pd/11 0x2180000 1 of 1
Logical disk vv63.adm.0 Check completed, logical disk consistent
```

## NOTES

- Repairing logical disks refers to making logical disks consistent.
- Consistency for RAID-1 means that all mirrors in the set have the same data.
- Consistency for RAID-5 or RAID-6 means that parity is consistent with the data in the set.
- The `checkld -n` command can be issued at any time as it reads only from the logical disk. If this command is issued for a started logical disk, false negative errors can be reported because of the contents of the physical disks changing during I/O and the difference in mirror update time.
- The `checkld -y` command can only be issued when the logical disk is in the `not started` state.

---

**COMMAND**

checkpassword

**DESCRIPTION**

The `checkpassword` command prompts for the <user>'s password and then displays the steps the system uses to authenticate the user and determine the user's privilege level. The information includes whether the user is local to the system or authenticated and authorized based on the configuration parameters set with `setauthparam` command for use with LDAP.

**SYNTAX**

checkpassword [ <user> ]

**AUTHORITY**

Super, Edit, Browse, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

None.

**SPECIFIER**

<user>

If the <user> parameter is not specified, then the current user is used. Only users with Super privileges with access to all domains can specify <user> names other than their own.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays the authentication and authorization of user User1:

```
cli% checkpassword User1
password:
+ attempting authentication and authorization using system-local data
user User1 is authenticated and authorized
```

- The output of the `checkpassword` command is a series of information statements, each starting with a plus sign (+) that indicates the steps the system is using to authenticate the user and determine the privilege level.
- The last line of output shows the summary of the steps with the user either being unauthenticated or authenticated and authorized.

## NOTES

- The output of the `checkpassword` command is based on current authentication and authorization parameters and might differ from the user's actual authorization level if the `setauthparam` command has been used to change parameters or data in the LDAP server has changed since the user last logged in.
- The `showuserconn` command can be used to verify the authorization levels assigned at login.

---

**COMMAND**

checkpd

**DESCRIPTION**

The `checkpd` command executes surface scans or diagnostics on physical disks.

**SYNTAX**

`checkpd scrub|diag [options <arg>] <pd_ID>...`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

The following options can only be used with the `scrub` specifier:

`-ch <number>`

To scan a specific chunklet rather than the entire disk.

`-count <number>`

To scan a number of chunklets starting from `-ch`.

The following options can only be used with the `diag` specifier:

`-path <p>`

Specifies a physical disk path as `a`, `b`, `both`, or `system`.

`-test <type>`

Specifies read, write, or verify test diagnostics. If no type is specified, the default is read.

`-iosize <size>`

Specifies I/O size, valid ranges are from 1s to 1m. If no size is specified, the default is 128k.

`-threads <num>`

Specifies of I/O threads, valid ranges are from 1 to 4. If the number of threads is not specified, the default is 1.



`-time <secs>`

Indicates the number of seconds to run, from 1 to 36000.

`-total <size>`

Indicates total bytes to transfer per disk. If a size is not specified, the default size is 1g.

`-retry <number>`

Specifies the total number of retries on an I/O error. If the number of retries is not specified, the default is 4.

`-range <size>`

Limits diagnostic regions to a specified size, from 2m to 2g.

## SPECIFIERS

`scrub`

Scans one or more chunklets for media defects.

`diag`

Performs read, write, or verifies test diagnostics.

`<pd_ID>...`

The ID of the physical disk to be checked. Only one `pd_ID` can be specified for the “scrub” test.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

In the following example, chunklet 500 on physical disk 1 is scanned for media defects:

```
cli% checkpd scrub -ch 500 1
{Tue Jul 10 17:11:26 PDT 2007} {31391} Open system device...
{Tue Jul 10 17:11:26 PDT 2007} {31391} Attach to system manager...
{Tue Jul 10 17:11:26 PDT 2007} {31391} Waiting for system manager ready...
{Tue Jul 10 17:11:26 PDT 2007} {31391} PD Scrubber 1.5 started
{Tue Jul 10 17:11:26 PDT 2007} {31391} Thread 16386 started for port 0:0:1
{Tue Jul 10 17:11:26 PDT 2007} {31391} Starting scan of pd 1 on port 0:0:1
{Tue Jul 10 17:11:27 PDT 2007} {31393} Scanning pd 1 ch 500
{Tue Jul 10 17:11:39 PDT 2007} {31393} Scan pd 1 ch 500 finished with 0
errors
{Tue Jul 10 17:11:39 PDT 2007} {31391} No media errors detected
```

## NOTES

- The <size> specifier can include a letter to indicate units:
  - ◆ g = gigabytes (2<sup>30</sup>)
  - ◆ t = terabytes (2<sup>40</sup>)
  - ◆ p = petabytes (2<sup>50</sup>)
  - ◆ m = 1048576 byte
  - ◆ k = 1024 bytes
  - ◆ s = 512 bytes
- I/O errors will be reported even if the eventual I/O succeeds due to retries.
- Up to 40 physical disk IDs can be specified for the diag test type.

---

**COMMAND**

checkport

**DESCRIPTION**

The `checkport` command performs a loopback test on Fibre Channel ports.

**SYNTAX**

`checkport [options <arg>] <N:S:P>`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-time <seconds_to_run>`

Specifies the number of seconds the test is to run. Use an integer from 0 to 300.

`-iter <iterations_to_run>`

Specifies the number of times the test is to run. Use an integer from 1 to 1000000.

**SPECIFIERS**

`<N:S:P>`

Specifies the port to be tested.

`node`

Specifies the node using a number from 0 through 7.

`slot`

Specifies the PCI bus slot in the specified node using a number from 0 through 5.

`port`

Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot using 1 through 4.

**RESTRICTIONS**

Access to all domains is required to run this command.

## EXAMPLES

In the following example, the loopback test is performed on port 0:0:1 a total of five times:

```
cli% checkport -iter 5 0:0:1
Starting loopback test on port 0:0:1
Port 0:0:1 completed 5 loopback frames in 0 seconds
Passed
```

## NOTES

- When both the `-time` and `-iter` options are specified, the first limit reached terminates the program. If neither are specified, the default is 1,000 iterations. The total run time is always limited to 300 seconds even when not specified.
- The default loopback is an ELS-ECHO sent to the HBA itself.
- QLogic HBAs do not support ELS-ECHOs to themselves, so a custom external loopback operation is used. This requires that no other devices are present on the port.

---

## COMMAND

checkvv

## DESCRIPTION

The `checkvv` command executes validity checks of Virtual Volume administration information in the event of an uncontrolled system shutdown and optionally repairs corrupted Virtual Volumes.

## SYNTAX

```
checkvv [options] <VV_name>...
```

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-y` | `-n`

Specifies that if errors are found they are either modified so they are valid (`-y`) or left unmodified (`-n`). If not specified, errors are left unmodified (`-n`).

## SPECIFIERS

`<VV_name> . . .`

Requests that the integrity of the specified Virtual Volume is checked. This specifier can be repeated to execute validity checks on multiple Virtual Volumes.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The `checkvv` command cannot be issued for Virtual Volumes that have been started.

## EXAMPLES

The following example displays a validity check of Virtual Volume `test1`:

```
cli% checkvv -n test1

Doing a consistency check only (no repairs)

Return PASS starting volume!!
```

## NOTES

None.

# 8

## CLI Command

---

In this chapter

cli

**8.2**

---

**COMMAND**

`cli`

**DESCRIPTION**

The `cli` command provides a means to set up your CLI session or to enter directly into a CLI shell.

**SYNTAX**

`cli [options] [<commands>]`

**AUTHORITY**

Super, Edit, Browse, Service

**OPTIONS**

`-v`

Displays the CLI client version.

`-b`

Displays the CLI client build level.

`-h`

Displays help for the `cli` command.

`-tpdportnum <portnum>`

Specifies the TCP port of the CLI server to which the CLI client connects. The default port number is 2540. The default SSL port number is 5783 and 5782 for an unsecure port.

`-sockssl`

Use SSL for a socket connection. SSL can be used when either this option is used or the `TPDSOCKSSL` environment variable has been set.

`-sys <sysname>`

Connect to system named `<sysname>`. When this option is not used, the CLI uses the value of the `TPDSYSNAME` environment variable. When the `TPDSYSNAME` environment variable is not set, the CLI prompts you for the system name.



`-pwf <passwdfile>`

Specifies the password file that contains the user name and encrypted password. If this option is not used, check the `TPDPWFILE` environment variable for the password filename. If the environment variable is not set, the CLI prompts you for the user name and password.

`-user <user_name>`

Specifies a user name to access the CLI. If used, this option must be issued with the `-password` option and overrides the `-pwf` option.

`-password <encrpw>`

Specifies an encrypted password to access the CLI. If used, this option must be issued with the `-user` option and overrides the `-pwf` option.

`-cmdloop`

Specifies that after commands are issued on the command line, an interactive command loop is entered.

`-hafter <nlines>`

Prints a header after `<nlines>` of data.

`-nohdtot`

Does not print header and totals.

`-csvtable`

Prints table data as Comma Separated Values (CSV).

`-listdom <domain_name>`

Indicates the domain associated with the storage server.

## SPECIFIERS

[ `<command>` ]

Any CLI command. This specifier is not required when issuing the `cli` command. If commands are specified, the CLI exits after executing the commands unless the `-cmdloop` option is specified. If no commands are specified in the command line, CLI enters the command loop.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays the CLI client version:

```
$ cli -v
CLI client version: 2.1.1
```

The following example identifies the storage server being accessed as `system1`:

```
$ cli -sys system1
```

The following example displays the current user `user1` and the user's password `pw2`:

```
$ cli -user user1 -password pw2
```

The following example places the CLI in an interactive command loop:

```
$ cli -cmdloop
```

## NOTES

- If the `cli` command is issued without any arguments, you are placed in a CLI shell.
- After commands are issued, the CLI exits unless the `-cmdloop` option was specified.

# 9

## Compact Commands

---

In this chapter

<code>compactcp</code>	<b>9.2</b>
<code>compactld</code>	<b>9.4</b>

---

**COMMAND**

`compactcpg`

**DESCRIPTION**

The `compactcpg` command consolidates logical disk space in Common Provisioning Groups (CPGs) into as few logical disks as possible, allowing unused logical disks to be removed and their space reclaimed.

**SYNTAX**

`compactcpg [options] <CPG_name|pattern>...`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-pat`

Compacts CPGs that match any of the specified patterns.

`-waittask`

Waits for any created tasks to complete.

`-trimonly`

Removes unused logical disks after consolidating the space. This option will not perform any region moves.

`-dr`

Specifies that the operation is a dry run, and the tasks will not actually be performed.

`-f`

Does not ask for confirmation before compacting the logical disks. Unless the `-f` option is specified, the command asks for confirmation before compacting each CPG.

## SPECIFIERS

<pattern>...

Specifies a glob-style pattern. This specifier can be repeated to compact multiple CPGs. If this specifier is not used, the <CPG\_name> specifier must be used. Refer to [Glob-Style Pattern](#) on page 2.4 for further information.

<CPG\_name>...

Specifies the name of the CPG. Multiple CPGs can be specified.

## RESTRICTIONS

None.

## EXAMPLES

The following example shows how to compact a single CPG named `testcpg` and remove any unused logical disks after consolidating the space:

```
cli% compactcpg -trimonly testcpg
Are you sure you want to compact CPG 'testcpg' ?
select q=quit y=yes n=no: y
Task 612 started
```

## NOTES

- Logical disks that are consolidated must have the same owner and backup nodes.
- If one logical disk exists with a different owner and backup node, that logical disk will be compacted individually.

---

**COMMAND**

`compactld`

**DESCRIPTION**

The `compactld` command consolidates space on the logical disks.

**SYNTAX**

`compactld [options <arg>] <LD_name>...|<pattern>...`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-pat`

Compacts the logical disks that match any of the specified patterns.

`-cons`

This option consolidates regions onto the fewest possible logical disks. When this option is not specified, the regions of each logical disk will be compacted within the same logical disk.

`-waittask`

Waits for any created tasks to complete.

`-taskname <taskname>`

Specifies a name for the task. When not specified, a default name is chosen.

`-dr`

Specifies that the operation is a dry run, and the tasks will not actually be performed.

`-f`

Does not ask for confirmation before compacting the logical disks. Unless the `-f` option is specified, the command asks for confirmation before compacting each logical disk.

`-trimonly`

Only unused logical disk space is removed. Regions are not moved.

## SPECIFIERS

<LD\_name> . . .

Specifies the name of the logical disk to be compacted. Multiple logical disks can be specified.

<pattern> . . .

Specifies a glob-style pattern. This specifier can be repeated to compact multiple logical disks. If this specifier is not used, the <LD\_name> specifier must be used. Refer to [Glob-Style Pattern](#) on page 2.4 for further information.

## RESTRICTIONS

None.

## EXAMPLES

In the following example, logical disk pdsld0.0 is compacted:

```
cli% compactld -f pdsld0.0
Compacting LDs:
  pdsld0.0
```

## NOTES

None.





# 10

## Control Commands

---

In this chapter

controliscsiport	<b>10.2</b>
controlmag	<b>10.5</b>
controlpd	<b>10.7</b>
controlport	<b>10.9</b>

---

## COMMAND

controliscsiport

## DESCRIPTION

The controliscsiport command is used to set properties of an iSCSI port.

## SYNTAX

The syntax for the controliscsiport command can be one of the following examples:

- controliscsiport addr <IP\_address> <netmask> [-f] <N:S:P>
- controliscsiport gw <gw\_address> [-f] <N:S:P>
- controliscsiport mtu <mtusz\_bytes> [-f] <N:S:P>
- controliscsiport dhcp on|off [-f] <N:S:P>
- controliscsiport isns <isns\_primary> [<isns\_secondary>][-f] <N:S:P>
- controliscsiport isnsport <isns\_port> [-f] <N:S:P>
- controliscsiport delete [-f] <N:S:P>
- controliscsiport ping [<count>] <ipaddr> <N:S:P>

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

addr

Sets the IP address and netmask of the iSCSI port.

gw

Sets the gateway address of the iSCSI port.

mtu

Sets the maximum transmission unit (MTU) size for the iSCSI port.

**dhcp**

Sets whether or not (on or off) the iSCSI target should use the DHCP to obtain its IP addresses.

**isns**

Sets the primary or secondary iSNS server IP addresses.

**isnsport**

Sets the TCP port number for the iSNS server. By default, the default iSNS port number is used.

**delete**

Deletes the iSCSI port configuration.

**ping**

Pings the specified IP address a specified number of times from the iSCSI port.

**OPTION****-f**

Do not ask for confirmation. The default is to ask for confirmation.

**SPECIFIERS****<IP\_address>**

Indicates the IP address of the iSCSI target.

**<netmask>**

Indicates the IP netmask of the iSCSI target.

**<gw\_address>**

Indicates the IP address of the gateway.

**<mtusz\_bytes>**

Indicates the MTU size in bytes.

**<isns\_primary>**

Indicates the IP address of the primary iSNS server.

**<isns\_secondary>**

Indicates the IP address of the secondary iSNS server.

**<isns\_port>**

Returns the TCP port number of the iSNS server. The default port number is 3205.

<count>

Indicates the number of ping packets to send. If a value is specified, it should be an integer from 1 through 64. If a value is not specified, the default is one packet.

<N:S:P>

The physical location of the iSCSI target port.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

In the following example, iSCSI port configuration is deleted:

```
cli% controliscsiport delete -f 104.64.98.1
```

## NOTES

None.

---

## COMMAND

controlmag

## DESCRIPTION

The `controlmag` command takes drive magazines, or disk drives within a magazine, either on loop or off loop. Use this command when replacing a drive magazine or disk drive within a drive magazine.

## SYNTAX

```
controlmag offloop|onloop [options] <cage_name> <magazine>
```

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

offloop|onloop

Specifies that the specified drive magazine or disk drive is either taken off loop or brought back on loop.

## OPTIONS

`-disk <disk_number>`

Specifies that the operation is performed on the disk as determined by its position within the drive magazine. If not specified, the operation is performed on the entire drive magazine.

`-port a|b|both`

Specifies that the operation is performed on port A, port B, or both A and B. If not specified, the operation is performed on both ports A and B.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

<age\_name>

Specifies the name of the drive cage. Drive cage information can be viewed by issuing the `showcage` command.

<magazine>

Specifies the drive magazine number within the drive cage. Valid formats are <drive\_cage\_number>.<drive\_magazine> or <drive\_magazine> (for example 1.3 or 3, respectively).

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example requests that drive magazine 1 in drive cage `cage0` be put on loop:

```
cli% controlmag onloop cage0 1
```

## NOTES

Taking a drive magazine off-loop has the following consequences:

- Relocation of chunklets.
- Affected logical disks are put into write-through mode.
- Momentary dip in throughput, but no loss of connectivity.

---

## COMMAND

`controlpd`

## DESCRIPTION

The `controlpd` command spins a physical disk up or down. This command is used when replacing a physical disk in a drive magazine.



**CAUTION:** Issuing the `controlpd` command puts the specified disk drive in a *not ready* state. Further, if this command is issued with the `spindown` subcommand, data on the specified drive becomes inaccessible.

## SYNTAX

`controlpd spinup|spindown [options] <WWN>...`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

`spinup`

Specifies that the physical disk is to spin up. If this subcommand is not used, then the `spindown` subcommand must be used.

`spindown`

Specifies that the physical disk is to spin down. If this subcommand is not used, then the `spinup` subcommand must be used.

## OPTIONS

`-ovrd`

Specifies that the operation is forced, even if the physical disk is in use.

## SPECIFIERS

<WWN> . . .

Specifies the World Wide Name (WWN) of the physical disk. This specifier can be repeated to identify multiple physical disks.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The spin down operation cannot be performed on a physical disk that is in use unless the `-ovrd` option is used.

## EXAMPLES

The following example instigates the spin-up of a physical disk identified by its WWN of 2000000087002078:

```
cli% controlpd spinup 2000000087002078
```

## NOTES

Issuing the `controlpd` command puts the specified disk drive in a `not ready` state. Further, if this command is issued with the `spindown` subcommand, data on the specified drive becomes inaccessible.



---

## COMMAND

`controlport`

## DESCRIPTION

The `controlport` command controls all aspects of a Fibre Channel or Remote Copy port, including the port's connection type and data rate.

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

`config`

Sets the specified connection mode and type on FC ports. Also allows setting the unique node WWN option for the port. When `unique_nwwn` is enabled, the port presents a unique node name on the connection and this is needed by certain initiators such as ONTAP.

`rst`

Resets a port.

`offline`

Holds the specified port offline indefinitely. Issue `controlport rst` to bring the port back online.

`lip`

Specifies that a Loop Initialization Primitive (LIP) command is issued from the port if there is a private loop topology. If the `-c` option is specified, then the LIP command is issued through the specified drive cage. If there is a point-to-point topology, then the link is reset. If there is a public loop or fabric topology, then a Registered State Change Notification is issued to the fabric controller.

`ct`

Sets the connection type. The specified port will be reset. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

This option has been deprecated and will be removed in a subsequent release.

## cl2

Specifies the Fibre Channel Class-2 parameter of the port. The specified port will be reset. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand. This option is deprecated and will be removed in a subsequent release.

## rate

Specifies the data rate of the Fibre Channel port. The specified port will be reset. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

## vcn

Sets the VLUN Change Notification (VCN) generation support (`enable` or `disable`). When VCN generation support is enabled with a public loop or fabric topology, a Registered State Change Notification (RSCN) message is issued to the fabric controller whenever a VLUN is created or removed. In addition, if enabled with a public loop topology, a Loop Initialization Primitive (LIP) is issued from the port whenever a VLUN is created or removed. See [Specifiers](#) on page 10.15 for additional information on parameters required to issue this subcommand.

## persona

Port personas have been removed and this command no longer sets them. As port personas have been replaced by host personas, this command will print the host persona commands to use instead. See `createhost` and `sethost` for setting host personas and `controlport config` for setting port parameters.

## nssync

Verifies current port database against the Name Server when a fabric is attached. Entries present in the database but missing from the Name Server are removed. Using this command is not required under normal circumstances.

## rcip addr

Sets the Remote Copy interface to use the specified IP address and netmask. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

## rcip gw

Sets the gateway for one or more Remote Copy interfaces. Only for RCIP ports. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

## rcip delete

Deletes the configuration for one or more specified Remote Copy interfaces. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

### `rcip mtu`

Sets the Maximum Transfer Unit (MTU) size for the specified Remote Copy interface(s), overriding the default of 1500. The largest supported value is 9000 and the smallest is 100. Only for RCIP ports. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

### `rcip state up|down`

Sets the specified Remote Copy interface(s) as either up or down. Only for RCIP ports. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

### `rcip speed`

Instructs the specified Remote Copy interface(s) to use the specified speed and duplex, or to auto negotiate speed and duplex. The default is `auto`. Only for RCIP ports. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

### `rcip ping`

Performs a ping from the specified interface to the specified address. Only for RCIP ports. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand. Use with the `-pf` option to prevent fragmentation of packets (see [Options](#) on page 10.13). Permitted for Super, Service, Browse, and Edit users.

### `rcfc init`

Sets the specified Remote Copy interface(s) on the local port. Only for RCFC ports.

### `rcfc delete`

Deletes the configuration for the Remote Copy interface on the local port. This will bring the interfaces down.

### `intcoal`

Enables or disables interrupt coalescing. The specified port will be reset. See [Specifiers](#) on page 10.15 for parameters required to issue this subcommand.

## SYNTAX

The syntax for the `controlport` command can be one of the following:

- `controlport rst [-m <mode>|-l|-i] [-f] <N:S:P>...`
- `controlport offline <N:S:P>`
- `controlport lip [-c <cagename>] [-f] <N:S:P>...`
- `controlport ct <ctval> [-f] <N:S:P>...` (deprecated usage)
- `controlport cl2 <cl2val> [-f] <N:S:P>...` (deprecated usage)
- `controlport rate <rateval> [-f] <N:S:P>...`
- `controlport vcn <vcnval> [-f] <N:S:P>...`
- `controlport nssync [-f] <N:S:P>...`
- `controlport intcoal <intcoalval> [-f] <N:S:P>...`
- `controlport rcip addr [-f] <IP_address> <netmask> <N:S:P>...`
- `controlport rcip gw [-f] <gateway_address> <N:S:P>...`
- `controlport rcip delete [-f] <N:S:P>...`
- `controlport rcip mtu <MTU_size> <N:S:P>...`
- `controlport rcip state up|down [-f] <N:S:P>...`
- `controlport rcip speed <value> half|full <N:S:P>...`
- `controlport rcip speed auto <N:S:P>...`
- `controlport rcip ping [-c <count>|-w <wait>|-s <size>|-pf] <IP_address> <N:S:P>...`
- `controlport rcfc init [-f] <N:S:P>`
- `controlport rcfc delete [-f] <N:S:P>`
- `controlport config <connmode> [-ct <ctval>] [-unique_nwwn enable|disable] [-f] <node:slot:port>...`

## OPTIONS

`-m <mode>`

This option can only be used with the `rst` subcommand. Resets the mode of the port. The port can be reset into a `target` or `initiator` mode. If not specified, the port is reset to its current mode. If the port's mode change value is prohibited, this command fails when attempting to reset to a different mode. Use `showport -c` to see whether mode change is allowed or prohibited for a particular port. See notes at the bottom of this section for additional information regarding port pair protection.

This option is deprecated and will soon be removed in a later release. Use `controlport config` to set `target` (host) or `initiator` (disk) modes.



**CAUTION:** The cable(s) must be disconnected from the port before resetting the mode to `target` or `initiator`.



**CAUTION:** Use caution when changing modes for ports in LSI Fibre Channel adapters. Changing the mode of one port in a pair (for example, from `initiator` to `target`) causes the other port in the pair to also undergo a mode change. In the case where one port in the pair is offline (and therefore a mode change can be allowed), but the partner port is online, changing the mode of the offline port causes the online partner port to undergo mode change as well. This results in loss of use of the partner port because it is no longer online.

`-l`

Forces the port to reload firmware. This option can only be used with the `rst` subcommand and cannot be used with the `-i` option.



**CAUTION:** Issuing the `controlport rst -l` command affects both ports of a port pair. Only use this command when irreversible damage has been done to a port or port pair. For more information about ports in your system, issue the `showport -i` command (see [showport](#) on page 22.121).

`-c`

When used with the Loop Initialization Primitive (LIP) subcommand:

`-c <cage_name>`

If using a private loop topology, a Loop Initialization Primitive (LIP) command is issued from the port. If a cage is specified using the `<cage_name>` argument, the LIP is issued through the cage controller.

- ◆ If using a point-to-point topology, the link is reset.
- ◆ If using a public loop or fabric topology, a Registered State Change Notification (RSCN) message is issued to the fabric controller.

When used with the `rcip ping` subcommand:

`-c <count>`

Specifies the number of replies accepted by the system before terminating the command. The default is 1; the maximum value is 25.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-i`

Forces a mode change for a port even if the port's mode change value is prohibited. The partner port's mode is changed with this option as well. Use with the `controlport rst` or `controlport persona` commands. This option cannot be used with the `-l` option.

`-w <wait>`

Specifies the maximum amount of time to wait for replies. The default is the number of requested replies plus 5. The maximum value is 30. If a number is not specified, the option can only be used with the `rcip ping` subcommand.

`-s <size>`

Specifies the packet size. If no size is specified, the option defaults to 64. This option and argument can only be used with the `rcip ping` subcommand.

`-pf`

Prevents packet fragmentation when issuing the `controlport rcip ping` command. This option can only be used with the `rcip ping` subcommand.

`-unique_nwwn <enable | disable>`

Enables or disables the use of a unique node WWN on the specified port.

`-ct`

Sets the connection type. The specified port will be reset.

## SPECIFIERS

`<ctval>`

Specifies the connection parameter of the port. Parameters can be `loop` or `point`. The `loop` parameter sets the port to arbitrated loop mode, the `point` parameter sets the port to point-to-point mode. This specifier must be provided when issuing the `ct` subcommand.

`<cl2val>`

Specifies the Fibre Channel Class-2 parameter of the port. Parameters can be one of `ack0`, `ack1`, or `disable`. This specifier must be provided when issuing the `cl2` subcommand.

`<rateval>`

Specifies the data rate of the Fibre Channel port. Rates can be one of `1`, `2`, `4`, or `auto`. `1` sets the data rate to 1 GBps, `2` sets the data rate to 2 GBps, and `4` sets the data rate to 4 GBps. The `auto` parameter sets the port to autodetect the data rate. This specifier must be used when issuing the `rate` subcommand.

`<vcnval>`

Specifies the value of the VCN. The VCN value can be set as `enable` or `disable`. This specifier must be used when issuing the `vcn` subcommand.

`<impval>`

Specifies the IMP port attribute. The IMP value can be set as `enable` or `disable`. This specifier must be used when issuing the `imp` subcommand.

`<intcoalval>`

Specifies if interrupt coalescing is enabled or disabled. The value can be set as `enable` or `disable`. This specifier must be used with the `intcoal` subcommand.

`<N:S:P>`

Specifies the port to be controlled.

`node`

Specifies the node using a number between 0 and 7.

`slot`

Specifies the PCI bus slot in the specified node using a number between 0 and 5.

`port`

Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot using 1 through 4.

`<IP_address>`

Specifies the IP address for a Remote Copy interface.

`<gateway_address>`

Specifies the gateway address for a Remote Copy interface.

`<netmask>`

Specifies the netmask for a Remote Copy interface.

`<MTU_size>`

Specifies the MTU size for a Remote Copy interface using an integer from 100 to 9000. If no integer is specified, the value defaults to 1500.

`<speed> half|full`

Use only with the `rcip speed` subcommand. Specifies the speed setting (10, 100, or 1000) and duplex setting (`half` or `full`) for a Remote Copy interface. In addition to this specifier, you must also specify an interface using `<N:S:P>`. If no speed or duplex settings are specified, or if you specify `auto` with the `rcip speed` subcommand, the system auto-negotiates the speed and duplex.

`<connmode> disk|host|rcfc`

Specifies whether the port is used for a disk, host, or RCFC connection. This specifier must be used when issuing the `config` command.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The `-m <mode>` option for the `rst` subcommand cannot be specified if there are active connections already using the port (that is the port online) except as noted in the following section under port pair protection.
- Port pair protection:
  - ◆ For dual-port LSI Fibre Channel adapters, both ports in the pair must use the same mode (initiator or target).
  - ◆ For quad-port LSI Fibre Channel adapters, each ports pair (ports 1 and 2, ports 3 and 4) must use the same mode.



- ◆ Changing the mode of one port in a pair (for example, from initiator to target) causes the other port in the pair to undergo a mode change as well.



**CAUTION:** Use caution when changing modes for ports in LSI Fibre Channel adapters. Changing the mode of one port in a pair (for example, from initiator to target) causes the other port in the pair to undergo a mode change. In the case where one port in the pair is offline (and therefore a mode change can be allowed), but the partner port is online, changing the mode of the offline port causes the online partner port to undergo a mode change as well. This results in loss of use of the partner port because it is no longer online.

- If there are active hosts or physical disks when issuing the `controlport rst` or `offline` commands, a warning is returned and you are prompted for confirmation to complete the execution of the commands.
- When issuing the `controlport ct`, `cl2`, `rate`, `persona`, or `rscn` commands, if there are active disks on the port, an error is returned. If there are active hosts on the port, the `-f` option is overridden (if specified), a warning is returned, and you are prompted for confirmation to complete the execution of the commands.
- The `controlport rcip addr` command is only allowed for node/slot/port combinations where there is an interface installed.

## EXAMPLES

The following example shows how to increase MTU to 9000 on Gigabit Ethernet port 1 in node 6, slot 3:

```
cli% controlport rcip mtu 9000 6:3:1
Remote Copy change successful.
```

The following example shows how to set Remote Copy interface 172.16.1.11 on a Gigabit Ethernet port 1 in node 6, slot 3 using a netmask of 255.255.255.0:

```
cli% controlport rcip addr 172.16.1.11 255.255.255.0 6:3:1
Are you sure you want to change the address for 6:3:1?
select q=quit y=yes n=no: y
Remote Copy interface change successful.
```

The following example shows how to set the gateway for Gigabit Ethernet port 1 in node 6, slot 3 using a gateway address of 172.16.1.1:

```
cli% controlport rcip gw 172.16.1.1 6:3:1
Are you sure you want to change the gateway for 6:3:1?
select q=quit y=yes n=no: y
Remote Copy interface change successful.
```

The following examples show usage of `controlport config`:

```
controlport config host -ct point 1:3:1
controlport config host -unique_nwwn enable 1:2:3
controlport config rcfc 0:0:1
controlport config disk 2:2:2
```

## NOTES

- The `ct`, `cl2`, or `rate` subcommands automatically resets the port for the changes to take effect.
- Issue the `showport` command with either the `-i` or `-par` options to verify the success of the `controlport` command. See [showport](#) on page 22.121.
- Resetting a port causes a momentary dip in throughput, but no loss of connectivity.
- The `-f` option forces the operation. Without the flag, the command prompts for confirmation. In some cases, unless the `TPDFORCE_OVERRIDE` environment variable is set, the command asks for confirmation even when the `-f` option is specified because the operation might disrupt the system operation.
- Use caution when changing modes for ports in LSI Fibre Channel adapters. Changing the mode of one port in a pair (for example, from initiator to target) causes the other port in the pair to also undergo a mode change. In the case where one port in the pair is offline (and therefore a mode change can be allowed), but the partner port is online, changing the mode of the offline port causes the online partner port to undergo mode change as well. This results in loss of use of the partner port because it is no longer online.
- Issuing the `controlport rst -l` command affects both ports of a port pair. Only use this command when irreversible damage has been done to a port or port pair.
- If specified, `<ctval>` will depend on the `<connmode>` value:
  - ◆ If `<connmode>` is `disk`, then `<ctval>` can only be `loop`

- ◆ If <connmode> is `rcfc`, then <ctval> can only be `point`
- ◆ If <connmode> is `host`, then <ctval> can be either `loop` or `point`



# 11

## Create Commands

---

### In this chapter

<code>createald</code>	<b>11.3</b>
<code>createaldvv</code>	<b>11.9</b>
<code>createavv</code>	<b>11.18</b>
<code>createcpg</code>	<b>11.22</b>
<code>createdomain</code>	<b>11.30</b>
<code>createdomainset</code>	<b>11.32</b>
<code>creategroupsv</code>	<b>11.34</b>
<code>creategroupvvcopy</code>	<b>11.37</b>
<code>createhost</code>	<b>11.40</b>
<code>createhostset</code>	<b>11.43</b>
<code>createld</code>	<b>11.45</b>
<code>creatercopygroup</code>	<b>11.48</b>
<code>creatercopytarget</code>	<b>11.50</b>
<code>createsched</code>	<b>11.52</b>
<code>createspare</code>	<b>11.56</b>

<code>createsv</code>	<b>11.58</b>
<code>createtemplate</code>	<b>11.61</b>
<code>createtpvv</code>	<b>11.71</b>
<code>createuser</code>	<b>11.77</b>
<code>createvlun</code>	<b>11.79</b>
<code>createvv</code>	<b>11.83</b>
<code>createvvcopy</code>	<b>11.90</b>
<code>createvvset</code>	<b>11.94</b>

---

**COMMAND**`createald`**DESCRIPTION**

This command was deprecated in the 2.2.3 release and will be changed or removed in a future release.

The `createald` command creates logical disks with automatic chunklet allocation for the disks.

**SYNTAX**

```
createald [options] <LD_name> <size>[g|G|t|T]
```

**AUTHORITY**

Super, Service, Edit



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-templ <template_name>`

Use the options defined in `template <tname>`. The template is created using the `createtemplate` command. Options specified in the template are read-only or read/write. The read/write options may be overridden with new options at the time of their creation, but read-only options may not be overridden at creation time. Options not explicitly specified in the template take their default values, and all of these options are either read-only or read/write (using the `-nro` or `-nrw` options of the `createtemplate` command).

`-t <RAID_type>`

Specifies the RAID type of the logical disk: `r0` for RAID-0, `r1` for RAID-1 or `r5` for RAID-5, or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklet>`

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row using an integer from 1 through 2147483647. If not specified, no row limit is imposed. The default is no limit.

`-ss <size_KB>`

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port|cage|mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the chunklet location preference characteristics, either *first* (attempt to use the lowest numbered available chunklets) or *last* (attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is *first*.

`-p <pattern>`

Specifies a pattern for disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to all Fibre Channel (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks matching the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL), FC, and Solid State Drive (SSD) drives.



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1,2,3). A range of nodes is separated with a hyphen (0-7). The primary path of the disks must be on the specified node number(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1,2,3). A range of slots is separated with a hyphen (0-7). The primary path of the disks must be on the specified PCI slot number(s).



`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–4). The primary path of the disks must be on the specified port number(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the *CagePos* column of *showpd* output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0–7). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0–3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0–3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create LDs based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that PDs with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that PDs with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that PDs with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that PDs with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that PDs identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <device_type>`

Specifies that PDs must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

`-domain <domain>`

Specifies the domain. The default is to create logical disk(s) in the current domain, or no domain if the current domain is not set.

`-wait <secs>`

If the `createald` command fails due to the lack of clean space, the `-wait` option specifies the number of seconds to wait for the system to clean the dirty space before returning.

If `-wait 0` is issued, the command returns immediately. If this option is not used, the command will keep waiting for dirty chunklets to be cleaned if enough space will be available with the dirty chunklets cleaned.

`-dr`

Specifies that the command is a dry run and that no logical disks or Virtual Volumes are created.

`-verbose on|off`

Specifies that verbose output is either enabled (`on`) or disabled (`off`). If not specified, verbose output is disabled.

`-cpsd <CPG_name>`

Specifies that the logical disk created is added to the specified Common Provisioning Group's (CPG's) snapshot data space.

`-cpsa <CPG_name>`

Specifies that the logical disk created is added to the specified CPG's snapshot administration space.

## SPECIFIERS

`<LD_name>`

Specifies the logical disk base name (up to 24 characters in length). The name is created by concatenating the name with an integer value - starting at zero, incrementing it by one for each logical disk generated, up to 999999. This field is required.

`<size>`

Specifies the minimum usable space in MB. Size should be an integer. An optional suffix (with no whitespace before the suffix) will modify the units to GB (`g` or `G` suffix) or TB (`t` or `T` suffix). If the size is to be taken from a template, this field should be `(-)`.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays the creation of a 256 MB logical disk named `testld.0`:

```
cli% createald testld 256
```

## NOTES

- Verify the creation of a logical disk by issuing the `showld` command. See [showld](#) on page 22.58 for more information.
- For this command:
  - ◆ MB = 1048576 bytes
  - ◆ GB = 1024MB
  - ◆ TB = 1024GB

---

## COMMAND

createaldevv

## DESCRIPTION

This command was deprecated in the 2.2.3 release and will be changed or removed in a future release. Please use the `createvv` command in the future.

The `createaldevv` command creates a virtual volume and its underlying logical disks and allows the system to automatically allocate resources to meet specified use requirements. The virtual volume and its logical disks can be created using either the listed options or by using preconfigured templates (created through the `createtemplate` command; see [page 11.61](#)).

## SYNTAX

```
createaldevv [options] <vvname> [.<index>] <size>[g|G|t|T]
```

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-templ <template_name>`

Use the options defined in template `<template_name>`. The template is created using the `createtemplate` command. Options specified in the template are read-only or read/write. The read/write options may be overridden with new options at creation time but read-only options may not be overridden at the time of creation.

Options not explicitly specified in the template take their default values, and all of these options are either read-only or read/write (using the `-nro` or `-nrw` options of the `createtemplate` command). If not included, the `-size` and `-cpg` options are automatically treated as read/write even if the other not included properties are marked read-only.

`-t <RAID_type>`

Specifies the RAID type of the logical disk: `r0` for RAID-0, `r1` for RAID-1, `r5` for RAID-5, or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklets>`

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row using an integer from 1 to 2147483647. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port|cage|mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the chunklet location preference characteristics: either first (attempt to use the lowest numbered available chunklets) or last (attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is *first*.

`-p <pattern>`

Specifies a pattern for disks. Patterns are used to select disks that are used for creating LDs. If no pattern is specified, the option defaults to Fibre Channel (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL), FC, and Solid State Drive (SSD) drives.

The following arguments can be specified as patterns for this option:

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0-7). The primary path of the disks must be on the specified node number(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0-7). The primary path of the disks must be on the specified PCI slot number(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0-5). The primary path of the disks must be on the specified port number(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0-3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-7). Disks must reside in the specified drive magazine(s).

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). Disks must reside in the specified position(s).

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that PDs with total chunklets greater than the number specified be selected.

`-tc_lt <number>`

Specifies that PDs with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel or NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.



`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

`-domain <domain>`

Specifies the domain. The default is to create it in the current domain, or no domain if the current domain is not set.

`-f`

Does not ask for confirmation before creating a RAID-0 volume.

`-wait <secs>`

If the command would fail due to the lack of clean space, the `-wait` option specifies the number of seconds to wait for the system to clean the dirty space before returning.

If `-wait 0` is issued, the command returns immediately. If this option is not used, the command will keep waiting for dirty chunklets to be cleaned if enough space will be available with the dirty chunklets cleaned.

`-dr`

Specifies that the command is a dry run and that no logical disks or virtual volumes are actually created.

`-verbose on|off`

Specifies that verbose output is either enabled (`on`) or disabled (`off`). If not specified, verbose output is disabled.

`-i <ID>`

Specifies the virtual volume ID. The default will be the next available virtual volume ID.

`-cnt <count>`

Specifies the number of identical virtual volumes to create. This must be an integer from 1 to 999. If `<id>` is specified with `-i` option, it is used as the ID of the first virtual volume and incremented by 1 for each subsequent volume.

`-shared`

Specifies that the system will try to share the logical disks among the virtual volumes. This option can only be used with the `-cnt` option.

`-szs <size>[g|G|t|T]`

Specifies the size of the snapshot volume. Defaults to megabytes. Size can be optionally specified in gigabytes (with `g` or `G` suffix) or terabytes (with `t` or `T` suffix). If not specified, snapshot space is 0. This option cannot be used with the `-pct` option.

`-pct <prc>`

Specifies the size of the snapshot volume as a percentage of the user volume. If not specified, the default value is 0. This option cannot be used with the `-szs` option. If used with the `-cpg <CPG_name>` option, the percentage value is 0.

`-cpg <CPG_name>`

Provision the SD and SA space from the CPG `<CPG_name>`.

`-exp <time>[d|D|h|H]`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.

`-f`

Do not ask for confirmation before creating a volume with retention time (`-retain`).

`-retain <time>[d|D|h|H]`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.



**NOTE:** If the volume is not in any domain, then its retention time cannot exceed the value of the system's `VVRetentionTimeMax`. The default value for the system's `VVRetentionTimeMax` is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's `VVRetentionTimeMax`, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the 3PAR Virtual Lock License. Contact your local 3PAR Authorized Service Provider for more information.

`-aw <percent>`

This option was deprecated in the 2.3.1 release and will be changed or removed in a future release.

Specifies the allocation warning threshold of the CPG. When the snapshot data space of the virtual volume exceeds the specified percentage of the virtual volume size, an alert is generated.

`-al <percent>`

Specifies the allocation limit threshold of the CPG. The snapshot data space is prevented from growing beyond the specified percentage of the virtual volume size.

`-spt <sectors_per_track>`

Defines the virtual volume geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is between 4 to 8192 and the default value is 304.

`-hpc <heads_per_cylinder>`

Defines the virtual volume geometry heads per cylinder value that is reported to the hosts through the SCSI mode pages. The valid range is between 1 to 1024 and the default value is 8.

`-pol <pol>[ ,<pol>...]`

Specifies VV policies.

`stale_ss`

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.

`no_stale_ss`

Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.

`one_host`

This constrains the export of a volume to one host or one host cluster (when cluster names can be used as a host name).

`no_one_host`

This policy should only be used when exporting a virtual volume to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

## SPECIFIERS

`<VV_name> [ .<index>]`

Specifies a VV name up to 31 characters in length. If the `-cnt` option is used, the optional decimal number `<index>` specifies the name of the first virtual volume (`<vvname>.<index>`). The `<index>` is incremented by 1 for each subsequent virtual volume. The `<index>` must be an integer from 0 to 999999. All virtual volume names have the same length constraint.

`<size>[g|G|t|T]`

Size for the user volume in MB (maximum of 16T). The volume size is rounded up to the next multiple of 256 MB. The size should be an integer. An optional suffix (with no whitespace before the suffix) will modify the units to GB (g or G suffix) or TB (t or T suffix). If the size is to be taken from a template, this field should be "-".

## RESTRICTIONS

When creating a logical disk, all physical disks must have the same device type.

## EXAMPLES

The following example creates 3 virtual volumes vv1.2, vv1.3, vv1.4:

```
cli% createaldvv -cnt 3 vv1.2 lg
```

The following example creates a RAID-5 virtual volume using disks in cages 0 and 1:

```
cli% createaldvv -t r5 -p -cg 0,1 vva lg
```

The following example creates 2 virtual volumes, vva.0 and vva.1 which may share logical disks:

```
cli% createaldvv -cnt 2 -shared vva lg
```

## NOTES

For this command:

- MB = 1048576 bytes
- GB = 1024MB
- TB = 1024GB

---

## COMMAND

`createavv`

## DESCRIPTION

This command was deprecated in the 2.2.3 release and will be changed or removed in a future release. Please use the `createvv` command in the future.

The `createavv` command creates a virtual volume where logical disk mapping is determined by the system.

## SYNTAX

The syntax for the `createavv` command can be one of the following:

- `createavv [options <arg>] <VV_name> <user_LD_ID> <usersize> <adminsace_LD_ID> <adminsace_size> <snapspace_LD_ID> <snapsize>`
- `createavv [options <arg>] <VV_name> <user_LD_ID> <usersize>`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-i <VV_ID>`

Specifies the ID of the virtual volume. If not specified, the next available virtual volume ID is assigned by the system.

`-verbose on|off`

Specifies that verbose output is either enabled (`on`) or disabled (`off`). The default is `off`.

`-dr`

Specifies that the command is a dry run and that no virtual volumes are actually created.

`-cpg <CPG_name>`

Specifies that the snapshot data space and snapshot administration space are provisioned from the indicated CPG (`<CPG_name>`). If this option is used, the `<adminsace_LD_ID>`, `<adminsace_size>`, `<snapspace_LD_ID>`, and `<snapsize>` specifiers cannot be used.

`-aw <percent>`

Specifies the percentage of used snapshot data space that when reached, results in a warning alert. To disable the warning, enter 0.

`-al <percent>`

Specifies the Virtual Volume's allocation limit. The snapshot data space of the Virtual Volume is prevented from growing beyond the specified percentage of the Virtual Volume size. After reached, new writes to the volume fail.

`-spt <sectors_per_track>`

Allows you to define the Virtual Volume geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is from 4 to 8192 and the default value is 304.

`-hpc <heads_per_cylinder>`

Allows you to define the Virtual Volume geometry heads per cylinder value that is reported to the hosts through the SCSI mode pages. The valid range is from 1 to 1024 and the default value is 8.

`-pol no_stale_ss | stale_ss | one_host | no_one_host`

Specifies the policy (rule) that the created virtual volume follows. If an argument is not specified, the policy defaults to `stale_ss`.

`stale_ss`

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume. However, the snapshot is considered invalid.

`no_stale_ss`

Specifies that invalid snapshot volumes are not permitted. A failure to update a snapshot is considered a failure to write to the base volume.

`one_host`

This constrains the export of a volume to one host or one host cluster (when cluster names can be used as a host name).

`no_one_host`

This policy should only be used when exporting a Virtual Volume to multiple hosts for use by a cluster-aware application, or when port presents VLUNs are used. This is the default policy setting.

## SPECIFIERS

<VV\_name>

Specifies the virtual volume name, using up to 31 characters. This option and argument are required on the command line.

<user\_LD\_ID>

Specifies the logical disks to be used as user space. The logical disks are identified by one or more integers (*item*). Integers can be provided as single number (1), a comma separated list of numbers (1, 2, 3), or a range of numbers separated with a dash (1–4).

<usersize>

Specifies the size of the user space in megabytes.

<adminsace\_LD\_ID>

Specifies the logical disks to be used as snapshot administrator space. The logical disks are identified by one or more integers (*item*). Integers can be provided as single number (1), a comma separated list of numbers (1, 2, 3), or a range of numbers separated with a dash (1–4).

<adminsize>

Specifies the size of the snapshot administrator space in megabytes.

<snapspace\_LD\_ID>

Specifies the logical disks to be used as snapshot data space. The logical disks are identified by one or more integers (*item*). Integers can be provided as single number (1), a comma separated list of numbers (1, 2, 3), or a range of numbers separated with a dash (1–4).

<snapsize>

Specifies the size of snapshot data space in megabytes.

<nb>

Specifies an integer.

<item>, <item>

Specifies the list of items.

<item>–<item>

Specifies a range of items.



## RESTRICTIONS

- Logical disks must be available to run the `createavv` command. Issue the `createald` command to create logical disks. See [createald](#) on page 11.3 for more information.
- The `<adminsace_LD_ID>`, `<adminsace>`, `<snapspace_LD_ID>`, and `<snapsize>` specifiers cannot be used if the `createavv -cpg` command is issued.

## EXAMPLES

The following example displays the creation of Virtual Volume `vv0` with its user space set to 256 MB and its administrator space set to 256 MB:

```
cli% createavv vv0 0 256 1 256 2 256
```

## NOTES

- Verify the creation of virtual volumes by issuing the `showvv` command. See [showvv](#) on page 22.201 for more information.
- A newly created logical disk is guaranteed to be clean. Chunklets of logical disks that are removed are cleaned before they are reused. However, regions of a logical disk that were previously used (for example, by another Virtual Volume) can contain data from its previous use. If these regions of the logical disk are mapped to your user space in a Virtual Volume, that data can be visible to the host that the Virtual Volume is exported.

If this is a concern, remove logical disks when the Virtual Volume is removed and use only newly created logical disks for your user space. Previous data in logical disks used for snapshot data space or snapshot admin space is not visible to you because these spaces are only visible after being written with new data.

- For this command, MB = 1048576 bytes.
- This command was deprecated in the 2.2.3 release and will be changed or removed in a future release.

---

**COMMAND**`createcpg`**DESCRIPTION**

The `createcpg` command creates a Common Provisioning Group (CPG).

**AUTHORITY**

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

```
createcpg [options <arg>] CPG_name
```

**OPTIONS**

`-templ <template_name>`

Use the options defined in the template `<template_name>`. The template is created using the `createtemplate` command. Options specified in the template are read-only or read/write. The read/write options may be overridden with new options at the time of their creation, but read-only options may not be overridden at creation time.

Options not explicitly specified in the template take their default values, and all of these options are either read-only or read/write (using the `-nro` or `-nrw` options of the `createtemplate` command).

`-aw <percent>`

Specifies the percentage of used snapshot administration or snapshot data space that results in a warning alert. A percent value of 0 disables the warning alert generation. The default is 0. This option is deprecated and will be removed in a subsequent release.



**NOTE:** The following options, `-sdgs`, `-sdgl`, and `-sdgw` control the auto logical disk creation for the common provisioning group's snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow (enlarge) size setting options (`-sdgs`, `-sdgl`).

`-sdgs <size> [g|G|t|T]`

Specifies the amount of logical disk storage created on each auto-grow operation. If `<size>` is non-zero it must be 8192 (8G) or bigger. A size of 0 disables the auto-grow feature. If the value is followed by a g or G, (no whitespace before g or G) the size is in GB. If the value is followed by a t or T, (no whitespace before t or T) the size is in TB. The default auto-grow size is fixed at 32G, but the minimum auto-grow is a function of the number of online nodes in the system:

Number of Nodes	Default	Minimum
1-2	32G	8G
3-4	64G	16G
5-6	96G	24G
7-8	128G	32G

`-sdgl <size> [g|G|t|T]`

Limits the auto-grow from exceeding this storage amount. A size of 0 means no limit is enforced. The storage amount can be specified in MB (default), GB (using g or G), or TB (using t or T). If the value is followed by a g or G, (no whitespace before g or G) the size is in GB. If the value is followed by a t or T, (no whitespace before t or T) the size is in TB. The default is 0.

`-sdgw <size>[g|G|t|T]`

Issues a warning alert when the used logical disk space exceeds this amount. A size of 0 means no warning limit is enforced. The default is 0. The size can be specified in MB (default), GB (using g or G), or TB (using t or T). If the value is followed by a g or G, (no whitespace before g or G) the size is in GB. If the value is followed by a t or T, (no whitespace before t or T) the size is in TB.

`-sa <LD_name>...`

Specifies that existing logical disks are added to the CPG and are used for snapshot admin (SA) space allocation. The `<LD_name>` argument can be repeated to specify multiple logical disks. This option is deprecated and will be removed in a subsequent release.

`-sd <LD_name> . . .`

Specifies that existing logical disks are added to the CPG and are used for snapshot data (SD) space allocation. The `<LD_name>` argument can be repeated to specify multiple logical disks. This option is deprecated and will be removed in a subsequent release.



**NOTE:** The following options, `-t`, `-szs`, `-rs`, `-ss`, `-ha`, `-ch`, and `-p` are used to control auto logical disk creation (if auto-grow is enabled).

`-domain <domain>`

Specifies the name of the domain with which the object will reside. The object must be created by a member of a particular domain with Edit or Super privileges. The default is created in the current domain, or no domain if the current domain is not set.

`-t <RAID_type>`

Specifies the RAID type of the logical disk: `r0` for RAID-0, `r1` for RAID-1, or `r5` for RAID-5 or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklets>`

Specifies the set size in the number of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row. The `<size>` is a positive integer. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port | cage | mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the chunklet location characteristics, either *first* (lowest numbered available chunklets) or *last* (highest numbered available chunklets). If no argument is specified, the default characteristic is *first*.

`-p <pattern>`

Specifies a pattern for candidate disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to all Fibre Channel (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL), FC, and Solid State Drive (SSD) drives:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0-7). The primary path of the disks must be on the specified node number(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0-7). The primary path of the disks must be on the specified PCI slot number(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0-4). The primary path of the disks must be on the specified port number(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0-3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-7). The specified magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified portion(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid` and `-devtype` are used to select the disks that are used to create common provisioning groups based on the characteristics of the disk.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified be selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified be selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified be selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified be selected.

`-devid <model>`

Specifies that PDs identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

`<CPG_name>`

Specifies the name of the common provisioning group being created.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the creation of common provisioning group `cpg1`, whose logical disk storage is set to autogrow at 16 GB, has a growth limit of 32 GB, and receives a growth warning at 24 GB:

```
cli% createcpg -sdgs 16g -sdgl 32g -sdgw 24g cpg1
```

The following example displays disks that satisfy all of the specified characteristics used. The example specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

```
cli% createcpg -p -fc_gt 60 -fc_lt 230 -nd 2
```

The following example specifies that all disks in cages 1 and 2 or magazine 4 of each cage will be used in the CPG.

```
cli% createcpg -p -cg 1-2 -p -mg 4 cpg3
```

## NOTES

- For this command:
  - ◆ KB = 1024 bytes
  - ◆ MB = 1048576 bytes
  - ◆ GB = 1073741824 bytes
- When using the `createcpg` command, the size of the logical disk space created is the first integer multiple of the RAID set size that is large enough to accommodate the requested virtual volume size. For example, with the default RAID-5 layout with a set size of 768 MB, a requested virtual volume size of 8192 MB causes the creation of logical disks with a total size rounded up to an integer multiple of 768 that is 8448 MB. The growth increment of CPGs is similarly rounded up because the growth is done by creating logical disks that must be created in units of the logical disk RAID set size. See the *InForm OS Concepts Guide* for further details.



- By default, logical disk are created using only physical disks with the same device type. (By default, the Fibre Channel device type is used). Use the `-p devtype NL` or `-p -devtype SSD` option to override this default. Use the `showpd` command to see the device types of physical disks in the system.
- The options `-sdgs`, `-sdgl` and `-sdgw` control the auto logical disk creation for the Common Provisioning Group's (CPG) snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow (enlarge) size setting options (`-sdgs`, `-sdgl`).
- If no device type is specified using the `-p -devtype` option, Fibre Channel is assumed.

---

**COMMAND**

`createdomain`

**DESCRIPTION**

The `createdomain` command creates system domains.

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

`createdomain [option <arg>] <domain_name>`

**OPTIONS**

`-comment <comment>`

Specifies any comments or additional information for the domain. The comment can be up to 511 characters long. Unprintable characters are not allowed.

`-vvretentiontimemax <value>[h|H|d|D]`

Specifies the maximum value that can be set for the retention time of a volume in this domain. `<time>` is a positive integer value in the range of 0 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

To disable setting the volume retention time in the domain, enter 0 for `<time>`.

**SPECIFIERS**

`<domain_name>`

Specifies the name of the domain you are creating. The domain name can be no more than 31 characters. The name `all` is reserved.

**RESTRICTIONS**

- Access to all domains is required to run this command.
- Use of this command requires a Domain license. Contact your local 3PAR representative for information.

## EXAMPLES

The following example displays the creation of domain `Engineering` with an informational comment:

```
cli% createdomain -comment "This is a test domain." Engineering
```

## NOTES

None.

---

**COMMAND**

createdomainset

**DESCRIPTION**

The `createdomainset` command defines a new set of domains and provides the option of assigning one or more existing domains to that set. The command also allows the addition of domains to an existing set by use of the `-add` option.

**SYNTAX**

```
createdomainset [options] <setname> [<domain | pattern>...]
```

**AUTHORITY**

Super, Edit

**OPTIONS**

`-add`

Specifies that the domains listed should be added to an existing set. At least one domain must be specified.

`-comment <comment>`

Specifies a comment relating to the set.

**SPECIFIERS**

`<setname>`

Specifies the name of the domain set to create or add to.

`<domain | pattern>...`

Optional list of domains or glob-style patterns matching domains that are to be included in the domain set. If no `<domain>` is specified a domain set with no domains is created.

**RESTRICTIONS**

None.

## EXAMPLES

To create a domain set with 2 domains in it:

```
cli% createdomainset domset domain1 domain2
```

To create an empty domain set with a comment:

```
cli% createdomainset -comment "A placeholder set" emptyset
```

## NOTES

None.

---

**COMMAND**

creategroupsv

**DESCRIPTION**

The `creategroupsv` command creates consistent group snapshots of a list of Virtual Volumes. Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent.

**SYNTAX**

`creategroupsv [option] <copy_of_VV>[:<snapshot_VV>[:<ID>]]...`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-ro`

Specifies that read-only snapshots are created. The default is read/write snapshots. To create read-only snapshots, all specified Virtual Volumes must be read/write. For read/write snapshots, all specified Virtual Volumes must be read-only. For information about snapshot rules, see the *InForm OS Concepts Guide*.

`-exp <time>[d|D|h|H]`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.

`-f`

Do not ask for confirmation before creating a volume with retention time (`-retain`).

`-retain <time>[d|D|h|H]`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.



**NOTE:** If the volume is not in any domain, then its retention time cannot exceed the value of the system's `VVRetentionTimeMax`. The default value for the system's `VVRetentionTimeMax` is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's `VVRetentionTimeMax`, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the 3PAR Virtual Lock License. Contact your local 3PAR Authorized Service Provider for more information.

## SPECIFIERS

`<copy_of_VV>`

Specifies the name of the Virtual Volume being copied (the parent volume).

`<snapshot_VV>`

Optional name of the snapshot Virtual Volumes. If `<snapshot_VV>` is not specified, the system tries to automatically generate the snapshot name of the form `<copy_of_VV>.<type><number>` where `<type>` is either `ro` or `rw` and `<number>` is either empty or the lowest number starting from 0 that does not conflict with an existing volume name. If the generated name is too long (because the `<copy_of_VV>` name is too long) the command will fail, and the user is required to specify `<snapshot_VV>` explicitly.

`[ <ID> ]`

Optional ID of the `<snapshot_VV>`. If the `<ID>` is not specified, an ID is chosen automatically.

## RESTRICTIONS

Use of this command requires a Virtual Copy license. Contact your local 3PAR representative for information.

## EXAMPLES

In the following example, Virtual Volumes VV1 and VV2, and their underlying logical disks are created using the `createaldivv` command. A read-only snapshot of the volumes is then taken and a list of the volumes is generated using the `creategroupsv` command:

```
cli% createaldivv -pct 10 VV1 4g
cli% createaldivv -pct 10 VV2 4g
cli% creategroupsv -ro VV1 VV2
CopyOfVV SnapshotVV
  VV1    VV1.ro
  VV2    VV2.ro
```

## NOTES

- Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent.
- A maximum of 256 writable virtual copies per volume are allowed.
- A maximum of 2048 virtual copies per volume are allowed.



---

## COMMAND

creategroupvvcopy

## DESCRIPTION

The `creategroupvvcopy` command creates consistent group physical copies of a list of virtual volumes.

## SYNTAX

The syntax for the `creategroupvvcopy` command can be one of the following:

- `creategroupvvcopy -p [options] <parent_VV>:<destination_VV>...`
- `creategroupvvcopy -r [options] <destination_VV>...`
- `creategroupvvcopy -halt <destination_VV>...`

## AUTHORITY

Super, Edit

## OPTIONS

`-p`

Starts a copy operation from the specified parent volume (as indicated using the `<parent_VV>` specifier) to its destination volume (as indicated using the `<destination_VV>` specifier). First a set of consistent group snapshots of all the `<parent_VV>` specifiers are taken and then each snapshot is copied to the corresponding `<destination_VV>` specifier. After the copies are complete, the snapshots are deleted unless the `-s` option is specified. Each `<parent_VV>` specifier must be a base virtual volume or a read/write snapshot.

`-r`

Resynchronizes the set of destination volumes (as indicated using the `<destination_VV>` specifier) with their respective parents using saved snapshots so that only the changes made since the last copy or resynchronization are copied. The resynchronization operation can only be performed if snapshots of the parents saved using the `-s` option in an earlier instance of the `creategroupvvcopy` command are present. These old snapshots are replaced by a new set of consistent group snapshots for the next resynchronization operation (such as the `-s` option need not be specified with the `-r` option).

**-halt**

Cancels an ongoing physical copy or snapshot promotion. This causes the destination volume (as indicated using the `<destination_VV> . . .` specifier) to be marked with the `copy failed` status, which will be cleaned up when they are promoted to base Virtual Volumes or when a new copy is started.

**-s**

Saves snapshots of the parent volume (as indicated with the `<parent_VV>` specifier) for quick resynchronization and to retain the parent-copy relationships between each parent and destination volume. The `-s` option is implied and need not be specified when the `-r` option is used. Each `<destination_VV>` specifier remains marked as a physical copy of its `<parent_VV>` specifier until it is promoted to a base virtual volume using the `promotevvcopy` command, which also removes the saved snapshot of the `<parent_VV>` specifier. The saved snapshots should not be removed manually. If the same `<parent_VV>` specifier is copied to different `<destination_VV>` specifiers with the `-s` option, a different snapshot of the `<parent_VV>` specifier is saved for each `<destination_VV>` specifier.

**-b**

Use this specifier to block until all the copies are complete. Without this option, the command completes before the copy operations are completed (use the `showvv` command to check the status of the copy operations).

**SPECIFIERS**

`<destination_VV>`

Indicates the destination virtual volume.

`<parent_VV>`

Indicates the parent virtual volume.

**RESTRICTIONS**

None.

**NOTES**

- The `creategroupvvcopy` command can be issued multiple times. However, the InServ system allows only two active physical copy tasks to run concurrently. Any additional physical copy tasks are queued, pending the completion of the active physical copy tasks.
- Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent.

---

## COMMAND

`createhost`

## DESCRIPTION

The `createhost` command creates or adds paths to a new system host and provides the option of assigning one or more paths to that host. Paths can be either Fibre Channel WWNs or iSCSI names. The command also provides options to annotate the host with descriptor information such as physical location, IP address, Operating System, Model, and so on.

## SYNTAX

The syntax for the `creathost` command can be one of the following:

- `createhost [options] <host_name> [<WWN>...]`
- `createhost -iscsi [options] <host_name> [<iSCSI_name>...]`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-add`

Adds the specified WWN(s) or `iscsi_name`(s) to an existing host (at least one WWN or `iscsi_name`, must be specified. See [Notes](#) for additional information about this option.

`-domain <domain | domain_set>`

Creates the host in the specified domain or domain set. The default is to create it in the current domain, or no domain if the current domain is not set. The domain set name must start with `set:`.

`-f`

Forces the tear down and removal of lower priority VLUN exports if necessary. See [Notes](#) for additional information about this option.

`-persona <hostpersonaval>`

Sets the host persona that specifies the personality for all ports which are part of the host set. This selects certain variations in SCSI command behavior which certain operating systems expect. The `<hostpersonaval>` is the host persona ID number with the desired capabilities. These can be seen with `showhost -listpersona`.

`-loc <location>`

Specifies the host's location.

`-ip <IP_address>`

Specifies the host's IP address.

`-os <OS>`

Specifies the operating system running on the host.

`-model <Model>`

Sets the host's model.

`-contact <contact>`

Sets the owner of the host and contact information.

`-comment <comment>`

Specifies any additional information for the host.

## SPECIFIERS

`<host_name>`

Specifies the name of the host, using up to 31 characters.

`<WWN>`

Specifies the World Wide Name (WWN) to be assigned or added to an existing host. This specifier can be repeated to specify multiple WWNs. This specifier is optional.

`<iSCSI_name>`

Host iSCSI name to be assigned or added to a host. This specifier is optional.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example creates system host `test01`:

```
cli% createhost test01 2000000087041F72
```

The following example creates an iSCSI host `test02`:

```
cli% createhost -iscsi test02 ign.1991-06.com.microsoft:dt-391-  
xp.hq.3par.com
```

The following example creates host `ibm520-13` with an 8/AIX-legacy host persona:

```
cli% createhost -persona 3 ibm520-13 10000000C97B142E
```

## NOTES

- If the host is running the 3PAR Host Agent then `createhost` may be called with no paths and the `hostname` reported by the agent to create the host with all reported paths. If no agent is running then `createhost` with no paths creates a host without assigning a host path. The `-add` option modifies an existing host by associating a host path with that host. If the `-add` option is not used, the `createhost` command defines a new host.
- If assigning paths to a host you are creating, specify the `-f` option to remove any existing VLUNs associated with those paths to avoid any conflicting VLUNs. See the *InForm OS Concepts Guide* for more information.
- Verify the creation of a host by issuing the `showhost` command. See [showhost](#) on page 22.47 for more information.
- The options that allow for adding descriptive information are for annotation purposes only; the storage server does not actively use the information provided here.
- Although it is optional to specify a host persona, a host must have one. If one is not specified, it will default to 1 (Generic). This will not operate correctly with all types of hosts. The host persona can be changed with `sethost -persona`.
- Host personas can not be specified with the `-add` argument.
- A port's FC parameters should be set with `controlport config`.

---

## COMMAND

createhostset

## DESCRIPTION

The `createhostset` command defines a new set of hosts and provides the option of assigning one or more existing hosts to that set. The command also allows the addition of hosts to an existing set by use of the `-add` option.

## SYNTAX

```
createhostset [options] <setname> [<host | pattern>...]
```

## AUTHORITY

Super, Edit

## OPTIONS

`-add`

Specifies that the hosts listed should be added to an existing set. At least one host must be specified.

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

`-domain <domain>`

Create the host set in the specified domain. For an empty set the default is to create it in the current domain, or no domain if the current domain is not set. A host set must be in the same domain as its members; if hosts are specified as part of the creation then the set will be created in their domain. The `-domain` option should still be used to specify which domain to use for the set when the hosts are members of domain sets. A domain cannot be specified when adding a host to an existing set with the `-add` option.

## SPECIFIERS

`<setname>`

Name of the host set to create/add to.

`<host | pattern>...`

Optional list of hosts or glob-style patterns matching hosts that are to be included in the host set. If no `<host>` is specified a host set with no hosts is created.

## RESTRICTIONS

None.

## EXAMPLES

To create an empty hostset:

```
cli% createhostset hostset
```

To add a host to the set:

```
cli% createhostset -add hostset hosta
```

To create a host set with a comment and a host in it:

```
cli% createhostset -comment "A host set" oraset oral
```

## NOTES

None.



---

## COMMAND

`createld`

## DESCRIPTION

This command was deprecated in the 2.2.3 release and will be changed or removed in a future release. Please use the `createvv` command in the future.

The `createld` command allows you to create a logical disk. Unlike the `createald` command, issuing the `createld` command requires that you allocate chunklets for the logical disk created.

## SYNTAX

`createld [options] <LD_name> <RAID_type> <row_size> <PD:ch>...`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-o <owner>`

Owner node for the logical disk.

`-b <backupowner>`

Backup owner node for the logical disk.

`-md <mirrordepth>`

Option `<raid_type>=1` only, the number of mirror copies in the set.

`-ps <parityset>`

Option `<raid_type>=5` or `6` only, the total number of chunklets in RAID set, including parity.

`-ss <stepsize>`

Step size is in KB and can take any value in the range of 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. The default is 256 KB for RAID-0 and RAID-1, 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-cpsd <CPG_name>`

Adds this logical disk to the specified CPG SD space.

`-cpsa <CPG_name>`

Adds this logical disk to the specified CPG SA space. Only RAID-1 logical disks may be used for the SA space.

`-ha port | cage | mag`

Defines the availability of the logical disk. The default is cage. The availability setting is used to select the destination chunklets during relocation of failed chunklets.

`-domain <domain>`

Creates the logical disk in the specified domain. The default is to create it in the current domain, or no domain if the current domain is not set.

## SPECIFIERS

`<LD_name>`

Specifies the logical disk name, using up to 31 characters.

`<RAID_type>`

Specifies the RAID type: 0 for RAID-0, 1 for RAID -1, 5 for RAID-5, or 6 for RAID-6.

`<row_size>`

Specifies the number of sets in a row.

`<PD:ch>`

Lists PDs and chunklet pairs.

## RESTRICTIONS

None.

## EXAMPLES

None.

## NOTES

- Verify the creation of a logical disk by issuing the `showld` command. See [showld](#) on page 22.58 for more information.

- To create a RAID-1 volume the mirror depth (`-md flag`) must be specified. For RAID-5 or RAID-6 volume the parityset (`-ps flag`) must be specified. The `owner` and `backup_owner` option should be a number between 0 and 7 and should be the node ID of a valid node in the system. The mirror depth can take any value between 2 and 4, inclusive.

---

**COMMAND**

creatercopygroup

**DESCRIPTION**

The `creatercopygroup` command creates a Remote Copy volume group.

**SYNTAX**

```
creatercopygroup [options] <group_name> <target_name>:<mode>
                  [<target_name>:<mode> ...]
```

**AUTHORITY**

Super, Edit

**OPTIONS**

-domain <domain>

Creates the Remote Copy group in the specified domain. The volume group must be created by a member of a particular domain with Edit or Super privileges.

**SPECIFIERS**

<group\_name>

Specifies the name of the volume group, using up to 22 characters if the `mirror_config` policy is set, or up to 31 characters otherwise. This name is assigned with this command.

<target\_name>

Specifies the target name associated with this group. This name should already have been assigned using the `creatercopytarget` command. See [creatercopytarget](#) on page 11.50 for details.

<mode>

Specifies that the mode of the created group is either kept in synchronization at all times (`sync`) or synchronized only periodically (`periodic`), either on command or by setting an automatic resynchronization period through the `setrcopygroup` command. See [setrcopygroup](#) on page 21.55 for details.

**RESTRICTIONS**

Functionality of this command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

## EXAMPLES

The following example creates an asynchronous periodic mode volume group named Group1 whose target system is InServ2 (target name InServ2\_out):

```
cli% creatercopygroup Group1 InServ2_out:periodic
```

If using domains, the syntax would appear as:

```
cli% creatercopygroup -domain domain2 Group1 InServ2_out:periodic
```

## NOTES

None.

---

## COMMAND

creatercopytarget

## DESCRIPTION

The `creatercopytarget` command creates a Remote Copy target definition.

## SYNTAX

The syntax for the `creatercopytarget` command can be one of the following:

- Syntax for Remote Copy over IP (RCIP) is as follows:

```
creatercopytarget [options] <target_name> IP  
                  [<node:slot:port:IP_address>...]
```

- Syntax for Remote Copy over Fibre Channel (RCFC) is as follows:

```
creatercopytarget [options] <target_name> FC <node_WWN>  
                  [<node:slot:port:WWN>...]
```

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

-disabled

Create the target in disabled mode.

## SPECIFIERS

<target\_name>

The name of the target definition to be created, specified by using up to 31 characters. See the *Remote Copy User's Guide* for suggested target naming conventions.

<node\_WWN>

The system's node World Wide Name (WWN) on the target system (Fibre Channel target only).

<node:slot:port:IP\_address>

Specifies the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.

<node:slot:port:WWN>

Specifies the node, slot, port of the Fibre Channel port on the local system and a World Wide Name (WWN) address on the target system.

## RESTRICTIONS

- Access to all domains is required to run this command.
- Functionality of this command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

## EXAMPLES

The following example sets up IP targets and links on `System2` (target name), node 0 and node 1. The IP addresses specify the addresses on the target system:

```
cli% creatercopytarget System2 IP 0:193.1.1.11 1:193.1.2.11
```

The following example sets up FC targets and links to `System2` (target name), node 0 and node 1:

```
cli%creatercopytarget System2 FC 2FF70002AC0000C3 0:1:1:20210002AC0000C3  
1:3:2:21210002AC0000C3
```

## NOTES

- An earlier version of this command provided the options `primary` and `secondary` for setting active and backup systems. These options were deprecated in InForm OS release 2.2.3.
- IP targets are made up of pairs composed of the node, slot, and port of the ethernet port on the local system and an IP address of the peer port on the target system.
- FC targets are made up of sets with the node, slot, and port of the FC port on the local system and WWN of the peer port on the target system.

---

## COMMAND

`createsched`

## DESCRIPTION

The `createsched` command allows users to schedule tasks that are periodically run by the scheduler.

## SYNTAX

`createsched [options] <cmd> <taskschedule> <schedname>`

## AUTHORITY

Super, Service

## OPTIONS

`-run_once`

The task will only run once at the specified time.

`-no_alert`

Tasks created with the `no_alert` option will not generate alerts if the task fails.

## SPECIFIERS

`<cmd>`

Specifies that it is either a CLI command or script that has been distributed by 3PAR. Can be up to 127 characters in length.

`<taskschedule>`

Specifies the use of a crontab-style schedule. Each field can be up to 127 characters in length.

`<schedname>`

Specifies the the name given to the schedule. The name can be up to 31 characters in length.

## RESTRICTIONS

- The use of **All** and **Upgrade** as schedule names are reserved.



- Only the following commands can be scheduled: `updatesnapspace`, `createsv`, `removevv`, `createvvcopy`, `creategroupsv`, `creategroupcopy`, `moverlocpd`, `checkhealth`, `compactcpv`, `compactld`, `syncrcopy`, `tunealddv`, `tunepd`, `tunetpvv`, `tunevv`, `updatevv`.

## EXAMPLES

The following example creates a snapshot every hour that expires in 2 hours for volume `vvname`:

```
cli% createsched "createsv -ro -exp 2h @vvname@.@s@ vvname" @hourly snp_vv
```

The following example creates a snapshot every hour that expires in 2 hours for volumes in `vvset`:

```
cli% createsched "createsv -ro -exp 2h @vvname@.@s@ set:vvset" @hourly snp_vvset
```

The following example creates a scheduled task that will run every hour on the hour ("0 \* \* \* \*") and will create a read-only (-ro) snapshot of sample\_vv, that will expire in two hours (-exp 2h). The created snapshot will be named sample\_vv\_set.HH.MMM where HH and MM refer to the hour and minute that the snapshot was created (@vvname@.@H@.@M@):

```
cli% createsched "createsv -ro -exp 2h @vvname@.@H@.@M@ vv" "0 * * * *" sv_task
```

The following example creates a scheduled task which will create a consistent group snapshot. The read-only snapshots will be created from the volume set set: `dbora` (see CLI Help `-sub objectsets`) that will expire in two hours (`-exp 2h`) and the snapshots will be named using the pattern `@vvnname@.@s@` (volume name specified in the set, with the epoch in seconds appended). The task will be named `snap_odbora_set` and run every hour on the hour during the working hours from 8am to 5pm ("`0 8-17 * * *`"):

```
cli createsched "createsv -ro -exp 2h @vname@.@s@ set:dbora" "0 8-17 * * *"
sv dbora_task
```

The following example creates a scheduled task which will create a consistent group snapshot for a single virtual volume:

```
createsched "createsv -ro -exp 2h @vname@.@s@ vname" "0 * * * *"
snap large tpvv a
```

## NOTES

- The <taskschedule> string has five fixed fields: minute, hour, day-of-month, month, day-of-week.
- Fields are separated by spaces. The allowed values for the fields are:

Field	Allowed Values
Minute	0-59
Hour	* or 0-23
Day-of-Month	* or 1-31
Month	* or 1-12
Day-of-Week	* or 0-6 (0 is Sunday)

- Scheduling on February 29 is not supported.
- Lists are allowed. A list is a set of numbers (or ranges) separated by commas. Examples: "1,2,5,9", "0-4,8-12".
- Ranges of numbers are allowed. Ranges are two numbers separated with a hyphen. The specified range is inclusive. For example, 8-11 for an hour entry specifies execution at hours 8, 9, 10 and 11. Ranges must be specified from low to high.

- One of six special strings may also appear instead of the specification above:

String	Meaning
@yearly	Run every January 1st, "0 0 1 1 *"
@monthly	Run the first day of every month, "0 0 1 * *"
@weekly	Run every Sunday, "0 0 * * 0"
@daily	Run every midnight, "0 0 * * *"
@hourly	Run every hour, on the hour, "0 * * * *"

- If some form of daylight savings or summer/winter time is in effect, then jobs scheduled during the switchover period could be executed once, twice, or not at all.

---

## COMMAND

`createspare`

## DESCRIPTION

The `createspare` command allocates chunklet resources as spares. Chunklets marked as spare are not used for logical disk creation and are reserved explicitly for spares, thereby guaranteeing a minimum amount of spare space.

## SYNTAX

`createspare [options <arg>] <chunklet_specifier>...`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-f`

Do not ask for confirmation. The default is to ask for confirmation.

`-p`

Specifies that partial completion of the command is acceptable. Additionally, specified chunklets are marked as spare only if they are not spare already.

## SPECIFIERS

`<chunklet_specifier>...`

The chunklet specifier is one of the following arguments:

`<PD_ID:chunklet_number>`

Specifies the identification of the physical disk and the chunklet number on the disk. This specifier can be repeated.

`<PD_ID:a>`

Specifies the identification of the physical disk and all chunklets (a) on the disk. This specifier can be repeated.

`a:<chunklet_num>`

Specifies a chunklet number on all physical disks. This specifier can be repeated.

`-pos <cage:mag:disk:chunklet_num>`

Specifies the position of a specific chunklet identified by its position in a drive cage, drive magazine, physical disk, and chunklet number. This specifier can be repeated.

`-pos <cage:mag:disk:a>`

Specifies that all chunklets on a physical disk, identified by drive cage number, drive magazine number, and physical disk number, are marked spare. This specifier can be repeated.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

- The following example marks chunklet 1 as spare for physical disk 15:

```
cli% createspare 15:1
```

- The following example specifies the position in a drive cage, drive magazine, physical disk, and chunklet number. `-pos 1:0.2:3:121`, where 1 is the drive cage, 0.2 is the drive magazine, 3 is the physical disk, and 121 is the chunklet number.

```
cli% createspare -pos 1:0.2:3:121
```

## NOTES

To verify the creation of a spare chunklet, issue the `showspare` command. See [showspare \[-used\]](#) on page 22.166 for more information.

---

**COMMAND**

`createsv`

**DESCRIPTION**

The `createsv` command creates a point-in-time (snapshot) copy of a virtual volume. Create snapshots to perform such tasks as backing up data on the base volume and allowing multiple copies of a base volume to be modified without affecting the original base volume.

**SYNTAX**

`createsv [options <arg>] <SV_name> <copy_of_VV | VV_set>`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-ro`

Specifies that the copied volume is read-only. If not specified, the volume is read/write.

`-i <VV_ID>`

Specifies the ID of the copied VV.

`-exp <time>[d|D|h|H]`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.

`-f`

Do not ask for confirmation before creating a volume with retention time (`-retain`).

`-retain <time>[d|D|h|H]`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.



**NOTE:** If the volume is not in any domain, then its retention time cannot exceed the value of the system's `VVRetentionTimeMax`. The default value for the system's `VVRetentionTimeMax` is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's `VVRetentionTimeMax`, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the 3PAR Virtual Lock License. Contact your local 3PAR Authorized Service Provider for more information.

## SPECIFIERS

`<SV_name>`

Specifies the snapshot name, using up to 31 characters in length.

`<copy_of_VV | VV_set>`

Specifies the parent volume name or volume set name, using up to 31 characters in length. The `VV_set` name must start with `set:`. If `copy_of_VV` is a `VV_set` then `svname` can contain various patterns that are used to form the snapshot volume name. See `Help sub,vvnamepat` for details.

## RESTRICTIONS

- Use of this command requires a Virtual Copy license. Contact your local 3PAR representative for information.
- A read-only snapshot of a read-only source volume is not allowed.
- A read/write snapshot of a read/write snapshot or base is not allowed.

## EXAMPLES

The following example creates a read-only snapshot volume `svro_vv0` from volume `vv0`:

```
cli% createsv -ro svr0_vv0 vv0
```

The following example creates snapshot volume `svrw_vv0` from the snapshot `svro_vv0`:

```
cli% createsv svrw_vv0 svro_vv0
```

The following example creates a snapshot volume for each member of the VV set `vvcopies`. Each snapshot will be named `svro-<name of parent volume>`:

```
cli% createsv -ro svro-@vvname@ set:vvcopies
```

If `vvcopies` contained VVs named `vv.0`, `vv.1` and `vv.2` this would result in snapshots named `svro-vv.0`, `svro-vv.1` and `svro-vv.2`.

## NOTES

- A maximum of 256 writable virtual copies per volume are allowed.
- A maximum of 500 virtual copies per volume are allowed.
- For the `VV_set` option, the `createsv` command creates consistent group snapshots of the VVs in the set. Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the set remain consistent.



---

## COMMAND

createtemplate

## DESCRIPTION

The `createtemplate` command creates Virtual Volume (VV), Logical Disk (LD), and Common Provisioning Group (CPG) templates. Templates are sets of predetermined parameters that are applied to the creation of VVs, LDs, and CPGs.

## SYNTAX

```
createtemplate VV|LD|CPG [options <arg>] <template_name>
```

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

- The following options are valid for VV, LD, and CPG templates:

`-nro`

Indicates that properties not specified in the template are set to read-only.

`-nrw`

Indicates that properties not specified in the template are set to read/write.

`-desc <description>`

Specifies a description for the created template.

`-rw`

Specifies that the list of property values in the remainder of the command line or until the next `-rw` or `-ro` option are read-write, which means that they can be modified when the template is used.

`-ro`

Specifies that the list of property values in the remainder of the command line or until the next `-rw` or `-ro` option are read-only, which means that they cannot be modified when the template is used. If neither `-ro` nor `-rw` is specified, the default is `-ro`.

`-t <RAID_type>`

Specifies the RAID type of the logical disk: `r0` for RAID-0, `r1` for RAID-1, `r5` for RAID-5, or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklets>`

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row using a positive integer. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port|cage|mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the chunklet characteristics, either `first` (attempt to use the lowest numbered available chunklets) or `last` (attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is `first`.

`-p <pattern>`

Specifies a pattern for candidate disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to all Fibre Channel (FC) type disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL) FC, and Solid State Drive (SSD) drives:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0–7). The primary path of the disks must be on the specified node number(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0–7). The primary path of the disks must be on the specified PCI slot number(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–4). The primary path of the disks must be on the specified port number(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the `CagePos` column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0–3). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0–3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create LDs based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

► The following options are used only for VV templates:

`-type cpvv|tpvv|none`

Specifies the type of VV for the template is being created. Valid arguments are `cpvv`, `tpvv`, or `none`. Use `cpvv` if the `-cpg <CPG_name>` option is used. If creating a template for creating TPVVs, use `tpvv`. If `cpvv` or `tpvv` are not specified, the default value is `none` meaning the created template is applicable for the creation of any volume type.

`-cpg <CPG_name>`

Specifies that the snapshot data space and snapshot administration space are provisioned from the indicated CPG (`<CPG_name>`). The template with this option can only be used with `createalddv` command. This option is deprecated and will be removed in a subsequent release.

`-aw <percent>`

Specifies the *allocation warning* threshold of the CPG. When the snapshot data space of the VV exceeds the specified percentage of the virtual volume size, an alert is generated. The template with this option can only be used with `createalddv` command. This option is deprecated and will be removed in a subsequent release.

`-al <percent>`

Specifies the *allocation limit* threshold of the CPG. The snapshot data space is prevented from growing beyond the specified percentage of the virtual volume size. The template with this option can only be used with `createalddv` command. This option is deprecated and will be removed in a subsequent release.

`-size <size>[g|G|t|T]`

Specifies the size of the virtual volume and logical disk in megabytes. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space).

`-szs <size>[g|G|t|T]`

By default, specifies the size of the snapshot volume in megabytes using an integer from 0 through 2096128 (2047 GB). Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). If not specified, snapshot space is 0. This option cannot be used with the `-pct` option.

The template with this option can only be used with the `createaldevv` command.

This option is deprecated and will be removed in a subsequent release.

`-pct <prc>`

Specifies the size of the snapshot volume as a percentage of the user volume. If not specified, the default value is 0 percent of your user volume. This option cannot be used with the `-szs` option. If used with the `-cpg <CPG_name>` option, the percentage value is 0. The template with this option can only be used with the `createaldevv` command. The template with this option This option is deprecated and will be removed in a subsequent release.

`-spt <sectors_per_track>`

Defines the virtual volume geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is between 4 to 8192 and the default value is 304.

`-hpc <heads_per_cylinder>`

Allows you to define the virtual volume geometry heads per cylinder value that is reported to the hosts through the SCSI mode pages. The valid range is between 1 to 1024 and the default value is 8.

`-pol <pol>[, <pol>...]`

Specifies the policy that the created virtual volume follows. If an argument is not specified, the option defaults to `stale_ss`.

**stale\_ss**

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.

**no\_stale\_ss**

Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.

**one\_host**

This constrains the export of a volume to one host or one host cluster (when cluster names can be used as a host name).

**no\_one\_host**

This policy should only be used when exporting a virtual volume to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.

**tp\_bzero**

Specifies that if a host write results in the allocation of a new data page that is only partially filled by the host write, then a zero-fill is performed on the unwritten portion of the data page. This ensures that the host cannot read data from deleted volumes or snapshot. The default allocation page size is 16 KB.

**no\_tp\_bzero**

Specifies that the zero-fill operation is bypassed on the allocation of partially written data pages.

**-usr\_cpg <usr\_cpg>**

Specifies the name of the CPG from which the user space will be allocated.

**-usr\_aw <percent>**

Indicates a user space allocation warning. Generates a warning alert when the user data space of the TPVV exceeds the specified percentage of the virtual volume size.

**-usr\_al <percent>**

Indicates the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the virtual volume size. After this limit is reached, any new writes to the virtual volume will fail.

`-snp_cpg <snp_cpg>`

The name of the CPG from which the snapshot space is allocated.

`-snp_aw <percent>`

Provides a snapshot space allocation warning. Generates a warning alert when the snapshot space of the virtual volume exceeds the specified percentage of the virtual volume size.

`-snp_al <percent>`

Provides a snapshot space allocation limit. The snapshot space of the virtual volume is prevented from growing beyond the indicated percentage of the virtual volume size.

► The following options are only used for CPG templates:

`-aw <percent>`

Specifies the percentage of used snapshot administration or snapshot data space that results in a warning alert. A percent value of 0 disables the warning alert generation. The default is 0.

This option is deprecated and will be removed in a subsequent release.

`-sdgs <size>[g|G|t|T]`

Specifies the growth increment, the amount of logical disk storage created on each auto-grow operation. The default growth increment is fixed at 32G, but the minimum growth increment varies according to the number of controller nodes in the system. If `<size>` is non-zero it must be 8G or bigger. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter or terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). A size of 0 disables the auto-grow feature. The following table displays the default and minimum growth increments per number of nodes:

Number of Nodes	Default	Minimum
1-2	32G	8G
3-4	64G	16G
5-6	96G	24G
7-8	128G	32G



`-sdgl <size>[g|G|t|T]`

Limits the auto-grow from exceeding this storage amount. A size of 0 means no limit is enforced. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter or terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). The default is 0.

`-sdgw <size>[g|G|t|T]`

Issues a warning alert when the used logical disk space exceeds this amount. A size of 0 means no warning limit is enforced. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter or terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). The default is 0.

## SPECIFIERS

`VV|LD|CPG`

Specifies that the template is for the creation of a Virtual Volume (VV), Logical Disk (LD), or Common Provisioning Group (CPG).

`<template_name>`

Specifies the name of the template, up to 31 characters in length.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example creates a virtual volume template called `vvrl`. The `-ha` and `-pol` option values are fixed. The RAID type can be changed when the virtual volume is created. Note that if patterns are specified, it is not possible to mix read-only and read/write pattern specifications.

```
cli% createtemplate vv -ha mag -rw -t r1 -ro -pol stale_ss vvrl
```

The command rejects the pattern.

```
cli% createtemplate vv -ro -p -nd 0 -rw -p -mg 1,2 vpatt vvrl
```

## NOTES

- By default, logical disk are created using only physical disks with the same device type. (By default, the Fibre Channel device type is used). Use the `-p devtype NL` or `-p -devtype SSD` option to override this default. Use the `showpd` command to see the device types of physical disks in the system.
- The value for `-ssz` and `-rs` can be set to `-`. This can be used in conjunction with the (default) read-only property for options to ensure that their value cannot be changed when a virtual volume or logical disk or CPG is created and the system is used to set the default values for these options.
- The options `-sdgs`, `-sdgl` and `-sdgw` control the auto logical disk creation for the CPG's snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow (enlarge) size setting options (`-sdgs`, `-sdgl`).
- For this command:
  - ◆ KB = 1024 bytes
  - ◆ MB = 1024 KB
  - ◆ GB = 1024 MB

---

**COMMAND**

`createtpvv`

**DESCRIPTION**

This command is deprecated and will be removed in a subsequent release. Please use the `createvv -tpvv` command in the future.

The `createtpvv` command creates a multi-space Thinly Provisioned Virtual Volume (TPVV). When the `-snp_cpg` option is specified, both of the user space and snapshot space of the TPVV are created.

When the `-snp_cpg` option is not specified, only the user space of the TPVV is created and the snapshot space is not created. The user can't create the snapshots of the volume without creating the snapshot space.

**AUTHORITY**

Super, Edit

**SYNTAX**

`createtpvv [options] <usr_cpg> <vvname> [.<index>] <virtsize>[g|G|t|T]`

**OPTIONS**

`-i <ID>`

Specifies the volume ID.

`-cnt <count>`

Specifies the number of identical Virtual Volumes (VVs) to create. This must be an integer from 1 to 999. If `<id>` is specified with `-i` option, it is used as the ID of the first VV and incremented by 1 for each subsequent volume.

- The following option specifies the name of the template used to create a TPVV.

`-templ <template_name>`

Use the options defined in template `<template_name>`. The template is created using the `createtemplate` command. Options specified in the template are read-only or read-write. The read-write options may be overridden with new options at the time of their creation, but read-only options may not be overridden at the time of creation. Options not explicitly specified in the template take their default values, and all of these options are either read-only or read-write (using the `-nro` or `-nrw` options of the `createtemplate` command). If not included, the `-size` and `-cpg` options are automatically treated as read-write even if the other unincluded properties are marked read-only.

- The following option is used to create the snapshot space of the TPVV:

`-snp_cpg <snp_cpg>`

Specifies the name of the Commonly Provisioned Group (CPG) from which the snapshot space will be allocated.

- The following options are used to set the allocation policies of the TPVV:

`-usr_aw <percent>`

This option enables user space allocation warning. It specifies that a warning alert is generated when the user space of the TPVV exceeds the specified percentage of the VV size.

`-usr_al <percent>`

This option sets the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the VV size. After this size is reached, any new writes to the VV will fail.

`-snp_aw <percent>`

Enables a snapshot space allocation warning. A warning alert is generated when the snapshot space of the VV exceeds the indicated percentage of the VV size.

`-snp_al <percent>`

Sets a snapshot space allocation limit. The snapshot space of the VV is prevented from growing beyond the indicated percentage of the VV size.

`-spt <sectors_per_track>`

Defines the VV geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is between 4 to 8192 and the default value is 304.

`-hpc <heads_per_cylinder>`

Defines the VV geometry heads per cylinder value that is reported to the hosts through the SCSI mode pages. The valid range is between 1 to 1024 and the default value is 8.

`-exp <time>[d|D|h|H]`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.

`-f`

Do not ask for confirmation before creating a volume with retention time (`-retain`).

`-retain <time>[d|D|h|H]`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.



**NOTE:** If the volume is not in any domain, then its retention time cannot exceed the value of the system's `VVRetentionTimeMax`. The default value for the system's `VVRetentionTimeMax` is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's `VVRetentionTimeMax`, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the 3PAR Virtual Lock License. Contact your local 3PAR Authorized Service Provider for more information.

- The following defines the volume policy:

`-pol <pol>[,<pol>...]`

Specifies VV policies.

`stale_ss`

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.

`no_stale_ss`

Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.

`one_host`

This constrains the export of a volume to one host or one host cluster (when cluster names can be used as a host name).

`no_one_host`

This policy should only be used when exporting a VV to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.

`tp_bzero`

Specifies that if a host write results in the allocation of a new data page that is only partially filled by the host write, then a zero-fill is performed on the unwritten portion of the data page. This ensures that the host cannot read data from deleted volumes or snapshot. The allocation page size is 16 KB. This is the default policy setting.

`no_tp_bzero`

Specifies that the zero-fill operation is bypassed on the allocation of partially written data pages.

`zero_detect`

This policy enables the InServ to scan for zeros in the incoming write data. This feature when used during physical copy to a TPVV will avoid allocating space for blocks containing zero. This feature when used with a Thin Persistence license will reclaim allocated space when zero blocks are written to the TPVV. This policy is only applicable for the base TPVV.



**NOTE:** Thin Persistence can have some performance implications during extremely busy system conditions. 3PAR recommends the zero\_detect policy only be enabled during Thin Persistence and Thin Conversion operations. The zero\_detect policy should be disabled during normal operation.

no\_zero\_detect

This policy disables the InServ to scan for zeros in the incoming write data to reclaim allocated space on the volume. This is the default policy setting.

## SPECIFIERS

<usr\_CPG>

Specifies the CPG name of the user space of the TPVV. If the CPG name is to be taken from a template, this field should be (-).

<vvname>[.<index>]

Specifies the VV name up to 31 characters in length. If the -cnt option is used, the optional decimal number <index> specifies the name of the first VV (<vvname>.<index>). The <index> is incremented by 1 for each subsequent VV. The <index> must be an integer from 0 to 999999. All VV names have the same length constraint.

<virtsize>[g|G|t|T]

Size for the VV in MB (maximum of 16T, minimum of 256 MB). The volume size is rounded up to the next multiple of 32 MB. The size should be an integer. An optional suffix (with no white space before the suffix) will modify the units to GB (g or G suffix) or TB (t or T suffix). If the size is to be taken from a template, this field should be (-).

## RESTRICTIONS

Use of this command requires a Thin Provisioning license. Contact your local 3PAR representative for information.

## EXAMPLES

The following example displays the creation of a TPVV named tppv1:

```
cli% createtppv -aw 50 -al 75 cpg1 tppv1 1g
```

In the preceding example, a 1 GB TPVV named `tpvv1` is allocated from a CPG named `cpg1` and its allocation warnings and limits are set at 50% and 75%, respectively.

The following example creates 3 VVs `vv1.2`, `vv1.3`, `vv1.4`:

```
cli% createtpvv -cnt 3 cpg1 vv1.2 1g
```

## NOTES

For this command:

- MB = 1048576 bytes
- GB = 1024 MB
- TB = 1024 GB



---

## COMMAND

createuser

## DESCRIPTION

The `createuser` command allows a Super user to create a new user account with a specified name and privilege level in a specified domain.

## SYNTAX

```
createuser [option <arg>] <user_name> <domain_name> <privilege_level>
```

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

Only one of the following options can be specified:

`-c <clear-text_password>`

Specifies the user's password in the clear-text format. The password must be six to eight characters in length.

`-e <encrypted_password>`

Specifies the user's password in the encrypted format. The password must be less than or equal to 31 characters in length.

## SPECIFIERS

`<user_name>`

Specifies the name of the user, using up to 31 characters in length. Valid characters are alphanumeric (letters and digits), a period ( . ), a dash ( - ), or an underscore ( \_ ). The first character must either be alphanumeric or an underscore for non-SSH users. To access the system via SSH, the first character of the `<user_name>` must be alphanumeric.

`<domain_name>`

Specifies the name of the domain in which the created user will belong. The domain name can be up to 31 characters in length.

<privilege\_level>

Specifies the privilege level you wish to assign to the created user. The privilege level can be specified as *Super*, *Service*, *Edit*, or *Browse*. For information about each privilege level, see the *InForm OS Concepts Guide*.

## RESTRICTIONS

- Access to all domains is required to run this command.
- Encrypted passwords are generated by the system. Therefore, you can only specify previously generated encrypted passwords. All passwords must be at least six characters long.

## EXAMPLES

The following example displays the successful creation of a new user `user1` with the clear-text password `123456`, with access to all service pools, and with edit level authority:

```
cli% createuser -c 123456 user1 testdomain edit
User created
```

## NOTES

- If no password is specified using the `-c` or `-e` options, then you will be prompted for a clear-text password.
- Verify the creation of a new user by issuing the `showuser` command. See [showuser](#) on page 22.188 for more information.



**Matched set:**

- `createvlun [options] <VV_name | VV_set> <LUN> <node:slot:port>  
<host_name>`
- `createvlun [options] <VV_name | VV_set> <LUN> <host_name>  
<node:slot:port>`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-f`

Specifies that the operation is forced and that the VLUN is created even if the specified VV has existing VLUNs. Unless the `-f` option is specified, the command asks for confirmation if a VV is already exported in a VLUN template.

`-cnt <number>`

Specifies that a sequence of VLUNs, as specified by the `number` argument, are exported to the same system port and host that is created. The `number` argument can be specified as any integer. For each VLUN created, the `.int` suffix of the `VV_name` specifier and LUN are incremented by one.

`-novcn`

Specifies that a VLUN Change Notification (VCN) not be issued after export. For direct connect or loop configurations, a VCN consists of a Fibre Channel Loop Initialization Primitive (LIP). For fabric configurations, a VCN consists of a Registered State Change Notification (RSCN) that is sent to the fabric controller.

`-ovrd`

Specifies that existing lower priority VLUNs will be overridden, if necessary. Can only be used when exporting to a specific host.

**SPECIFIERS**

`<VV_name | VV_set>`

Specifies the VV name or VV set name, using up to 31 characters in length. The volume name is provided in the syntax of `basename.int`. The `VV_set` name must start with `set :`.

<LUN>

Specifies the LUN as an integer from 0 through 16383. Alternatively <n>+ can be used to indicate a LUN should be auto assigned, but be a minimum of n, or n-m to indicate that a LUN should be chosen in the range n to m. In addition the keyword `auto` may be used and is treated as 0+.

<host\_name>

Specifies the host where the LUN is exported, using up to 31 characters.

<host\_set>

Specifies the host set where the LUN is exported, using up to 31 characters.

<N:S:P>

Specifies the system port of the virtual LUN export.

node

Specifies the system port, where the node is a number between 0 and 7.

slot

Specifies the PCI bus slot in the node, where the slot is a number between 0 and 5.

port

Specifies the port number on the FC card, where the port number between 1 and 4.

## RESTRICTIONS

None.

## EXAMPLES

The following example exports VV test on LUN 2:

```
cli% createvln test 2 testhost
```

## NOTES

- If a volume is already exported as a VLUN, you will be prompted for confirmation if a new export of the same volume is attempted.
- The host and port can both be supplied when issuing this command (matched set). This is the most restrictive access as both the host name and port must match before access to the VLUN is granted.

- Verify the creation of VLUNs by issuing the `showvln` command. See [showvln](#) on page 22.195 for more information.
- Conflicts between overlapping VLUN templates are resolved by a priority order among templates with matched set being the highest and port presents the lowest.

---

## COMMAND

`createvv`

## DESCRIPTION

The `createvv` command creates volumes that are provisioned from one or more Common Provisioning Groups (CPGs). Volumes can be fully provisioned from a CPG or can be thinly provisioned. You can optionally specify a CPG for snapshot space for fully provisioned volumes.

## SYNTAX

`createvv [options] <usr_CPG> <VV_name> [.<index>] <size>[g|G|t|T]`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command with logical disks specified.

## OPTIONS

`-templ <template_name>`

Use the options defined in template `<tname>`. The template is created using the `createtemplate` command. Options specified in the template are read-only or read/write. The read/write options may be overridden with new options at creation time but read-only options may not be overridden their default values, and all of these options are either read-only or read/write (using the `-nro` or `-nrw` options of the `createtemplate` command). If not included, the `-size` and `-usr_cpg` options are automatically treated as read-write even if the other unincluded properties are marked read-only.

`-i <ID>`

Specifies the ID of the volume. The default is the next available ID.

`-cnt <count>`

Specifies the number of identical VVs to create. This must be an integer from 1 through 999. If `<id>` is specified with `-i` option, it is used as the ID of the first VV and incremented by 1 for each subsequent volume.

`-shared`

Specifies that the system will try to share the Logical Disks (LDs) among the VVs. This option can only be used with the `-cnt` option. This option cannot be used with the `-tpvv` option.

`-wait <secs>`

If the command would fail due to the lack of clean space, the `-wait` option specifies the number of seconds to wait for the system to clean the dirty space before returning.

If `-wait 0` is issued, the command returns immediately. If this option is not used, the command will keep waiting for dirty chunklets to be cleaned if enough space will be available with the dirty chunklets cleaned. This option cannot be used with the `-tpvv` option.

`-exp <time>[d|D|h|H]`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters in length for the volume.

`-f`

Do not ask for confirmation before creating a volume with retention time (`-retain`).

`-retain <time>[d|D|h|H]`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.





**NOTE:** If the volume is not in any domain, then its retention time cannot exceed the value of the system's `VVRetentionTimeMax`. The default value for the system's `VVRetentionTimeMax` is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's `VVRetentionTimeMax`, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the 3PAR Virtual Lock License. Contact your local 3PAR Authorized Service Provider for more information.

`-spt <sectors_per_track>`

Defines the VV geometry `sectors_per_track` value that is reported to the hosts through the SCSI mode pages. The valid range is from 4 to 8192 and the default value is 304.

`-hpc <heads_per_cylinder>`

Defines the VV geometry `heads_per_cylinder` value that is reported to the hosts through the SCSI mode pages. The valid range is from 1 to 1024 and the default value is 8.

`-pol <pol>[, <pol>...]`

Specifies the VV policies.

`stale_ss`

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.

`no_stale_ss`

Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.

`one_host`

This constrains the export of a volume to one host or one host cluster (when cluster names can be used as a host name).

### `no_one_host`

This policy should only be used when exporting a VV to multiple hosts for use by a cluster-aware application, or when `port_presents` VLUNs are used. This is the default policy setting.

### `tp_bzero`

Specifies that if a host write results in the allocation of a new data page that is only partially filled by the host write, then a zero-fill is performed on the unwritten portion of the data page. This ensures that the host cannot read data from deleted volumes or snapshot. The allocation page size is 16 KB. This is the default policy setting.

### `no_tp_bzero`

Specifies that the zero-fill operation is bypassed on the allocation of partially written data pages.

### `zero_detect`

This policy enables the InServ to scan for zeros in the incoming write data. This feature when used during physical copy to a TPVV will avoid allocating space for blocks containing zero. This feature when used with a Thin Persistence license will reclaim allocated space when zero blocks are written to the TPVV. This policy is only applicable for the base TPVV.



**NOTE:** Thin Persistence can have some performance implications during extremely busy system conditions. 3PAR recommends the `zero_detect` policy only be enabled during Thin Persistence and Thin Conversion operations. The `zero_detect` policy should be disabled during normal operation.

### `no_zero_detect`

This policy disables the InServ to scan for zeros in the incoming write data to reclaim allocated space on the volume. This is the default policy setting.

### `-snp_cpg <snp_CPG>`

Specifies the name of the CPG from which the snapshot space is allocated.

### `-snp_aw <percent>`

Enables a snapshot space allocation warning. A warning alert is generated when the snapshot space of the VV exceeds the indicated percentage of the VV size.

`-snp_al <percent>`

Sets a snapshot space allocation limit. The snapshot space of the VV is prevented from growing beyond the indicated percentage of the VV size.

The following options can be used when creating Thinly Provisioned Virtual Volumes (TPVVs):

`-tpvv`

Specifies that the volume should be a TPVV.

`-usr_aw <percent>`

This option enables user space allocation warning. It specifies that a warning alert is generated when the user space of the TPVV exceeds the specified percentage of the VV size.

`-usr_al <percent>`

This option sets the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the VV size. After this size is reached, any new writes to the VV will fail.

## SPECIFIERS

`<usr_CPG>`

Specifies the name of the CPG from which the volume user space will be allocated. If the `-tpvv` option is specified, the volume is thinly provisioned. Otherwise, the volume is fully provisioned from the specified CPG. If the name of the CPG is to be taken from a template, this field should be (-).

`<VV_name> [.<index>]`

Specifies a VV name up to 31 characters in length. If the `-cnt` option is used, the optional decimal number `<index>` specifies the name of the first VV (`<vvname>.<index>`).

The `<index>` is incremented by 1 for each subsequent VV. The `<index>` must be an integer from 0 through 999999. All VVs have the same length constraint.

`<size>[g|G|t|T]`

Specifies the size for the volume in MB. The volume size is rounded up to the next multiple of 256 MB. The size should be an integer. An optional suffix (with no whitespace before the suffix) will modify the units to GB (g or G suffix) or TB (t or T suffix). If the size is to be taken from a template, this field should be (-).

## RESTRICTIONS

- Access to all domains is required to to run this command with logical disks specified.
- Use of this command with the `-tpvv` option requires a Thin Provisioning license. Contact your local 3PAR representative for information.
- The `-templ` option is not valid for volumes that are fully provisioned.
- The `-shared` and `-wait` option cannot be used when creating Thinly Provisioned Virtual Volumes with the `-tpvv` option.

## EXAMPLES

The following example creates a 10G TPVV named `tpvv1` whose user space is allocated from the common provisioning group `cpgr1`:

```
cli% createvv -tpvv -usr_aw 50 -usr_al 75 cpgr1 tpvv1 10G
```

The following example creates 3 VVs `vv1.2`, `vv1.3`, `vv1.4`:

```
cli% createvv -cnt 3 cpgr1 vv1.2 1g
```

The following example creates 2 VVs `vva.0` and `vva.1` which may share LDs:

```
cli% createvv -cnt 3 cpgr1 vv1.2 1g
```

## NOTES

- Using the `-shared` option can be more efficient when creating a large number of small volumes since the system will attempt to share the underlying LDs. However, if volumes are subsequently removed, this can lead to unused space in the LDs that cannot be reclaimed. Using `-cnt` option without the `-shared` option avoids this problem, but can be inefficient when creating small volumes due to the overhead of creating a larger number of unshared LDs for the volumes.
- Verify the creation of VV by issuing the `showvv` command. See [showvv](#) on page 22.201 for more information.
- The maximum VV size is 16 Terabytes.

- For this command:
  - ◆ MB = 1048576 bytes
  - ◆ GB = 1024 MB
  - ◆ TB = 1024 GB



**-halt**

Specifies that an ongoing physical copy or snapshot promotion should be stopped. This will cause the destination volume to be marked with the `copy` status, which will be cleaned up when it is promoted to a base volume or when a new copy is started.

**-s**

Saves the snapshot of the source volume after the copy of the volume is completed. This enables a fast copy for the next resynchronization. If not specified, the snapshot is deleted and the association of the destination volume as a copy of the source volume is removed. The `-s` option is implied when the `-r` option is used and need not be explicitly specified.

**-b**

Specifies that this command blocks until the operation is completed. If not specified, the `createvvcopy` command operation is started as a background task.

**SPECIFIERS**

`<destination_volume | destination_volumeset>`

Specifies the destination volume name or volume set name for the copy operation using up to 31 characters in length. The volume set name must start with `set:`. The destination volume (or each member of the destination volume set) must be a writable base volume (not a snapshot) of equal or greater size than a parent volume (if specified) and it must not be exported as a VLUN.

**RESTRICTIONS**

- The source and destination volumes must be writable.
- The destination volume cannot be exported as a VLUN before or during the `createvvcopy` command process.
- The destination volume must be greater than or equal in size to the source volume.
- If both the destination and source are volume sets then they must contain the same number of members.
- If the source is a volume set then the destination name can contain various patterns that are used to form the copy volume name. See `Help sub, vvnamepat` for details. Note that the time/date related patterns are not supported in this command.
- If the `-s` option is specified to save a snapshot for fast resynchronization and the snapshot goes stale, the copy fails.

- A physical copy of a VV fails in any situation that a snapshot copy fails or when there is insufficient snapshot space or I/O errors.

## EXAMPLES

The following example displays the creation of a copy of the VV `vv1`:

```
cli% createvvcopy -p vv1 vv2
Started copy. child=vv2 parent=vv1
```

Create a set of copys for the volumes in set `vvcopyset`, keeping snapshots around for quick resynchronization:

```
cli% createvvcopy -s -p set:vvcopyset set:copys
Child Parent Status TaskID
s1      foo      started    14
s2      bar      started    15
s3      baz      queued     16
```

Resynchronise the volumes in the volume set `copys`:

```
createvvcopy -r set:copys
Child Parent Status TaskID
s1          started    17
s2          started    18
s3          queued     19
```

Create a set of copys for the volumes in set `vvcopyset`, mapping the parent volumes to destination volumes with a pattern:

```
cli% createvvcopy -p set:vvcopyset @vvname@-copy
Child      Parent Status TaskID
foo-copy   foo      started    20
bar-copy   bar      started    21
baz-copy   baz      queued     22
```

## NOTES

- The `createvvcopy` command can be issued multiple times. However, the InServ system allows only two active physical copy tasks to run concurrently. Any additional physical copy tasks are queued, pending the completion of the active physical copy tasks.



- Multiple physical copy operations can occur simultaneously. Host-initiated I/O operations and those operations initiated by issuing the `createvvcopy` command are executed at the same priority level. As a result, noticeable performance degradation from a host perspective can be observed.
- Issuing the `createvvcopy` command results in the creation of a temporary snapshot and, in the case of `-r`, a resynchronization (resync) snapshot.
  - ◆ The temporary and resynchronization snapshots cannot be deleted while the copy is in progress.
  - ◆ Upon completion of the copy, the temporary snapshot is automatically deleted if the `-s` option is not specified.
  - ◆ If the resynchronization snapshot is saved, it can later be manually deleted. If the resynchronization snapshot is deleted, later resynchronization is not possible.
- If the `-s` option is not specified, the relationship between the destination volume and source volume is not retained.
- Issue the `showvv` command to verify that a virtual volume copy has been made.
- Issue the `showvv -d` command to display the number of blocks remaining to be copied.
- If the source or destination are volume sets then the `createvvcopy` command creates consistent group snapshots at the same point in time so that related structures on different volumes in the group remain consistent.
- On systems that support `zero_detect`, if destination volume is a TPVV, `zero_detect` policy is enabled at start of physical copy. After physical copy completes, the policy is changed to `no_zero_detect`, even when the destination TPVV had `zero_detect` policy enabled before the start of physical copy operation.
- When Thin Persistence license is present, `createvvcopy` command will reclaim allocated space when zero blocks are written to the destination TPVV.

---

## COMMAND

`createvvset`

## DESCRIPTION

The `createvvset` command defines a new set of Virtual Volumes (VV) and provides the option of assigning one or more existing VVs to that set. The command also allows the addition of VVs to an existing set by use of the `-add` option.

## SYNTAX

```
createvvset [options] <setname> [<VV | pattern>...]
```

## AUTHORITY

Super, Edit

## OPTIONS

`-add`

Specifies that the VVs listed should be added to an existing set. At least one VV must be specified.

`-cnt <num>`

Add a sequence of `<num>` VVs starting with `vvname`. The `vvname` should be of the format `<basename>.<int>`. For each VV in the sequence, the `.<int>` suffix of the `vvname` is incremented by 1.

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters in length. Unprintable characters are not allowed.

`-domain <domain>`

Create the VV set in the specified domain. For an empty set the default is to create it in the current domain, or no domain if the current domain is not set. A VV set must be in the same domain as its members; if VVs are specified as part of the creation then the set will be created in their domain. A domain cannot be specified when adding a VV to an existing set with the `-add` option.

## SPECIFIERS

<setname>

Specifies the name of the VV set to create or add to. The name can be up to 27 characters in length.

<VV | pattern>...

Optional lists of VVs or glob-style pattern matching VVs which are to be included in the VV set. If no <VV> or <pattern> is specified a VV set with no VVs is created. If the `-cnt` option is used then a VV must be specified.

## RESTRICTIONS

None.

## EXAMPLES

To create an empty vvset:

```
cli% createvvset
```

To add a VV to the set:

```
cli% createvvset -add vvset vv1
```

To create a VV set with a comment and a collection of VVs in it:

```
cli% createvvset -comment "Our Oracle VVs" -cnt 10 oravv oravv.0
```

## NOTES

None.



# 12

## Dismiss Commands

---

In this chapter

<code>dismisspd</code>	<b>12.2</b>
<code>dismissrcopylink</code>	<b>12.3</b>
<code>dismissrcopytarget</code>	<b>12.5</b>
<code>dismissrcopyvv</code>	<b>12.6</b>

---

**COMMAND**

`dismisspd`

**DESCRIPTION**

The `dismisspd` command removes Physical Disk (PD) definitions from system use.

**SYNTAX**

`dismisspd <PD_ID>...`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

None

**SPECIFIERS**

`<PD_ID>...`

Specifies the PD(s), identified by integers, to be removed from system use.

**RESTRICTIONS**

- Access to all domains is required to run this command.
- A PD that is in use cannot be removed.

**EXAMPLES**

The following example removes a physical disk with ID 1:

```
cli% dismisspd 1
```

**NOTES**

Verify the removal of a physical disk by issuing the `showpd` command. See [showpd](#) on page 22.91 for more information.

---

## COMMAND

`dismissrcopylink`

## DESCRIPTION

The `dismissrcopylink` command removes one or more links (connections) to a target system.

## SYNTAX

- Syntax for Remote Copy over IP (RCIP) is as follows:

```
dismissrcopylink <target_name> <N:S:P:IP_address>...
```

- Syntax for Remote Copy over Fibre Channel (RCFC) is as follows:

```
dismissrcopylink <target_name> <N:S:P:WWN>...
```

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

None.

## SPECIFIERS

`<target_name>`

Specifies the name of the target from which the link(s) will be dismissed (See [creatercopytarget](#) on page 11.50).

`<N:S:P:IP_address>...`

Specifies the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.

`<N:S:P:WWN>...`

Specifies the node, slot, and port of the Fibre Channel port on the local system and a World Wide Name (WWN) of the peer port on the target system. This specifier can be repeated.

## RESTRICTIONS

- Access to all domains is required to run this command.
- Functionality of this command requires a 3PAR Remote Copy license. Contact your local service provider for more information.
- This command cannot be used to remove the last link of a target system with started groups.
- This command terminates with a list of one or more links to be dismissed.

## EXAMPLES

The following example removes the link from node 1 of System2:

```
cli% dismissrcopylink System2 1:193.1.2.11
```

## NOTES

- See the *Remote Copy User's Guide* for more examples and for recommended link naming conventions.
- IP targets are made up of pairs composed of the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.
- FC targets are made up of sets with the node, slot, and port of the fibre channel port on the local system and a WWN of the peer port on the target system.



---

**COMMAND**

`dismissrcopytarget`

**DESCRIPTION**

The `dismissrcopytarget` command removes a Remote Copy target from a Remote Copy volume group.

**SYNTAX**

`dismissrcopytarget <target_name> <group_name>`

**AUTHORITY**

Super, Edit

**OPTIONS**

None.

**SPECIFIERS**

`<target_name>`

Specifies the name of the target to be removed.

`<group_name>`

The name of the group that currently includes the target.

**RESTRICTION**

The `dismissrcopytarget` command cannot be used to remove the last target.

**EXAMPLES**

The following example removes target Target1 from Group1:

```
cli% dismissrcopytarget Target1 Group1
```

**NOTES**

None.

---

**COMMAND**

`dismissrcopyvv`

**DESCRIPTION**

The `dismissrcopyvv` command removes a Virtual Volume (VV) from a Remote Copy volume group.

**SYNTAX**

`dismissrcopyvv <VV_name> <group_name>`

**AUTHORITY**

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

None.

**SPECIFIERS**

`<VV_name>`

The name of the volume to be removed. Volumes are added to a group with the `admitrcopyvv` command.

`<group_name>`

The name of the group that currently includes the VV.

**RESTRICTIONS**

- This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- A volume cannot be removed from a group that is not currently stopped.

**EXAMPLES**

The following example removes virtual volume `vv1` from `Group1`:

```
cli% dismissrcopyvv vv1 Group1
```

## NOTES

- The `dismissrcopyvv` command removes any Remote Copy synchronization snapshots affiliated with the removed volume.
- If a group's target has the `mirror_config` policy set and the group is a primary group, then this command is mirrored to that target and the volume is removed from the corresponding secondary group. If the policy is set and the group is a secondary, then this command fails.



# 13

## Free Command

---

In this chapter

`freespace`

**13.2**

---

**COMMAND**

freespace

**DESCRIPTION**

The `freespace` command frees snapshot administration and snapshot data spaces from a Virtual Volume (VV) if they are not in use.

**SYNTAX**

`freespace [options] <VV_name>...|<pattern>...`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-pat`

Remove the snapshot administration and snapshot data spaces from all the VV that match any of the specified glob-style patterns.

`-f`

Suppresses the prompt for confirmation that appears before removing the snapshot administration and snapshot data space of each volume.

**SPECIFIERS**

`<VV_name>...`

Specifies the VV name, using up to 31 characters.

`<pattern>...`

Specifies a glob-style pattern. This specifier can be repeated to compact multiple volumes. If this specifier is not used, the `VV_name` specifier must be used. See [Glob-Style Pattern](#) on page 2.4 for more information.

**RESTRICTIONS**

None.

## EXAMPLES

The following example demonstrates how to remove SA and SD space from virtual volume testd:

```
cli% freespace testd
```

## NOTES

This command fails if the virtual volume is an old-style thin provisioned VV (created on a 2.2.4 release or earlier) or it has snapshots.





# 14

## Grow Commands

---

In this chapter

<code>growaldvv</code>	<b>14.2</b>
<code>growavv</code>	<b>14.10</b>
<code>growtpvv</code>	<b>14.12</b>
<code>growvv</code>	<b>14.14</b>

---

## COMMAND

growaldivv

## DESCRIPTION

This command was deprecated in the 2.2.4 release and will be changed or removed in a future release. Please use the `growvv` command in the future. If virtual volumes were created with deprecated commands then only deprecated commands can be used to modify virtual volumes.

The `growaldivv` command automatically enlarges (grows) the size of a virtual volume by adding underlying logical disks.

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## SYNTAX

The syntax for the `growaldivv` command can be any of the following:

- `growaldivv -szu <size> [options <arg>] <VV_name>`
- `growaldivv -szs <size> [options <arg>] <VV_name>`
- `growaldivv -sza <size> [options <arg>] <VV_name>`
- `growaldivv -szu <size> -szs <size> [options <arg>] <VV_name>`
- `growaldivv -szu <size> -sza <size> [options <arg>] <VV_name>`
- `growaldivv -szu <size> -szs <size> -sza <size> [options <arg>] <VV_name>`

## OPTIONS

`-szu <size>[g|G|t|T]`

Specifies the size of the user volume in megabytes using an integer from 0 through 16777216 (16384 GB). Size can optionally be set in gigabytes or terabytes by providing either the `g` or `G` options (for gigabytes) and `t` or `T` options (for terabytes) following (no whitespace between) the entered size value.

`-szs <size>[g|G|t|T]`

Specifies the size of the snapshot volume in megabytes using an integer between 0 and 67108864 (65536 GB). Size can optionally be specified in gigabytes or terabytes by providing either the `g` or `G` options (for gigabytes) and `t` or `T` options (for terabytes) following (no whitespace between) the entered size value. This option cannot be used with the `-pct` option. This option is deprecated and will be removed in a subsequent release.

`-sza <size>[g|G]`

Specifies the size of the administration volume in megabytes using an integer between 0 and 524288 (512 GB). Size can optionally be specified in gigabytes by providing either the `g` or `G` options following (no whitespace between) the entered size value. This option cannot be used with the `-pct` option. This option is deprecated and will be removed in a subsequent release.

`-f`

Suppresses the requested confirmation before growing a virtual volume with a different RAID type than the last region of the existing virtual volume.

`-t <RAID_type>`

Specifies the RAID type with `r0`, `r1`, `r5` or `r6`. If not specified, the default RAID type is the same as the last region of the existing virtual volume. Specifying a different RAID type than the existing virtual volume results in a warning message and a prompt for confirmation unless the `-f` option is specified.

`-ssz <size_number_chunklet>`

Specifies the set size in terms of chunklets. The default depends on the RAID type specified. If the RAID type is not specified or if the same RAID type as the last region of the existing virtual volume is specified, the default is the same set size as the last region of the existing VV. Otherwise, the default is 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row for each logical disk using an integer between 1 and 2147483647. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size in kilobytes using 32, 64, 128, 256, or 512. The default depends on whether the RAID type is specified. If the RAID type is not specified or is the same RAID type as the last region of the existing virtual volume, the default is the same step size as the last region of the virtual volume. Otherwise, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6 the step size is a function of the set size.

`-ha <port> | <cage> | <mag>`

Specifies that the RAID layout must support a failure of one port pair, one cage, or one mag. The default depends on whether the RAID type is specified. If the RAID type is not specified or if the same RAID type as the last region of the existing virtual volume is specified, the default is the same as the last region of the existing virtual volumes. Otherwise the default is cage. This option has no meaning for RAID-0. Note that snap admin space is always created with cage availability.

`-ch first | last`

Specifies the chunklet location preference characteristics, either first (the lowest numbered chunklets) or last (the highest numbered chunklets). If no argument is specified, the default characteristic is `first`.

`-pct <prc>`

Specifies the required growth size of the snapshot volume as a percentage of the required growth size of your volume (as specified with the `-szu` option). This option must be used with the `-szu` option, and cannot be used with the `-szs` option. If not specified, the default value is zero percent of your user volume.

`-wait <secs>`

If the `growaldivv` command fails due to the lack of clean space, the `-wait` option specifies the number of seconds to wait for the system to clean the dirty space before returning. If `-wait 0` is issued, the command returns immediately. If this option is not used, the command will keep waiting for dirty chunklets to be cleaned if enough space will be available with the dirty chunklets cleaned.

`-dr`

Specifies that the command is a dry run and that no logical disks or virtual volumes are created.

`-verbose on | off`

Specifies that verbose output is either enabled (`on`) or disabled (`off`). If not specified, verbose output is disabled.

`-p <pattern>`

Specifies a pattern for disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to all disks of device type specified at creation time. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The following arguments can be specified as patterns for this option:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a comma (1, 2, 3). A range of nodes is separated with a hyphen (0–7). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are each separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0–7). The primary path of the disks must be on the specified PCI slot(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are each separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–4). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are each separated with a comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-7). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are each separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are each separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks depending on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devId <ID>`

Specifies that physical disks identified by their device IDs be selected for logical disk creation. Device IDs can be specified in a comma-separated list. Device IDs can be displayed by issuing the `showpd -i` command.

`-devtype <model>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the *Speed* column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.



**NOTE:** Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

`<VV_name>`

Specifies the name of the virtual volume, using up to 31 characters, that you are enlarging.

## RESTRICTIONS

- Access to all domains is required to run this command.
- Option `-pct` can only be used if option `-szu` is used.

- Options `-pct` and `-szs` cannot be used in the same instance of issuing the `growalddv` command.
- One or more of the `-szu`, `-sza`, and `-szs` options must be specified.
- If enlarging either the administration or snapshot space on the VV and both administration and snapshot space are zero MB, the administration and snapshot space must be enlarged together. Both, the `-sza` and `-szs` options must be specified on the command line. Specifying `-sza` or `-szs` alone is invalid.
- If the volume being enlarged has snapshot data space and snapshot administration space with sizes of 0 bytes, the `-sza` and `-szs` options must be specified together.
- This command cannot be used to grow a volume in such a way that a volume that currently only has logical disks of a particular device type (Fibre Channel, Nearline, or Solid State Drive) will have logical disks of multiple device types after the grow operation. If a virtual volume currently has logical disks of multiple device types, it is possible to grow the volume using logical disks of any device type. Use `showpd` to see the device types of physical disks in the system.

## EXAMPLES

The following example displays the enlarging of a RAID-5 virtual volume with 3+1 parity ratio named `vv01` by 10 GB:

```
cli% growalddv -szu 10g -t r5 -ssz 4 vv01
```

## NOTES

- A volume can be enlarged while I/O is occurring.
- When issuing the `growalddv` command, the specified options only apply to the newly enlarged sections of the volume. This can result in a VV that has different characteristics in its base and grown sections.
- If options were used in the creation of the original virtual volume (`createalddv`) that were not the default option values, and not specified in the same manner when issuing the `growalddv` command, `growalddv` uses the `createalddv` default option values. This can result in a virtual volume that has different characteristics in its base and grown sections.
- By default, LDs are created using only physical disks with the same device type (by default, the Fibre Channel device type is used). Use the `-p devtype NL` option or `-p -devtype SSD` to override this default. Use the `showpd` command to see the device types of PDs in the system.



- If no device type is specified using the `-p -devtype` option, Fibre Channel is assumed.
- This command is only used for a non-provisioned VV. The command `growvv` should be used instead for a virtual volume that has its user space provisioned from a CPG.
- For this command:
  - ◆ KB = 1024 bytes
  - ◆ MB = 1024 KB
  - ◆ GB = 1024 MB

---

## COMMAND

growavv

## DESCRIPTION

This command was deprecated in the 2.2.4 release and will be changed or removed in a future release. Please use the `growvv` command in the future. If virtual volumes were created with deprecated commands then only deprecated commands can be used to modify virtual volumes.

The `growavv` command enlarges a Virtual Volume (VV) with automatic mapping to the Logical Disks (LDs).

## SYNTAX

```
growavv [options] <VV_name> <user_LD_ID> <LD_size> <sa_LD_ID>  
<admin_size> <sd_LD_ID> <snap_size>
```

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-verbose on|off`

Specifies that verbose output is either enabled (`on`) or disabled (`off`). If not specified, verbose output is disabled.

`-dr`

Specifies that the command is a dry run and that no logical disks or virtual volumes are actually created.

## SPECIFIERS

`<VV_name>`

Specifies the virtual volume name, using up to 31 characters.

<user\_LD\_ID>

Specifies the LDs to be used as user space. The LDs are identified by one or more integers (*item*). Integers can be provided as a single number (1), a comma separated list of numbers (1, 2, 3), or a range of numbers separated with a dash (1–4).

<LD\_size>

Specifies the size of your area in megabytes.

<sa\_LD\_ID>

Specifies the LDs to be used as snapshot administrator space. The LDs are identified by one or more integers (*item*). Integers can be provided as a single number (1), a comma separated list of numbers (1, 2, 3), or a range of numbers separated with a dash (1–4).

<admin\_size>

Specifies the size of the administrator space in megabytes.

<sd\_LD\_ID>

Specifies the LDs to be used as snapshot data space. The LDs are identified by one or more integers (*item*). Integers can be provided as a single number (1), a comma separated list of numbers (1, 2, 3), or a range of numbers separated with a dash (1–4).

<snap\_size>

Specifies the size of the snapshot space in megabytes.

## RESTRICTIONS

None.

## EXAMPLES

None.

## NOTES

- For this command:
  - ◆ MB = 1048576 bytes

---

**COMMAND**

growtpvv

**DESCRIPTION**

This command was deprecated in the 2.2.3 release and will be changed or removed in a future release. Please use the `growvv` command in the future. If virtual volumes were created with deprecated commands then only deprecated commands can be used to modify virtual volumes.

The `growtpvv` command enlarges a Thinly Provisioned Virtual Volume (TPVV).

**AUTHORITY**

Super, Edit

**SYNTAX**

growtpvv <size\_increment>[g|G|t|T] <TPVV\_name>

**OPTIONS**

<size\_increment>[g|G|t|T]

Specifies the size in MB to be added to the TPVV user size. Should be in the range 1 to 16T. The volume size is rounded up to the next multiple of 32 MB. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space).

**SPECIFIERS**

<TPVV\_name>

Specifies the TPVV's name, using up to 31 characters.

**RESTRICTIONS**

None.

**EXAMPLES**

None.

**NOTES**

For this command:

- MB = 1048576 bytes
- GB = 1024 MB
- TB = 1024 GB

---

## COMMAND

growvv

## DESCRIPTION

The `growvv` command increases the size of a Virtual Volume (VV) created with the `createvv` or `createtpvv` commands.

## SYNTAX

```
growvv <VV_name> <size>[g|G|t|T]
```



**CAUTION:** For Remote Copy configuration, before issuing the `growvv` command on the primary InServ, ensure the secondary InServ has sufficient space for an equal increase in Virtual Volume size. If there is insufficient space on the secondary InServ, the volume group cannot be started.

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains with logical disks specified in order to run this command.

## OPTIONS

None.

## SPECIFIERS

<VV\_name>

Specifies the name of the VV to grow.

<size>[g|G|t|T]

Specifies the size in MB to be added to the volume user space. The size must be an integer in the range from 1 to 16 TB. The volume size is rounded up to the next multiple of 256 MB. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space).

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the enlarging of the of virtual volume `vv0` by 10 G:

```
cli% growvv vv0 10g
```

## NOTES

- For the command:
  - ◆ MB = 1048576 bytes
  - ◆ GB = 1024 MB
  - ◆ TB = 1024 GB
- For Remote Copy configuration after issuing the `growvv` command, the affected Virtual Volume cannot be reduced back to the original size on the primary InServ.





# 15

## Hist Commands

---

In this chapter

histch	<b>15.2</b>
histld	<b>15.6</b>
histpd	<b>15.10</b>
histport	<b>15.18</b>
histvlun	<b>15.23</b>
histvv	<b>15.28</b>

---

**COMMAND**

histch

**DESCRIPTION**

The `histch` command displays a histogram of service times in a timed loop for individual chunklets.

**SYNTAX**

`histch [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-ld <LD_name>`

Specifies the Logical Disk (LD), identified by name, from which chunklet statistics are sampled.

`-ch <chunklet_num>`

Specifies that statistics are limited to only the specified chunklet, identified by number.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- ◆ `both` - (Default) Displays both I/O time and I/O size histograms.
- ◆ `time` - Displays only the I/O time histogram.
- ◆ `size` - Displays only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, show the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`) (default value of 6).
- ◆ The last column (`<lcol>`) must be less than or equal to 31 (default value of 15).

- ◆ The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, show the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

- ◆ The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).
- ◆ The last column (<lcol>) must be less than or equal to 15 (default value of 11).

`-pct`

Shows the access count in each bucket as a percentage.



**NOTE:** For the following options, by default the histogram shows data from the start of the command.

`-prev|-begin`

Specifies that the histogram displays data either from a previous sample (`-prev`) or from when the system was last started (`-begin`). If no option is specified, the histogram shows data from the beginning of the command's execution.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col> [,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, you can specify the direction of sorting (<dir>) as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple can specify multiple columns on the command line by using a colon (:). Rows having the same column(s) are sorted by the values in the latter specified columns.

`-filt <fspec>`

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>,<count>`.

`<op>`

The `<op>` argument can be specified as one of the following:

- ◆ `r` - Specifies read statistics.
- ◆ `w` - Specifies write statistics.
- ◆ `t` - Specifies total statistics.
- ◆ `rw` - Specifies total read and write statistics.

`<val_ms>`

Specifies the threshold service time in milliseconds.

`<count>`

Specifies the minimum number of access above the threshold service time.

`-ni`

Specifies that histograms for only nonidle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays one iteration of a histogram of service times for system chunklets:

```
cli% histch -iter 1
12:42:57 10/20/04 -----Time (millisec)-----
Ldid      Ldname  LdCh  Pdid  PdCh  0.26  0.53  1.05  2.1  4.2  8.4  17  34  67  135
2  tp-0-sa-0.0    1   21    0    0    0    0    0    0    0    0    0    0    0
2  tp-0-sa-0.0    0   45    0    0    0    0    0    0    0    0    0    0    0
1  tp-0-sa-0.1    1   20    0    0    0    0    0    0    1    0    0    0    0
1  tp-0-sa-0.1    0   42    0    0    0    0    0    0    1    0    0    0    0
0  admin.usr.0    1   22    3    0    0    0    1    0    3    2    0    0    0
0  admin.usr.0    0   44    3    0    0    0    0    1    3    2    0    0    0
-----
total                0    0    0    1    1    8    4    0    0    0
```

For the previous example, before the `histch` command was issued, the `setstatch start` command was issued for chunklets 0 and 1 on logical disks `admin.usr.0`, `tp-0-sa-0.1`, and `tp-0-sa-0.0`.

## NOTES

- The `setstatch` command must be issued to enable statistics collection on chunklets before issuing the `histch` command. See [setstatch](#) on page 21.72 for additional information.
- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column).
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).

---

**COMMAND**

histld

**DESCRIPTION**

The `histld` command displays a histogram of service times for Logical Disks (LDs) in a timed loop.

**SYNTAX**

`histld [options <arg>] [<LD_name_or_pattern>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-vv <VV_name>...|<pattern>...`

Show only LDs that are mapped to Virtual Volumes (VVs) with names that match any of the names or patterns specified. Multiple volumes or patterns can be repeated using a comma separated list.

`-domain <domain_name>...|<pattern>...`

Shows only LDs that are in domains with names that match any of the names or patterns specified. Multiple volumes or patterns can be repeated using a comma separated list.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- ◆ `both` - (Default) Displays both I/O time and I/O size histograms.
- ◆ `time` - Displays only the I/O time histogram.
- ◆ `size` - Displays only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, shows the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`) (default value of 6).
- ◆ The last column (`<lcol>`) must be less than or equal to 31 (default value of 15).

- ◆ The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, show the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

- ◆ The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).
- ◆ The last column (<lcol>) must be less than or equal to 15 (default value of 11).

`-pct`

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.



**NOTE:** For the following options, by default the histogram shows data from the start of the command.

`-prev|-begin`

Histogram displays data either from a previous sample (`-prev`) or from when the system was last started (`-begin`). If no option is specified, the histogram shows data from the beginning of the command's execution.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col> [ ,<dir>][:<col>[ ,<dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting `<dir>` as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

`-filt <fspec>`

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>,<count>`.

`<op>`

The `<op>` argument can be specified as one of the following:

- ◆ `r` - Specifies read statistics.
- ◆ `w` - Specifies write statistics.
- ◆ `t` - Specifies total statistics.
- ◆ `rw` - Specifies total read and write statistics.

`<val_ms>`

Specifies the threshold service time in milliseconds.

`<count>`

Specifies the minimum number of access above the threshold service time.

`-ni`

Specifies that histograms for only nonidle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

## SPECIFIERS

`[<LD_name_or_pattern>]...`

Specifies the LD(s) or pattern(s) for which the histogram data is collected.

## RESTRICTIONS

None.



## EXAMPLES

The following example displays one iteration of a histogram of service times for all LDs:

```
cli% histld -iter 1
12:38:49 10/20/04 -----Time (millisec)-----
      Ldame 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
tp-0-sa-0.0    0    0    0    0    0    0    0    0    0    0
tp-0-sd-0.1    0    0    0    0    0    0    0    0    0    0
tp-0-sd-0.0    0    0    0    0    0    0    0    0    0    0
admin.usr.0    0    0    0    0    1    6    6    1    0    0
tp-0-sa-0.1    0    0    0    0    0    0    0    0    0    0
tp-0-sd-0.3    0    0    0    0    0    0    0    0    0    0
tp-0-sd-0.2    0    0    0    0    0    0    0    0    0    0
-----
      total    0    0    0    0    1    6    6    1    0    0
```

## NOTES

- Each service time histogram column shows the number of accesses with service times between the time shown in the column's heading to the time shown in the heading of the column to its right (or infinity for the last column).
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- If the <LD\_name\_or\_pattern> specifier is used, then logical disks with names that match any of the patterns are listed, otherwise all LDs are listed. These patterns are glob-style patterns (see help on `sub,globpat`).
- Patterns are specified as regular expressions. See CLI Help on `sub,regexpat` for help on regexps. Issuing `histld -n LD_name.*` displays histogram data for all LDs whose name begins with `LD_name`.

---

**COMMAND**

histpd

**DESCRIPTION**

The `histpd` command displays a histogram of service times for Physical Disks (PDs).

**SYNTAX**

`histpd [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-w <WWN>`

Specifies the World Wide Name (WWN) of the physical disk for which service times are displayed.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified port slots and physical disks connected to those port slots. The port list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the port list is not specified, all disks on all ports are displayed.

`-devinfo`

Indicates the device disk type and speed.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- ◆ `both` - (Default) Displays both I/O time and I/O size histograms.
- ◆ `time` - Displays only the I/O time histogram.
- ◆ `size` - Displays only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, shows the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`). The default value of `<fcol>` is 6.
- ◆ The last column (`<lcol>`) must be less than or equal to 31. The default value of `<lcol>` is 15.
- ◆ The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, show the columns from the first column (`<fcol>`) through the last column (`<lcol>`). Available columns range from 0 through 15.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`). The default value of `<fcol>` is 3.
- ◆ The last column (`<lcol>`) must be less than or equal to 15. The default value of `<lcol>` is 15.

`-pct`

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.



**NOTE:** For the following options, by default the histogram shows data from the start of the command.

`-prev|-begin`

Histogram displays data either from a previous sample (`-prev`) or from when the system was last started (`-begin`). If no option is specified, the histogram shows data from the beginning of the command's execution.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col> [,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting `<dir>` as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

`-filt <fspec>`

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>,<count>`.

`<op>`

The `<op>` argument can be specified as one of the following:

- ◆ `r` - Specifies read statistics.
- ◆ `w` - Specifies write statistics.
- ◆ `t` - Specifies total statistics.
- ◆ `rw` - Specifies total read and write statistics.

`<val_ms>`

Specifies the threshold service time in milliseconds.

`<count>`

Specifies the minimum number of access above the threshold service time.

`-ni`

Specifies that histograms for only nonidle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

`-p <pattern>`

Specifies the pattern of PDs displayed in the histogram. Patterns are used to filter and select the disks displayed in the histogram. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (`<item>`). Multiple nodes are separated with a single comma (1,2,3). A range of nodes is separated with a hyphen (0-7). The primary path of the disks must be on the specified node number(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (`<item>`). Multiple slots are separated with a single comma (1,2,3). A range of slots is separated with a hyphen (0-7). The primary path of the disks must be on the specified PCI slot number(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (`<item>`). Multiple ports are separated with a single comma (1,2,3). A range of ports is separated with a hyphen (0-5). The primary path of the disks must be on the specified port number(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (<item>). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0-3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (<item>). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-3). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (<item>). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (<item>). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devId`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Device) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd -i` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays a histogram of service times for all physical disks:

```
cli% histpd
12:36:53 10/20/04 -----Time (millisec)-----
Pdid      Port 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
  0      0:1:1    0    0    0    0    0    1    0    0    0    0
  1      0:1:1    0    0    0    0    0    0    0    0    0    0
  2      1:5:2    0    0    0    0    1    1    0    0    0    0
  3      0:1:1    0    0    0    0    0    0    0    0    0    0
  4      1:5:2    0    0    0    1    1    0    0    0    0    0
  5      0:1:1    0    0    0    0    0    4    0    0    0    0
  6      1:5:2    0    0    0    0    0    0    0    0    0    0
  7      0:1:1    0    0    0    0    1    0    0    0    0    0
  8      1:5:2    0    0    0    0    0    2    1    0    0    0
  9      0:1:1    0    0    0    0    0    2    4    0    0    0
 10      1:5:2    0    0    0    0    1    1    1    0    0    0
 11      0:1:1    0    0    0    0    0    0    0    0    0    0
  (... )
 32      1:5:1    0    0    0    1    0    2    0    0    0    0
 33      0:1:2    0    0    0    0    1    4    1    0    0    0
 34      1:5:1    0    0    0    0    2    0    0    0    0    0
 35      0:1:2    0    0    0    0    0    0    0    0    0    0
 36      1:5:1    0    0    0    1    1    0    0    0    0    0
 37      0:1:2    0    0    0    0    1    2    0    0    0    0
 38      1:5:1    0    0    0    0    1    1    2    0    0    0
 39      0:1:2    0    0    0    0    0    0    0    0    0    0
 40      1:5:1    0    0    0    0    2    4    5    3    0    0
 41      0:1:2    0    0    0    0    0    0    0    0    0    0
 42      1:5:1    0    0    0    1    0    1    0    0    0    0
 43      0:1:2    0    0    0    0    1    0    0    0    0    0
 44      1:5:1    0    0    0    0    2    3    2    0    0    0
 45      0:1:2    0    0    0    0    0    1    0    0    0    0
-----
total          0    0    0    9   30   48  25    5    0    0
Press the enter key to stop...
```

## NOTES

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column).
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).



- Previous options `-n <node_list>`, `-s <slot_list>`, and `-p <port_list>` have been replaced with `-nodes <node_list>`, `-slots <slot_list>`, and `-ports <port_list>`.

---

**COMMAND**

histport

**DESCRIPTION**

The `histport` command displays a histogram of service times for ports within the system.

**SYNTAX**

`histport [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-both|-ctl|-data`

Specifies that both control and data transfers are displayed (`-both`), only control transfers are displayed (`-ctl`), or only data transfers are displayed (`-data`). If this option is not specified, only data transfers are displayed.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and Physical Disks (PDs) connected to those nodes. The node list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and PDs connected to those PCI slots. The slot list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the slot list is not specified, all disks on all nodes are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified ports and PDs connected to those ports. The port list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the port list is not specified, all disks on all nodes are displayed.

`-host|disk|-rcfc`

Specifies to display only host ports (target ports), only disk ports (initiator ports), or only Fibre Channel Remote Copy configured ports. If no option is specified, all ports are displayed.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- ◆ `both` - (Default) Displays both I/O time and I/O size histograms.
- ◆ `time` - Displays only the I/O time histogram.
- ◆ `size` - Displays only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, show the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`) (default value of 6).
- ◆ The last column (`<lcol>`) must be less than or equal to 31 (default value of 15).
- ◆ The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, show the columns from the first column (`<fcol>`) through the last column (`<lcol>`). Available columns range from 0 through 15.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`) (default value of 3).
- ◆ The last column (`<lcol>`) must be less than or equal to 15 (default value of 11).

`-pct`

Shows the access count in each bucket as a percentage.



**NOTE:** For the following options, by default the histogram shows data from the start of the command.

`-prev|-begin`

Histogram displays data either from a previous sample (`-prev`) or from when the system was last started (`-begin`). If no option is specified, the histogram shows data from the beginning of the command's execution.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting `<dir>` as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

`-filt <fspec>`

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>,<count>`.

`<op>`

The `<op>` argument can be specified as one of the following:

- ◆ `r` - Specifies read statistics.
- ◆ `w` - Specifies write statistics.

- ◆ `t` - Specifies total statistics.
- ◆ `rw` - Specifies total read and write statistics.

`<val_ms>`

Specifies the threshold service time in milliseconds.

`<count>`

Specifies the minimum number of access above the threshold service time.

`-ni`

Specifies that histograms for only nonidle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays a histogram of service times for reads and writes to ports:

```
li% histport
12:35:24 10/20/04 -----Time (millisec)-----
Port      D/C 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
0:0:1     data   0   0   0   0   0   0   0   0   0   0
0:0:2     data 315 778   2   0   0   0   0   0   0
0:1:1     data   0   0   0   5  24  51  25   4   0   0
0:1:2     data   0   0   0   5  27  53  23   1   0   0
1:5:1     data   0   0   0   2  19  38  28  11   0   0
1:5:2     data   0   0   0   5  20  36  29   7   0   0
-----
total     data 315 778   2  17  90 178 105  23   0   0
Press the enter key to stop...
```

## NOTES

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column).

- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- Previous options `-n <node_list>`, `-s <slot_list>`, and `-p <port_list>` have been replaced with `-nodes <node_list>`, `-slots <slot_list>`, and `-ports <port_list>`.

---

**COMMAND**

histvlun

**DESCRIPTION**

The `histvlun` command displays Virtual Volume Logical Unit Number (VLUN) service time histograms.

**SYNTAX**

`histvlun [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-domain <domain_name>... | <pattern>...`

Shows only VLUNs whose Virtual Volumes (VVs) are in domains with names that match one or more of the specified domain names or patterns. Multiple domain names or patterns can be repeated using a comma-separated list.

`-host <hostname>... | <pattern>...`

Shows only VLUNs exported to the specified host(s) or pattern(s). Multiple host names or patterns can be repeated using a comma-separated list.

`-v <VV_name>... | <pattern>...`

Requests that only LDs mapped to VVs that match and of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list.

`-l <LUN>... | <pattern>...`

Specifies that VLUNs with LUNs matching the specified LUN(s) or pattern(s) are displayed. Multiple LUNs or patterns can be repeated using a comma-separated list.

`-nodes <node_list>`

Specifies that only exports from the specified nodes are to be displayed. The node list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that only exports from the specified slots are to be displayed. The slot list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that only exports to the specified ports are to be displayed. The port list is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the port list is not specified, all disks on all ports are displayed.

`-lw`

Lists the host's World Wide Name (WWN) or iSCSI name. This is especially useful when multiple WWNs or iSCSI names belonging to the same host are visible on the same port.

`-domainsum`

Specifies that sums for VLUNs are grouped by domain in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a – host name.

`-vvsum`

Specifies that sums for VLUNs of the same VV are displayed.

`-hostsum`

Specifies that sums for VLUNs are grouped by host in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a nameless host.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- ◆ both - (Default) Displays both I/O time and I/O size histograms.
- ◆ time - Displays only the I/O time histogram.
- ◆ size - Displays only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, show the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

- ◆ The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 6).
- ◆ The last column (<lcol>) must be less than or equal to 31 (default value of 15).



- ◆ The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, show the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

- ◆ The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).
- ◆ The last column (<lcol>) must be less than or equal to 15 (default value of 11).

`-pct`

Shows the access count in each bucket as a percentage.



**NOTE:** For the following options, by default the histogram shows data from the start of the command.

`-prev|-begin`

Histogram displays data either from a previous sample (`-prev`) or from when the system was last started (`-begin`). If no option is specified, the histogram shows data from the beginning of the command's execution.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on column number <col>. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, you can specify the direction of sorting <dir> as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

`-filt <fspec>`

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>,<count>` with no space after each comma.

`<op>`

The argument can be specified as one of the following:

- ◆ `r` - Specifies read statistics.
- ◆ `w` - Specifies write statistics.
- ◆ `t` - Specifies total statistics.
- ◆ `rw` - Specifies total read and write statistics.

`<val_ms>`

Specifies the threshold service time in milliseconds.

`<count>`

Specifies the minimum number of access above the threshold service time.

`-ni`

Specifies that histograms for only nonidle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays two iterations of a histogram of service times for all VLUNs:

cli% histvlun -iter 2													
12:48:50 10/20/04 -----Time (millisec)-----													
Lun	VVname	Host	Port	0.26	0.53	1.05	2.1	4.2	8.4	17	34	67	135
0	tpvv	queasy09	0:0:2	0	1	20	0	0	0	0	0	0	0
-----													
total				0	1	20	0	0	0	0	0	0	0
12:48:52 10/20/04 -----Time (millisec)-----													
Lun	VVname	Host	Port	0.26	0.53	1.05	2.1	4.2	8.4	17	34	67	135
0	tpvv	queasy09	0:0:2	0	1	30	0	0	0	0	0	0	0
-----													
total				0	1	30	0	0	0	0	0	0	0

## NOTES

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column).
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- Previous options `-n <node_list>`, `-s <slot_list>`, and `-p <port_list>` have been replaced with `-nodes <node_list>`, `-slots <slot_list>`, and `-ports <port_list>`.

---

**COMMAND**

histvv

**DESCRIPTION**

The `histvv` command displays Virtual Volume (VV) service time histograms in a timed loop.

**SYNTAX**

`histvv [options <arg>] [<VV_name>|<pattern>]...`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-domain <domainname>...|<pattern>...`

Shows only the VVs that are in domains with names that match the specified domain name(s) or pattern(s).

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- ◆ `both` - (Default) Displays both I/O time and I/O size histograms.
- ◆ `time` - Displays only the I/O time histogram.
- ◆ `size` - Displays only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, show the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

- ◆ The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`) (default value of 6).
- ◆ The last column (`<lcol>`) must be less than or equal to 31 (default value of 15).
- ◆ The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, show the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

- ◆ The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).
- ◆ The last column (<lcol>) must be less than or equal to 15 (default value of 11).

`-pct`

Shows the access count in each bucket as a percentage.



**NOTE:** For the following options, by default the histogram shows data from the start of the command.

`-prev`

Histogram displays data from a previous sample.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on column number <col>. Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting <dir> as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

`-filt <fspec>`

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>,<count>`.

`<op>`

The argument can be specified as one of the following:

- ◆ `r` - Specifies read statistics.
- ◆ `w` - Specifies write statistics.
- ◆ `t` - Specifies total statistics.
- ◆ `rw` - Specifies total read and write statistics.

`<val_ms>`

Specifies the threshold service time in milliseconds.

`<count>`

Specifies the minimum number of access above the threshold service time.

`-ni`

Specifies that histograms for only nonidle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays two iterations of a histogram of service times for all VVs:

```
%cli histvv -iter 2
12:53:03 10/20/04 -----Time (millisec)-----
      VVname 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
      admin   0   0   0   0   0   0   0   0   0   0
      tpvv    29   0   0   0   0   0   0   0   0   0
-----
      total   29   0   0   0   0   0   0   0   0   0

12:53:05 10/20/04 -----Time (millisec)-----
      VVname 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
      admin   0   0   0   0   0   0   0   0   0   0
      tpvv   105   0   0   0   0   0   0   0   0   0
-----
      total  105   0   0   0   0   0   0   0   0   0
```

## NOTES

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column).
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- If a <VV\_name> or <pattern> are specified, then VVs with names that match any of the patterns are listed, otherwise all VVs are listed. These patterns are glob-style patterns (see help on sub,globpat).
- Virtual volumes may be accessed externally by hosts and internally by the prefetcher. Virtual volume data measured by this command include accesses by the prefetcher.
- In addition to external accesses by hosts, VVs can be read internally by the system read-ahead prefetcher. The histvv data includes read-ahead accesses from the prefetcher that can cause the read data to appear more than seen by the hosts. Use the histv lun -vvsum command to see data for only accesses from the host.





# 16

## Locate Commands

---

In this chapter

locateage	<b>16.2</b>
locatesys	<b>16.4</b>

---

**COMMAND**

`locatecage`

**DESCRIPTION**

The `locatecage` command allows system administrators to locate a drive cage, drive magazine, or port in the system using the devices' blinking LEDs.

**SYNTAX**

The syntax for the `locatecage` command can be one of the following:

- To locate an entire drive cage:

```
locatecage [option <arg>] <cage_name>
```

- To locate a drive magazine:

```
locatecage [option <arg>] <cage_name> <mag>
```

- To locate a port:

```
locatecage [option <arg>] <cage_name> <port_name>
```

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-t <sec>`

Specifies the number of seconds, from 0 through 255 seconds, to blink the LED. If the argument is not specified, the option defaults to 60 seconds.

**SPECIFIERS**

`<cage_name>`

Specifies the drive cage name as shown in the `Name` column of `showcage` command output.

<port\_name>

Indicates the port specifiers. Accepted values are A0 | B0 | A1 | B1 | A2 | B2 | A3 | B3. If a port is specified, the port LED will oscillate between green and off. The <port\_name> specifier is not supported for DC3 drive cages.

<mag>

Indicates the drive magazine by number.

- ◆ For DC1 drive cages, accepted values are 0 through 4.
- ◆ For DC2 and DC4 drive cages, accepted values are 0 through 9.
- ◆ For DC3 drive cages, accepted values are 0 through 15.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The <port\_name> specifier is not supported for DC3 drive cages.

## EXAMPLES

The following example causes the Fibre Channel LEDs on the drive cage `cage0` to blink for 20 seconds:

```
cli% locatecage -t 20 cage0
```

## NOTES

- Issue the `showcage` command for a list of cage names.
- If no port or magazine is specified, all LEDs in the cage are set as amber or oscillate (depending on the cage type).
- If the port is specified, it turns green (not available in all cage types).

---

## COMMAND

locatesys

## DESCRIPTION

The `locatesys` command helps locate a storage system by blinking the node status LEDs on all nodes of a storage system alternating between amber and green. By default, the LEDs in all connected cages are also set as amber or oscillating (depending on the cage).

## SYNTAX

```
locatesys [options <arg>]
```

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-t <seconds>`

Specifies the number of seconds to blink the LEDs. The default is 60 seconds; the maximum is 255 seconds.

`-nodes <node_list>`

Specifies a comma-separated list of nodes on which to blink LEDs. The default is all nodes.

`-nocage`

Specifies that LEDs on the drive cages should not blink. The default is to blink LEDs for all cages in the system.

## SPECIFIERS

None.

## RESTRICTIONS

Access to all domains is required to run this command.

**EXAMPLES**

In the following example, an InServ Storage Server is identified by blinking the LEDs on all drive cages in the system for 90 seconds.

```
cli% locatesys -t 90
```

**NOTES**

None.



# 17

## Move Commands

---

In this chapter

<code>movech</code>	<b>17.2</b>
<code>movechtospare</code>	<b>17.6</b>
<code>movepdtospare</code>	<b>17.9</b>
<code>moverelocpd</code>	<b>17.12</b>
<code>movetodomain</code>	<b>17.18</b>

---

**COMMAND**

`movech`

**DESCRIPTION**

The `movech` command moves a list of chunklets from one physical disk to another.

**SYNTAX**

`movech [options] <fd:fp-td:tp>...`

**AUTHORITY**

Super, Service, Edit



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-nowait`

Specifies that the command returns before the operation is completed.

`-dr`

Specifies that the operation is a dry run. No chunklets are actually moved.

`-devtype`

Permits the moves to happen to different device types.

`-perm`

Specifies that chunklets are permanently moved and the chunklets' original locations are not remembered. If the `-perm` option is not specified, the chunklets' original locations are retained, thereby allowing the chunklets to be returned to their original locations through the `moverelocpd` and `servicemag resume` commands.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.



`-ovrd`

Permits the moves to happen to a destination even when there will be a loss of quality because of the move. This option is only necessary when the target of the move is not specified and the `-perm` flag is used.

## SPECIFIERS

`<fd:fp>[<-td:tp>]...`

Specifies that the chunklet located at the specified PD (`<fd>`) and the chunklet's position on that disk (`<fp>`) be moved to either the specified destination disk (`<td>`) and chunklet position (`<tp>`), or a location determined by the system if a destination (`<-td:tp>`) is not specified. This specifier must be used at least once on the command line. Repeated use of this specifier allows multiple chunklets to be moved.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example moves the chunklet in position 0 on disk 24, to position 50 on disk 64 and chunklet in position 0 on disk 25, to position 1 on disk 27:

```
cli% movech 24:0-64:50 25:0-27:1
Are you sure you want to move the chunklets?
select q=quit y=yes n=no: y

Move      -State- --Detailed_State---
24:0-64:50 normal    normal
25:0-27:1  degraded disks_reused_in_row
```

The columns in the previous example are identified as follows:

- Status. The overall status of the move.
  - ◆ normal. Indicates the chunklet has successfully moved with no loss of quality.
  - ◆ degraded. Indicates the chunklet has successfully moved with loss of quality.
  - ◆ failed. Indicates the chunklet was not moved.
- Detailed\_State. Provides additional details of the move with General Information, Error Information, or Degraded Information.

**General Information:**

- ◆ `dest_unknown`. Destination is unknown.
- ◆ `dest_unknown`. Destination is known.
- ◆ `valid_move`. Source and Destination are valid.
- ◆ `ch_moved`. Chunklet was moved.
- ◆ `ch_move_pending`. Chunklet move is pending due to `-nowait` option.
- ◆ `move_error`. Error moving chunklet.

**Error Information:**

- ◆ `error_type_unknown`. An unknown error occurred.
- ◆ `spares_not_found`. No suitable spares available.
- ◆ `duplicate_src`. Duplicate source input.
- ◆ `duplicate_dest`. Duplicate destination in chunklet list.
- ◆ `src_relocating`. Source is relocating.
- ◆ `ch_synching`. Chunklet is synchronizing.
- ◆ `src_no_ld`. Source not associated with an LD.
- ◆ `src_ld_invalid`. Source associated with invalid LD.
- ◆ `src_set_invalid`. Source is in invalid set.
- ◆ `src_not_found`. Source chunklet does not exist.
- ◆ `invalid_dest`. Destination is invalid.
- ◆ `move_failed`. Chunklet move failed.
- ◆ `disk_relocating`. Disk is relocating.

**Degraded Information:**

- ◆ `disks_reused_in_row`. Disks reused in row.
- ◆ `connection_degraded`. Connection is degraded.
- ◆ `remote_pds_used`. Remote PDs used.
- ◆ `ldpattern_not_obeyed`. LD pattern was not obeyed.

- ◆ `raid_availability_reduced`. RAID availability reduced.

## NOTES

- Chunklets moved through the `movech` command are only moved temporarily. Issuing either the `moverelocpd` or `servicemag resume` command can move the chunklet back to its original position.
- Specifying the `-dr` option can be used to see if the specified moves succeed and what the results (quality) of the moves are.

---

**COMMAND**

`movechtospare`

**DESCRIPTION**

The `movechtospare` command moves data from specified Physical Disks (PDs) to a temporary location selected by the system.

**SYNTAX**

`movechtospare [options] <fd:fp>`

**AUTHORITY**

Super, Service, Edit



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-nowait`

Specifies that the command returns before the operation is completed.

`-dr`

Specifies that the operation is a dry run. No chunklets are actually moved.

`-devtype`

Permits the moves to happen to different device types.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<fd:fp>`

Indicates that the move takes place from the specified PD (`<fd>`) and chunklet position (`<fp>`).

**RESTRICTIONS**

Access to all domains is required to run this command.

## EXAMPLES

In the following example, chunklet 0 from physical disk 66 is moved to spare:

```
cli% movechtospare 66:0
      Are you sure you want to move the chunklet to spare?
      select q=quit y=yes n=no: y

Move      -Status- -Detailed_State--
66:0-54:49 normal    ch_moved,normal
```

The columns in the previous example are identified as follows:

- **Status.** The overall status of the move.
  - ◆ **normal.** Indicates the chunklet has successfully moved with no loss of quality.
  - ◆ **degraded.** Indicates the chunklet has successfully moved with loss of quality.
  - ◆ **failed.** Indicates the chunklet was not moved.
- **Detailed\_State.** Provides additional details of the move with General Information, Error Information, or Degraded Information.

### General Information:

- ◆ **dest\_unknown.** Destination is unknown.
- ◆ **dest\_known.** Destination is known.
- ◆ **valid\_move.** Source and Destination are valid.
- ◆ **ch\_moved.** Chunklet was moved.
- ◆ **ch\_move\_pending.** Chunklet move is pending due to `-nowait` option.
- ◆ **move\_error.** Error moving chunklet.

### Error Information:

- ◆ **error\_type\_unknown.** An unknown error occurred.
- ◆ **spares\_not\_found.** No suitable spares available.
- ◆ **duplicate\_src.** Duplicate source input.
- ◆ **duplicate\_dest.** Duplicate destination in chunklet list.
- ◆ **src\_relocating.** Source is relocating.

- ◆ `ch_synching`. Chunklet is synchronizing.
- ◆ `src_no_ld`. Source not associated with an LD.
- ◆ `src_ld_invalid`. Source associated with invalid LD.
- ◆ `src_set_invalid`. Source is in invalid set.
- ◆ `src_not_found`. Source chunklet does not exist.
- ◆ `invalid_dest`. Destination is invalid.
- ◆ `move_failed`. Chunklet move failed.
- ◆ `disk_relocating`. Disk is relocating.

**Degraded Information:**

- ◆ `disks_reused_in_row`. Disks reused in row.
- ◆ `connection_degraded`. Disks reused in row.
- ◆ `remote_pds_used`. Connection Degraded.
- ◆ `ldpattern_not_obeyed`. LD pattern was not obeyed.
- ◆ `raid_availability_reduced`. RAID availability reduced.

**NOTES**

None.

---

**COMMAND**

`movepdtospare`

**DESCRIPTION**

The `movepdtospare` command moves data from specified Physical Disks (PDs) to a temporary location selected by the system.

**SYNTAX**

`movepdtospare [options] <PD_ID>...`

**AUTHORITY**

Super, Service, Edit



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-nowait`

Specifies that the command returns before the operation is completed.

`-dr`

Specifies that the operation is a dry run. No PDs are actually moved.

`-devtype`

Permits the moves to happen to different device types.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<PD_ID>...`

Specifies the PD ID. This specifier can be repeated to move multiple PDs.

**RESTRICTIONS**

Access to all domains is required to run this command.

## EXAMPLES

The following example displays a dry run of moving the data on PD 0 to free or spare space:

```
cli% movepdto spare -dr 0
Are you sure you want to move the pd to spare?
select q=quit y=yes n=no: y
Move      -Status- -Detailed_State--
0:0-92:49 normal    valid_move,normal
0:1-62:50 normal    valid_move,normal
0:2-42:51 normal    valid_move,normal
0:3-88:52 normal    valid_move,normal
0:4-62:53 normal    valid_move,normal
0:5-90:55 normal    valid_move,normal
0:6-88:55 normal    valid_move,normal
```

The columns in the previous example are identified as follows:

- **Status.** The overall status of the move.
  - ◆ normal. Indicates the chunklet has successfully moved with no loss of quality.
  - ◆ degraded. Indicates the chunklet has successfully moved with loss of quality.
  - ◆ failed. Indicates the chunklet was not moved.
- **Detailed\_State.** Provides additional details of the move with General Information, Error Information, or Degraded Information.

### General Information:

- ◆ dest\_unknown. Destination is unknown.
- ◆ dest\_known. Destination is known.
- ◆ valid\_move. Source and Destination are valid.
- ◆ ch\_moved. Chunklet was moved.
- ◆ ch\_move\_pending. Chunklet move is pending due to `-nowait` option.
- ◆ move\_error. Error moving chunklet.

### Error Information:

- ◆ error\_type\_unknown. An unknown error occurred.
- ◆ spares\_not\_found. No suitable spares available.



- ◆ `duplicate_src`. Duplicate source input.
- ◆ `duplicate_dest`. Duplicate destination in chunklet list.
- ◆ `src_relocating`. Source is relocating.
- ◆ `ch_synching`. Chunklet is synchronizing.
- ◆ `src_no_ld`. Source not associated with an LD.
- ◆ `src_ld_invalid`. Source associated with invalid LD.
- ◆ `src_set_invalid`. Source is in invalid set.
- ◆ `src_not_found`. Source chunklet does not exist.
- ◆ `invalid_dest`. Destination is invalid.
- ◆ `move_failed`. Chunklet move failed.
- ◆ `disk_relocating`. Disk is relocating.

**Degraded Information:**

- ◆ `disks_reused_in_row`. Disks reused in row.
- ◆ `connection_degraded`. Disks reused in row.
- ◆ `remote_pds_used`. Connection Degraded.
- ◆ `ldpattern_not_obeyed`. LD pattern was not obeyed.
- ◆ `raid_availability_reduced`. RAID availability reduced.

## NOTES

- The destination PDs do not need to be specified as the system automatically determines the spare locations.
- Specifying the `-dr` option can be used to see if the specified moves succeeds and the results (quality) of the moves.

---

**COMMAND**

moverelocpd

**DESCRIPTION**

The moverelocpd command moves chunklets that were on a physical disk to the target of relocation.

**SYNTAX**

moverelocpd [options] [<fd>[-<td>...]

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

-nowait

Specifies that the command returns before the operation is completed.

-dr

Specifies that the operation is a dry run. No physical disks are actually moved.

-partial

Move as many chunklets as possible. If this option is not specified, the command fails if not all specified chunklets can be moved.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-p <pattern>

Specifies a pattern to select <fd> disks. The <td> specifier cannot be used with this -p option. If no <fd> are specified, then all disks that match the pattern are selected as the <fd> disks. If <fd> disks are specified along with -p then only disks that match the pattern are selected as <fd> disk. The following arguments can be specified as patterns for this option:

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (e.g. 1, 2, 3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (e.g. 1, 2, 3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (e.g. 1, 2, 3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (e.g. 1, 2, 3). A range of drive cages is separated with a hyphen (e.g. 0-3). Disks must reside in the specified drive cage(s).

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the `CagePos` column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (e.g. 1, 2, 3). A range of drive magazines is separated with a hyphen (e.g. 0-7). Disks must reside in the specified drive magazine(s).

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (e.g. 1, 2, 3). A range of disk positions is separated with a hyphen (e.g. 0-3). Disks must reside in the specified position(s).

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (e.g. 1, 2, 3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

`<fd>[-<td>]...`

Specifies that the chunklets that were relocated from specified disk (`<fd>`), are moved to the specified destination disk (`<td>`). If destination disk (`<td>`) is not specified then the chunklets are moved back to original disk (`<fd>`). The `<fd>` specifier is not needed if `-p` option is used, otherwise it must be used at least once on the command line. If this specifier is repeated then the operation is performed on multiple disks.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example moves chunklets that were on PD 8 that were relocated to another position, back to PD 8:

```
cli% moverelocpd 8
Are you sure you want to move the chunklets ?
select q=quit y=yes n=no: y
Move      -Status- -Detailed_State-
42:51-8:2  normal  normal
62:50-8:1  normal  normal
62:53-8:4  normal  normal
62:57-8:7  normal  normal
64:49-8:0  normal  normal
88:52-8:3  normal  normal
88:55-8:6  normal  normal
90:55-8:5  normal  normal
```

The columns in the previous example are identified as follows:

- Status. The overall status of the move.
  - ◆ normal. Indicates the chunklet has successfully moved with no loss of quality.
  - ◆ degraded. Indicates the chunklet has successfully moved with loss of quality.
  - ◆ failed. Indicates the chunklet was not moved.

- **Detailed\_State.** Provides additional details of the move with General Information, Error Information, or Degraded Information.

**General Information:**

- ◆ **dest\_unknown.** Destination is unknown.
- ◆ **dest\_known.** Destination is known.
- ◆ **valid\_move.** Source and Destination are valid.
- ◆ **ch\_moved.** Chunklet was moved.
- ◆ **ch\_move\_pending.** Chunklet move is pending due to `-nowait` option.
- ◆ **move\_error.** Error moving chunklet.

**Error Information:**

- ◆ **error\_type\_unknown.** An unknown error occurred.
- ◆ **spares\_not\_found.** No suitable spares available.
- ◆ **duplicate\_src.** Duplicate source input.
- ◆ **duplicate\_dest.** Duplicate destination in chunklet list.
- ◆ **src\_relocating.** Source is relocating.
- ◆ **ch\_synching.** Chunklet is synchronizing.
- ◆ **src\_no\_ld.** Source not associated with an LD.
- ◆ **src\_ld\_invalid.** Source associated with invalid LD.
- ◆ **src\_set\_invalid.** Source is in invalid set.
- ◆ **src\_not\_found.** Source chunklet does not exist.
- ◆ **invalid\_dest.** Destination is invalid.
- ◆ **move\_failed.** Chunklet move failed.
- ◆ **disk\_relocating.** Disk is relocating.

**Degraded Information:**

- ◆ **disks\_reused\_in\_row.** Disks reused in row.
- ◆ **connection\_degraded.** Disks reused in row.

- ◆ `remote_pds_used`. Connection Degraded.
- ◆ `ldpattern_not_obeyed`. LD pattern was not obeyed.
- ◆ `raid_availability_reduced`. RAID availability reduced.

## NOTES

- Chunklets moved from physical disks `fd` are treated as if they originated on disk `td`. Disk `td` can be the same as disk `fd`.
- Specifying the `-dr` option can be used to see if the specified moves succeeds and what the results (quality) of the moves are.
- If the `-partial` option is used, the command relocates as many chunklets as possible and prints messages for the chunklets it could not move.

---

## COMMAND

`movetodomain`

## DESCRIPTION

The `movetodomain` command moves Common Provisioning Groups (CPGs) from one domain to another.

## SYNTAX

`movetodomain [options] <object_name> <domain_name | domain_set>`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-vv`

Specifies that the object is a virtual volume.

`-cpg`

Specifies that the object is a CPG.

`-host`

Specifies that the object is a host.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

`<object_name>`

Specifies the name of the object to be moved.

`<domain_name | domain_set>`

Specifies the domain or domain set to which the specified object is moved. The domain set name must start with `set:`. To remove the domain from an object from any domain, specify the string `-unset` for the `<domain_name>` or `<domain_set>` specifier.



## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLE

The following example displays the movement of VV vv1 to domain SampleDomain:

```
cli% movetodomain -vv vv1 SampleDomain
The following volumes will have their domain modified:
vv1
The following hosts will have their domain modified:
thehost
The following CPGs will have their domain modified:
SampleCPG
14 associated LDs will also have their domain changed.
Do you want to proceed with moving the above to domain SampleDomain?
select y=yes n=no: y
```

The following example displays the removal of the host testhost from any domain:

```
cli% movetodomain -host testhost -unset

The following hosts will have their domain modified:

Id Name      Persona -WWN/iSCSI_Name- Port
2 testhost   Generic

Do you want to proceed with removing the domain of the above?
select y=yes n=no: y
```

## NOTES

- The `movetodomain` command moves all objects that are directly or indirectly related to the specified object into the specified domain. Possible relationships include, but are not limited to, Virtual Volume Logical Unit Numbers (VLUNs) between hosts and VVs, VVs using a CPG for snapshot space, and two VVs sharing the same logical disk. If the `-f` option is not used, a list of all objects that will be modified is shown before the confirmation prompt.
- Only hosts are permitted to be members of domain sets.



# 18

## Promote Commands

---

In this chapter

<code>promotesv</code>	<b>18.2</b>
<code>promotevvcopy</code>	<b>18.4</b>

---

## COMMAND

`promotesv`

## DESCRIPTION

The `promotesv` command copies the differences of a snapshot back to its base volume, allowing you to revert the base volume to an earlier point in time.

## SYNTAX

`promotesv [options] <virtual_copy_name | VV_set>`

## AUTHORITY

Super, Edit

## OPTIONS

`-target`

Copy the differences of the virtual copy to the specified RW parent in the same VV family tree. The default is to copy the differences to the base volume. This option cannot be used with the `-halt` option.

`-rcp`

Allows the promote operation to proceed even if the RW parent volume is currently in a Remote Copy volume group, if that group has not been started. If the remote copy group has been started, this command fails. This option cannot be used in conjunction with the `-halt` option.

`-halt`

Cancels an ongoing snapshot promotion. Marks the RW parent volume with the `cpf` status, which can be cleaned up using the `promotevvcopy` (see [promotevvcopy](#) on page 18.4) command or by issuing a new instance of the `promotesv` command. This option cannot be used in conjunction with any other option.

## SPECIFIERS

`<virtual_copy_name | VV_set>`

Specifies the name of the virtual copy volume or set of virtual copy volumes to be promoted, using up to 31 characters in length. The virtual volume set name must start with `set:.`

## RESTRICTIONS

- The virtual copy and its base volume must not be exported.
- Only one promote operation is allowed at a time within a VV family tree.

## EXAMPLES

The following example promotes the differences from VV Dub and its base volume Dublin:

```
cli% promotesv Dub
```

## NOTES

- Issue the `showvv` command to verify that differences in the snapshot volume are promoted to its target volume.
- Issue the `showvv -d` command to display the number of remaining blocks to be copied.

---

**COMMAND**

`promotevvcopy`

**DESCRIPTION**

The `promotevvcopy` command promotes a physical copy back to a regular base volume.

**SYNTAX**

`promotevvcopy <physical_copy_name>`

**AUTHORITY**

Super, Edit

**OPTIONS**

None.

**SPECIFIERS**

`<physical_copy_name>`

Specifies the name of the physical copy to be promoted, using up to 31 characters.

**RESTRICTIONS**

The physical copy must have completed the copy from the base volume.

**EXAMPLES**

The following example promotes virtual volume Belfast to a base volume:

```
cli% promotevvcopy Belfast
```

**NOTES**

- The saved snapshot of the parent of `<physical_copy_name>` is also removed.
- The `promotevvcopy` command can also be used to clean up a failed physical copy.
- Issue the `showvv` command to verify that promoted volume is a base volume.
- After a physical copy has been promoted, the association between it and its parent volume is broken; the physical copy and base volume can no longer resync. The saved snapshot of the parent of physical copy is also removed.

# 19

## Remove Commands

---

### In this chapter

<code>removealert</code>	<b>19.3</b>
<code>removecpg</code>	<b>19.5</b>
<code>removedomain</code>	<b>19.7</b>
<code>removedomainset</code>	<b>19.8</b>
<code>removeeventlog</code>	<b>19.10</b>
<code>removehost</code>	<b>19.12</b>
<code>removehostset</code>	<b>19.14</b>
<code>removeld</code>	<b>19.16</b>
<code>removercopygroup</code>	<b>19.18</b>
<code>removercopytarget</code>	<b>19.20</b>
<code>removesched</code>	<b>19.22</b>
<code>removesnmpmgr</code>	<b>19.24</b>
<code>removesnmppw</code>	<b>19.26</b>
<code>removespare</code>	<b>19.28</b>
<code>removesshkey</code>	<b>19.30</b>

removetask	<b>19.31</b>
removetemplate	<b>19.33</b>
removeuser	<b>19.35</b>
removeuserconn	<b>19.37</b>
removevlun	<b>19.39</b>
removevv	<b>19.43</b>
removevvset	<b>19.46</b>



---

## COMMAND

`removealert`

## DESCRIPTION

The `removealert` command removes one or more alerts from the system.



**CAUTION:** Use care when removing alerts. Alerts that have not been fixed or acknowledged should NOT be removed.

## SYNTAX

`removealert [option] -a | <alert_ID>...`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-a`

Specifies all alerts from the system and prompts removal for each alert. If this option is not used, then the `<alert_ID>` specifier must be used.

`-f`

Specifies that the command is forced. If this option is not used and there are alerts in the new state, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

`<alert_ID>...`

Indicates a specific alert be removed from the system. This specifier can be repeated to remove multiple alerts. If this specifier is not used, the `-a` option must be used.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays all alerts from the system with the option to remove individual alerts:

```
cli% removealert -a

Id 120 - New
  Occurred 4 times, last at Tue May 03 22:45:47 PDT 2005
  Message code: 196609
  Tue May 03 22:23:17 PDT 2005
  Node: 0 Severity: Minor
  Firmware coredump event
  Firmware COREDUMP: recovered file /var/core/hba/fwcore.n00.s02.p01.20050503.224547

Alert 120 is marked as "New".
Are you sure you want to remove it?
select q=quit y=yes n=no: y

Id 131 - New
  Message code: 1114115
  Thu May 05 00:11:25 PDT 2005
  Node: 0 Severity: Minor
  Too many events are being logged
  Too many events are being generated. 2 event files were rolled over in less
  than 1800 seconds. Current event files could not be archived because too many
  have been archived already.

Alert 131 is marked as "New".
Are you sure you want to remove it?
select q=quit y=yes n=no: y

Id 133 - New
  Message code: 1966081
  Thu May 05 00:25:27 PDT 2005
  Node: 0 Severity: Degraded
  Cage log event
  cage2-A, loop 1:0:2, cage time Thu May 5 00:25:29 2005. Fan at position 1 is
  running at high speed. Internal parameters: 0x0003 0x0109 01 01 00 00 00 00 00
  00 00 00 00 00.

Alert 133 is marked as "New".
Are you sure you want to remove it?
select q=quit y=yes n=no:
```

## NOTES

- Use care when removing alerts. Alerts that have not been fixed or acknowledged should NOT be removed.

---

## COMMAND

`removecpg`

## DESCRIPTION

The `removecpg` command removes Common Provisioning Groups (CPGs) from the system or removes specific Logical Disks (LDs) from CPGs.

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## SYNTAX

```
removecpg [options <arg>] {<CPG_name>|<pattern>}...
```

## OPTIONS

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-ld`

Specifies that after the CPG is removed, all LDs that were part of the CPG are also removed. This is the default behavior.

`-sa <LD_name>`

Specifies that the LD, as identified with the `<LD_name>` argument, used for snapshot administration space allocation is removed. The `<LD_name>` argument can be repeated to specify multiple LDs. This option is deprecated and will be removed in a subsequent release.

`-sd <LD_name>`

Specifies that the LD, as identified with the `<LD_name>` argument, used for snapshot data space allocation is removed. The `<LD_name>` argument can be repeated to specify multiple LDs. This option is deprecated and will be removed in a subsequent release.

`-pat`

The specified patterns are treated as glob-style patterns and that all common provisioning groups matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the `-f` option is specified. This option must be used if the `pattern` specifier is used.

## SPECIFIERS

`<CPG_name>`

Specifies the name of the CPG that is either being removed or losing LDs.

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to remove multiple common provisioning groups. If this specifier is not used, the `<CPG_name>` specifier must be used. See [Glob-Style Pattern](#) on page 2.4 for more information.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The `removecpg` command fails if any of the LDs, or the entire CPG, is in use by a Thinly Provisioned Virtual Volume (TPVV).

## EXAMPLES

The following example displays the removal of CPG `cpg1`:

```
cli% removecpg cpg1
```

## NOTES

- By default, this command deletes any unused LDs. This is equivalent to using the `-ld` option. The `-ld` option still exists for backward compatibility.
- If neither the `-sa` or `-sd` options are specified, the entire CPG is removed, including all LDs.
- The operation fails if any of the LDs are in use.

---

**COMMAND**

removedomain

**DESCRIPTION**

The `removedomain` command removes an existing domain from the system.

**SYNTAX**

`removedomain [option] <domain_name>`

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-f`

When using this option, the command does not ask for confirmation before removing the domain.

**SPECIFIERS**

`<domain_name>`

Specifies the domain that is removed.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example removes the domain named `sample_domain` from the system:

```
cli% removedomain -f sample_domain
```

**NOTES**

None.

---

**COMMAND**

`removedomainset`

**DESCRIPTION**

The `removedomainset` command removes a domain set or removes domains from an existing set.

**SYNTAX**

`removedomainset [options] <setname> [<domain>...]`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<setname>`

Specifies the name of the domain set.

`<domain>...`

Optional list of domain names that are members of the set. If no `<domain>` is specified, the domain set is removed, otherwise the specified `<domain>` is removed from the domain set.

**RESTRICTIONS**

None.

**EXAMPLES**

To remove a domain set:

```
cli% removedomainset domainset
```

To remove a single domain from a set:

```
cli% removedomainset domainset domain1
```

## NOTES

None.

---

**COMMAND**`removeeventlog`**DESCRIPTION**

The `removeeventlog` command removes all Debug events from the system event log.

**SYNTAX**`removeeventlog [option]`**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

None.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example removes Debug events from the system event log:

```
cli% removeeventlog -f
```



**NOTES**

- Without the `-f` flag, the command prompts for confirmation.
- Verify the removal of event logs by issuing the `showeventlog` command. See [showeventlog](#) on page 22.41 for more information.

---

## COMMAND

removehost

## DESCRIPTION

The `removehost` command removes a system host or paths to a host. If one or more paths are specified, the command removes only those paths, otherwise the entire host definition is removed.

## SYNTAX

```
removehost [options] <hostname> [<WWN>...|<ISCSI_name>...]
```

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-rvl`

Remove World Wide Name(s) (WWN(s)) or iSCSI names even if there are Virtual Volume Logical Unit Numbers (VLUNs) exported to the host. This option cannot be used if the entire host definition is being removed.

`-iscsi`

Specifies that the paths are iSCSI names. If this option is not specified, the paths are WWNs.

## SPECIFIERS

`<hostname>`

Specifies the host name, using up to 31 characters.

`<WWN>`

Only the specified WWN(s) path to the specified host is removed. This specifier is not required on the command line. If a WWN is not specified, the entire host definition is removed.

<ISCSI\_name>

Specifies the host iSCSI name to be removed from the specified host. If no iSCSI name is specified, the entire host definition is removed.

## RESTRICTIONS

- Access to all domains is required to run this command.
- A host that has one or more VLUNs exported on it cannot be removed.

## EXAMPLES

The following example removes host `test01`:

```
cli% removehost test01
```

## NOTES

- Removing an entire host definition by issuing the `removehost <hostname>` command is not equivalent to removing all of the paths associated with a host. The latter leaves a host definition with no paths associated to it, whereas the former removes the entire host definition.
- Verify the removal of hosts by issuing the `showhost` command.

---

**COMMAND**

removehostset

**DESCRIPTION**

The `removehostset` command removes a host set or removes hosts from an existing set.

**SYNTAX**

`removehostset [options] <setname> [<host>...]`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<setname>`

Name of the host set to remove.

`<host>...`

Optional list of host names that are members of the set. If no `<host>` is specified, the host set is removed, otherwise the specified `<host>` is removed from the host set.

**RESTRICTIONS**

None.

**EXAMPLES**

To remove a host set:

```
cli% removehostset hostset
```

To remove a single host from a set:

```
cli% removehostset hostset host1
```

## NOTES

None.

---

**COMMAND**

removeld

**DESCRIPTION**

The `removeld` command removes a specified Logical Disk (LD) from the system service group.

**SYNTAX**

`removeld [options] {<LD_name>|<pattern>}...`

**AUTHORITY**

Super, Service, Edit

**OPTIONS**

`-pat`

Specified patterns are treated as glob-style patterns and all LDs matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the `-f` option is specified. This option must be used if the `pattern` specifier is used.

`-dr`

Specifies that the operation is a dry run and no LDs are removed.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-rmsys`

Specifies that system resource LDs such as logging LDs and preserved data LDs are removed. See the *InForm OS Concepts Guide* for information on logging logical disks and preserved data logical disks.



**CAUTION:** System resource logical disks are required for correct operation of the InServ storage system. Removal of system resource logical disks should be performed by qualified service personnel. Incorrect use of the `-rmsys` option can result in data loss.

## SPECIFIERS

<LD\_name>

Specifies the LD name, using up to 31 characters. Multiple LDs can be specified.

<pattern>

Specifies a glob-style pattern. Multiple LDs can be specified. If this specifier is not used, the <LD\_name> specifier must be used. See [Glob-Style Pattern](#) on page 2.4 for more information.

## RESTRICTIONS

- System resource LDs are required for correct operation of the InServ Storage Server. Removal of system resource LDs should be performed by qualified service personnel. Incorrect use of the `-rmsys` option can result in data loss.
- If the LD is mapped to a virtual volume, the LD is not removed and the command stops at that logical disk.
- Issuing the `removeld` command with the `-pat` option specified returns a request for confirmation to remove LDs, unless the `-f` option is specified.
- Do not issue the `removeld` command while a `movech`, `movech2spare`, `movepd2spare`, or `moverelocpd` operation is being executed. Issue the `showldch` command to view operations currently running on the system. See [showldch](#) on page 22.66 for additional information.

## EXAMPLES

The following example removes LD `test0`:

```
cli% removeld -f test0
```

## NOTES

Verify the removal of LDs by issuing the `showld` command. See [showld](#) on page 22.58 for additional information.

---

**COMMAND**

removercopygroup

**DESCRIPTION**

The `removercopygroup` command removes a Remote Copy volume group.

**SYNTAX**

`removercopygroup <group_name>`

**AUTHORITY**

Super, Edit

**OPTIONS**

None.

**SPECIFIERS**

`<group_name>`

The name of the group to be removed.

**RESTRICTIONS**

- This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- This command is not allowed if Remote Copy is in progress; the system generates an error. The group must be stopped using the `stoprcopygroup` command.

**EXAMPLES**

The following example removes `Group1` from the local system:

```
cli% removercopygroup Group1
```

**NOTES**

- If the `mirror_config` option policy is set for this group's target and the group is a primary group, then this command is mirrored to the target and the corresponding group is also removed. If the policy is set and the group is a secondary group, then this command fails.



- The `removercopygroup` command removes all the associations configured in the specified group and removes the group name and any Remote Copy synchronization snapshots affiliated with volumes in the group.

---

## COMMAND

`removercopytarget`

## DESCRIPTION

The `removercopytarget` command removes target designation from a Remote Copy system and removes all links affiliated with that target.

## SYNTAX

`removercopytarget [options] <target_name>`

## AUTHORITY

Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-cleargroups`

Remove all groups that have no other targets or dismiss this target from groups with additional targets. All groups that contain this target must be stopped before this can be issued.

`-f`

Do not ask for confirmation when `-cleargroups` is specified.

## SPECIFIERS

`<target_name>`

The target name for the target definition to be removed.

## RESTRICTIONS

- This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- A target definition cannot be deleted if it is being used by any group.

**EXAMPLES**

The following example removes target `InServ1_in` from a Remote Copy system:

```
cli% removercopytarget InServ1_in
```

**NOTES**

None.

---

**COMMAND**

removesched

**DESCRIPTION**

The `removesched` command removes a scheduled task from the system.

**SYNTAX**

`removesched [options] <schedname|pattern>`

**AUTHORITY**

Super, Service

**OPTIONS**

`-pat`

Specifies that certain patterns are treated as glob-style patterns and that all scheduled tasks matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the `-f` option is specified. This option must be used if the pattern specifier is used.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<schedname>`

Specifies the schedule that is removed. Can be at most 31 characters in length.

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to remove multiple scheduled tasks. If this specifier is not used, the `<schedname>` specifier must be used. The `<schedname>` can be up to 31 characters in length.

**RESTRICTIONS**

None.

## EXAMPLES

None.

## NOTES

None.

---

**COMMAND**

```
removesnmpmgr
```

**DESCRIPTION**

The `removesnmpmgr` command removes preregistered SNMP managers from receiving alerts (traps).

**SYNTAX**

```
removesnmpmgr [option <arg>] <manager_IP>
```

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

```
-p <port_number>
```

Specifies the port number where the manager receives traps. If not specified, the port number defaults to 162.

**SPECIFIERS**

```
<manager_IP>
```

Specifies the IPv4 or IPv6 address of the host where the manager runs. IPv6 address is in hexadecimal, case insensitive, and is separated by colons. An example would be:

```
5def:2008:0:0:abcd:0:0:161a
```

In addition, a double colon (: :) can be used once in an address to replace multiple fields of zeros. For example:

```
5def:2008:0:0:abcd::161a
```

**RESTRICTIONS**

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the removal of a manager with the IPv4 address 123.45.67.89 from the list of registered managers:

```
cli% removesnmpmgr 123.45.67.89
```

The following example displays the removal of a manager with the IPv6 address 5def:2008:abcd::161a from the list of registered managers:

```
cli% removesnmpmgr 5def:2008:abcd::161a
```

## EXIT STATUS

The following codes are returned indicating success or failure:

- 0 indicates that the trap manager host was removed and the command was successful.
- 1 indicates that the command failed.
- 2 indicates that the host is not on the list of registered hosts.

## NOTES

- SNMP managers are registered by issuing the `addsnmpmgr` command. See [addsnmpmgr](#) on page 4.2 for additional information.
- Verify the removal of SNMP managers by issuing the `showsnmpmgr` command. See [showsnmpmgr](#) on page 22.157 for more information.

---

**COMMAND**

`removesnmppw`

**DESCRIPTION**

The `removesnmppw` command allows a user to remove SNMP access community string passwords.

**SYNTAX**

`removesnmppw [options]`

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-rw` | `-r` | `-w`

Removes the read/write (`-rw`), read-only (`-r`), or write-only (`-w`) password. If not specified, the read/write community string password is removed.

`-f`

Forces the operation so that the command does not require confirmation before proceeding.

**SPECIFIERS**

None.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example displays the removal of the read/write SNMP access password:

```
cli% removesnmppw -f
```



## EXIT STATUS

The following codes are returned indicating success or failure:

- 0 indicates that the password was removed and the command was successful.
- 1 indicates that the command failed.
- 2 indicates that a password does not exist.

## NOTES

- After a password has been removed, the system manager can no longer use that password to send requests to the SNMP agent.
- Verify the removal of SNMP passwords by issuing the `showsnmppw` command. See [showsnmppw](#) on page 22.159 for additional information.

---

**COMMAND**

removespare

**DESCRIPTION**

The `removespare` command removes chunklets from the spare chunklet list.

**SYNTAX**

`removespare [options] <chunklet_specifier>...`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-f`

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-p`

Specifies that partial completion of the command is acceptable. Chunklets specified that are not on the current service group spare list are ignored.

**SPECIFIERS**

`<chunklet_specifier>...`

The chunklet specifier can be issued in the following formats:

`<PD_ID:chunklet_num>...`

Specifies the identification of the physical disk (`PD_ID`) and the position number of the chunklet on the disk (`chunklet_num`). This specifier can be repeated.

`<PD_ID>:a...`

Specifies the identification of the physical disk (`PD_ID`) and all (a) chunklets on the disk. This specifier can be repeated.

`a:<chunklet_num>...`

Specifies the chunklet number on all physical disks. This specifier can be repeated.

`a:a`

Removes all spare chunklets that are not currently used.

`-pos <cage:mag:disk:chunklet_num>`

Specifies the position of a specific chunklet identified by its position in a drive cage, drive magazine, physical disk, and chunklet number. For example `-pos 1:0.2:3:121`, where 1 is the drive cage, 0.2 is the drive magazine, 3 is the physical disk, and 121 is the chunklet number.

`-pos <cage:mag:disk:a>`

Specifies that all chunklets on a physical disk, identified by drive cage number, drive magazine number, and disk number, are marked to be removed.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example removes a spare chunklet from position 3 on physical disk 1:

```
cli% removespare 1:3
```

## NOTES

- Verify the removal of spare chunklets by issuing the `showspare` command. See [showspare \[-used\]](#) on page 22.166 for more information.
- If a wildcard (a) is used or the `-p` flag is specified, prints the number of spares removed. Otherwise, if all the explicitly specified spares could not be removed, prints an error message.

---

**COMMAND**

removesshkey

**DESCRIPTION**

The `removesshkey` command removes your Secure Shell (SSH) public key to disable key authentication.

**SYNTAX**

`removesshkey [ <user_name>... ]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

`<user_name>...`

Specifies the name of the user whose SSH key is removed. Multiple users can be specified. If not specified, the SSH key for the current user is removed.

**RESTRICTIONS**

Only Super level users are allowed to remove other users' keys.

**EXAMPLES**

The following example displays the removal of your SSH public key:

```
cli% removesshkey
```

**NOTES**

- After removing the user's SSH public key on the InServ Storage Server, the user cannot use the SSH key authentication to log in. The user must use name and password to log in.
- The `showuser -k` command can be used to display users that have SSH keys.

---

## COMMAND

removetask

## DESCRIPTION

The `removetask` command removes information about one or more completed tasks and their details.

## SYNTAX

The syntax of the `removetask` command can be one of the following:

- `removetask [options <arg>] -a`
- `removetask [options <arg>] -t <hours>`
- `removetask [options <arg>] <task_ID>...`

## AUTHORITY

Super, Edit

## OPTIONS

`-a`

Removes all tasks including details.

`-d`

Remove task details only.

`-f`

Specifies that the command is to be forced. You are not prompted for confirmation before the task is removed.

`-t <hours>`

Removes tasks that have not been active within the past `<hours>`, where `<hours>` is an integer from 1 through 240.

## SPECIFIERS

`<task_ID>...`

Allows you to specify tasks to be removed using their task IDs.

## RESTRICTIONS

None.

## EXAMPLES

The following example shows how to remove a task based on the task ID.

```
cli% removetask 2
Remove the following tasks?
2
select q=quit y=yes n=no: y
```

The following example shows how to remove all tasks, including details.

```
cli% removetask -a
Remove all tasks?
select q=quit y=yes n=no: y
```

## NOTES

- See the *InForm OS Concepts Guide* and *InForm OS CLI Administrator's Manual* for additional information and examples regarding task management and task management commands.
- With this command, the specified task ID and any information associated with it are removed from the system. However, task IDs are not recycled, so the next task started on the system uses the next whole integer that has not already been used. Task IDs roll over at 9999. The system stores information for the most recent 1000 tasks.

---

## COMMAND

removetemplate

## DESCRIPTION

The `removetemplate` command removes one or more Virtual Volume (VV), Logical Disk (LD), and Common Provisioning Group (CPG) templates.

## SYNTAX

```
removetemplate [options]{<template_name>|<pattern>}...
```

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-pat`

The specified patterns are treated as glob-style patterns and that all templates matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the `-f` option is specified. This option must be used if the `pattern` specifier is used.

## SPECIFIERS

`<template_name>`

Specifies the name of the template to be deleted, using up to 31 characters. This specifier can be repeated to remove multiple templates.

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to remove multiple templates. If this specifier is not used, the `<template_name>` specifier must be used. See [Glob-Style Pattern](#) on page 2.4 for more information

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example displays the forced removal of template vv1:

```
cli% removetemplate -f vv1
```

**NOTES**

None.



---

## COMMAND

`removeuser`

## DESCRIPTION

The `removeuser` command removes a user account from the system.

## SYNTAX

`removeuser [option] <username>...`

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

`<username>...`

Specifies a login name using any combination of letters and numbers. This argument can be repeated to specify multiple user names.

## RESTRICTIONS

- Access to all domains is required to use this command.
- Do not remove users `3parsvc` and `3paradm`.
- A user cannot remove oneself. The last user on the system cannot be removed.

## EXAMPLES

The following example displays the forced removal of `user1` from the system:

```
cli% removeuser -f user1
User removed.
```

**NOTES**

- Verify the removal of users by issuing the `showuser` command. See [showuser](#) on page 22.188 for additional information.
- The `removeuser` command does not affect currently connected users. If an attempt is made to remove a user that is currently connected, an error message will be returned.

---

**COMMAND**

`removeuserconn`

**DESCRIPTION**

The `removeuserconn` command removes user connections to the current system.

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

`removeuserconn [options] <user_ID> <user_name> <IP_address>`

**OPTIONS**

`-pat`

Specifies that the `<user_ID>`, `<user_name>`, and `<IP_address>` specifiers are treated as glob-style (shell-style) patterns and all user connections matching those patterns are removed. By default, confirmation is required to proceed with removing each connection unless the `-f` option is specified.

`-dr`

Specifies that the operation is a dry run and no connections are removed.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<user_ID>`

Specifies the ID of the user to be removed.

`<user_name>`

Specifies the name of the user to be removed.

`<IP_address>`

Specifies the IP address of the user to be removed.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the forced removal of user `user1` at IP address `127.0.0.1`:

```
cli% removeuserconn -f 2315 user1 127.0.0.1  
Removing user connection Id:2315 Name:user1 Addr:127.0.0.1
```

## NOTES

Because user connections can disappear from the time they are listed and the time they are removed, the `removeuserconn` command continues past errors while removing individual connections if the `-pat` option is specified.

---

## COMMAND

removevln

## DESCRIPTION

The `removevln` command removes a Virtual Volume's (VVs) SCSI Logical Unit Number (LUN) export definition from the system.

## SYNTAX

The syntax for the `removevln` command can be one of the following:

- `removevln [options] <VV_name | VV_set> <LUN>`  
`<N:S:P> | <host_name | host_set>`
- `removevln [options] <VV_name | VV_set> <LUN>`  
`<host_name | host_set> <N:S:P>`
- `removevln [options] <VV_name | VV_set> <LUN>`  
`<N:S:P> <host_name | host_set>`

## AUTHORITY

Super, Edit

## OPTIONS

`-novcn`

Specifies that a VLUN Change Notification (VCN) not be issued after removal of the VLUN.

- ◆ For direct connect or loop configurations, a VCN consists of a Fibre Channel Loop Initialization Primitive (LIP).
- ◆ For fabric configurations, a VCN consists of a Registered State Change Notification (RSCN) that is sent to the fabric controller.

`-pat`

Specifies that the `<VV_name>`, `<LUN>`, `<N:S:P>`, and `<host_name>` specifiers are treated as glob-style patterns and that all VLUNs matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the `-f` option is specified.

`-dr`

Specifies that the operation is a dry run and no VLUNs are removed.

–f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

<VV\_name | VV\_set>

Specifies the volume or volume set name, using up to 31 characters in length. The volume set name must start with `set :`.

<LUN>

Specifies the LUN ID to remove.

<N:S:P>

Specifies that exports to the specified port are removed. If this specifier is not used, the `host_name` specifier must be used.

`node`

Specifies the system port where `node` is a number from 0 through 7.

`slot`

Specifies the PCI bus slot in the node where `slot` is a number from 0 through 5.

`port`

Specifies the FCS port number of the card in PCI bus slot using 1 through 4.

<host\_name | host\_set>

Specifier requests that exports to the specified host or host set, named using up to 31 characters in length, be removed. The host set name must start with `set :`. If this specifier is not used, the `N:S:P` specifier must be used.

## RESTRICTIONS

- To remove a specific VLUN, you must supply the same specifiers and options that you used when that VLUN was created. Use the `showvlun -t` command to view all created VLUN specifiers and options.
- Issuing the `removevlun` command with the `-pat` option specified returns a request for confirmation to remove VLUNs, unless the `-f` option is specified.

## EXAMPLES

The following example deletes VLUNs for volume `vv0`, LUN 0, host `host1`:

```
cli% removevlun -f vv0 0 host1
```

**NOTES**

- If the VLUN to be removed is a matched-set, use the `port` and `host_name` specifiers.
- Verify the removal of VLUNs by issuing the `showvln` command. See [showvln](#) on page 22.195 for more information.



---

## COMMAND

`removevv`

## DESCRIPTION

The `removevv` command removes Virtual Volumes (VVs) from the system.

## SYNTAX

The syntax of the `removevv` command can be as follows:

- `removevv [options] <VV_name>|<pattern>...`
- `removevv -expired [options] [VV_name>|<pattern>...]`

## AUTHORITY

Super, Edit

## OPTIONS

`-pat`

Specifies that specified patterns are treated as glob-style patterns and that all VVs matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the `-f` option is specified. This option must be used if the `pattern` specifier is used.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-dr`

Specifies that the operation is a dry run and no VVs are removed.

`-stale`

Specifies that all stale VVs can be removed.

`-expired`

Remove specified expired volumes. This option cannot be used with the `-stale` option.

`-snaonly`

Remove the snapshot copies only.

`-cascade`

Remove all the descendent volumes as long as none has an active VLUN. It will remove any VLUN templates as long as there were no active VLUNs. It will remove the volumes from all the volume sets. If the `-expired` option is specified, all expired volumes and their descendent volumes will be removed regardless if they are expired or not. If the `-stale` option is specified, all stale volumes and their descendent volumes will be removed regardless if they are stale or not.

## SPECIFIERS

`<VV_name>`

Specifies the VV name, using up to 31 characters in length. This specifier can be repeated to remove multiple VVs. If this specifier is not used, the `pattern` specifier must be used.

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to remove multiple VVs. If this specifier is not used, the `<VV_name>` specifier must be used. See [Glob-Style Pattern](#) on page 2.4 for more information.

## RESTRICTIONS

- Issuing the `removevv` command with invalid VV names causes the command to exit without removing any VVs.
- Any VVs exported as Virtual Volume Logical Unit Numbers (VLUNs) are not removed and the command stops at that VV.
- Any VV that contains snapshots cannot be removed.

## EXAMPLES

The following example removes virtual volume `vv0`:

```
cli% removevv -f vv0
```

The following example removes all the expired volumes:

```
cli% removevv -f expired
```

The following example removes the volumes that start with test and are snapshot:

```
cli% removevv -f -snaonly -pat test*
```

The following example removes vv1\_snap if it is a snapshot and all its descendents:

```
cli% removevv -f snaonly -cascade vv1_snap
```

## NOTES

- By default, this command deletes any unused LDs that are not also part of a CPG. This is equivalent to using the `-ld` option. The `-ld` option still exists for backward compatibility.
- A newly created LD is guaranteed to be clean. Chunklets of LDs that are removed are cleaned before they are reused. However regions of an LD that were previously used (for example by another VV) can contain data from its previous use. If these regions of the LD are mapped to your user space of a VV, that data can be visible to the host to which the VV is exported.

If this is a concern, remove LDs when the VV is removed and use only newly created LDs for your user space. Previous data in LDs used for snapshot data space or snapshot admin space is not visible to you because these spaces are only visible after being written with new data.

- Verify the removal of VVs by issuing the `showvv` command. See [showvv](#) on page 22.201 for additional information.

---

**COMMAND**

removevvset

**DESCRIPTION**

The `removevvset` command removes a VV set or removes VVs from an existing set.

**SYNTAX**

`removevvset [options] <setname> [<vv>...]`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

**SPECIFIERS**

`<setname>`

Specifies the name of the VV set to remove.

`<VV>...`

Optional list of VV names that are members of the set. If no `<VV>` is specified, the VV set is removed, otherwise the specified `<VV>` is removed from the VV set.

**RESTRICTIONS**

None.

**EXAMPLES**

To remove a VV set:

```
cli% removevvset vvset
```

To remove a single VV from a set:

```
cli% removevvset vvset vv1
```

## NOTES

None.



# 20

## Service Commands

---

In this chapter

<code>servicecage</code>	<b>20.2</b>
<code>servicehost</code>	<b>20.5</b>
<code>servicemag</code>	<b>20.8</b>

---

## COMMAND

`servicecage`

## DESCRIPTION

The `servicecage` command is necessary when executing removal and replacement actions for a drive cage FC-AL module. The `startfc` or `unstartfc` subcommands are used to initiate service on a cage, and the `endfc` subcommand is used to indicate that service is completed.

## SYNTAX

The syntax for the `servicecage` command can be one of the following:

- `servicecage remove [-f] <cagename>`
- `servicecage startfc|unstartfc|endfc|resetfc|hresetfc|clearlog [-f] [-ovrd] [a|b]|[0|1] <cage_name>`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

Can vary for each subcommand as noted in the following section.

## SUBCOMMANDS

`startfc`

Prepare an FC-AL module for removal.

`unstartfc`

Stop the `startfc` subcommand. See [Notes](#) on page 20.4 for additional information about the `unstartfc` subcommand.

`endfc`

Indicates that service on the drive cage is completed and allows the system to resume use of the FC-AL module. Both `startfc` and `unstartfc` need this action to complete the service operation. See [Notes](#) on page 20.4 for additional information about the `endfc` subcommand. Permitted for Edit user in addition to Super and Service users.



### resetfc

- ◆ For DC1 and DC3 drive cages, soft resets the FC-AL module.
- ◆ For DC2 and DC4 drive cages, resets the cage. FC-AL specifier is ignored.

### hresetfc

- ◆ For DC1 and DC3 drive cages, hard reset the FC-AL module.
- ◆ For DC2 and DC4 drive cages, hard reset the cage. FC-AL specifier is ignored.

### remove

Removes the indicated drive cage (indicated with the <cage\_name> specifier) from the system. If this subcommand is used, the a|b or 0|1 port specifier is not required. This command fails when the cage has active ports or is in use. Only permitted for a Super user.

### clearlog

- ◆ For DC2 and DC4 drive cages, clear the log in the cage. FC-AL specifier is ignored.
- ◆ For DC3 drive cages, clearlog is not supported.

## OPTIONS

### -f

Forces the operation. When this option is not used, the command requires confirmation before proceeding.

### -ovrd

Forces the specified physical disk path offline even if it is the last remaining path.

## SPECIFIERS

### a|b

Specifies the side of the specified DC1 or DC3 drive cage to be serviced. This specifier is not required.

### 0|1

Specifies the side of the specified DC2 or DC4 drive cage to be serviced. This specifier is not required.

### <cage\_name>

Specifies the name of the drive cage to be serviced.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the commencement of Fibre Channel hot-plugging for drive cage cage0:

```
cli% servicecage startfc -f -ovrd cage0
```

## NOTES

- Issuing the `servicecage` command results in chunklet relocation, causing a dip in throughput.
- The `unstartfc` subcommand is provided if a mistake was made when issuing the `servicecage` command. The `unstartfc` subcommand stops the original command.
- After issuing the `startfc` or `unstartfc` subcommands, the `endfc` subcommand must be issued to indicate that service is completed and to restore the drive cage to its normal state.

---

## COMMAND

servicehost

## DESCRIPTION

The `servicehost` command executes removal and replacement actions for a host connecting to an InServ Storage Server port.

## SYNTAX

The syntax for the `servicehost` command can be one of the following:

- `servicehost list`
- `servicehost remove [-f] [<N:S:P> [<WWN_or_iSCSI_name>...]]`
- `servicehost copy [-f] <src_N:S:P> <WWN_or_iSCSI_name_pattern>  
<dest_N:S:P>`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

Can vary for each subcommand as noted in the following section.

## SUBCOMMANDS

`list`

Displays a list of hosts, ports, and Virtual Volume Logical Unit Numbers (VLUNs) for all inactive hosts. Permitted for all users.

`remove`

Removes an inactive host, as specified with the `<WWN_or_iSCSI_name>` specifier, from the indicated port (`<N:S:P>`) and its associated VLUNs. If the `<WWN_or_iSCSI_name>` specifier is not issued with the `remove` subcommand, all inactive hosts on the specified port and their LUNs are removed. If the `<N:S:P>` specifier is not used with the `remove` subcommand, all inactive hosts in the system and their LUNs are removed. See [Restrictions](#) on page 20.6 for additional information about the `remove` subcommand.

## copy

Copies all active VLUNs from the specified source port (as specified with `<src_N:S:P>`) from host WWNs or iSCSI names matching the specified pattern (`<WWN_or_iSCSI_name_pattern>`) to the destination port (as specified with `<dest_N:S:P>`). If necessary, the port is reset to target mode.

## OPTIONS

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

`<N:S:P>`

Specifies the host to be removed which is visible to the specified port. This specifier can only be used with the `remove` subcommand.

`<WWN_or_iSCSI_name>`

Specifies a host's World Wide Name (WWN). This specifier can only be used with the `remove` and `copy` subcommands. When used with the `remove` subcommand, this specifier can be repeated.

`<WWN_or_iSCSI_name_pattern>`

Specifies that the indicated WWN or iSCSI name is treated as a glob-style pattern. See [Glob-Style Pattern](#) on page 2.4 for more information.

`<src_N:S:P>`

Specifies the source port when exporting VLUNs. This specifier can only be used with the `copy` subcommand.

`<dest_N:S:P>`

Specifies the destination port when exporting VLUNs. This specifier can only be used with the `copy` subcommand.

## RESTRICTIONS

- Access to all domains is required to run this command.
- If the `<WWN_or_iSCSI_name>` specifier is not issued with the `remove` subcommand, all inactive hosts on the specified port and their LUNs are removed.

- If the <N:S:P> specifier is not used with the `remove` subcommand, all inactive hosts in the system and their LUNs are removed.

## EXAMPLES

The following example displays the creation of a host on port 0:2:1 for the export of VLUNs from port 2:1:1:

```
cli% servicehost copy 2:1:1 20000200000CF790 0:2:1
Are you sure you want to run servicehost?
select q=quit y=yes n=no: y
```

The following example displays the removal of an inactive host from port 2:1:1:

```
cli% servicehost remove 2:1:1 20000200000CF790
Removing inactive host 20000200000CF790 on port 2:1:1
Are you sure?
select q=quit y=yes n=no: y
```

## NOTES

None.

---

## COMMAND

`servicemag`

## DESCRIPTION

The `servicemag` command executes service on a drive magazine or disk.

## SYNTAX

The syntax for the `servicemag` command can be one of the following:

- `servicemag start [options] <cage_ID> <magazine>`
- `servicemag start [options] -pdid <PD_ID_0>...<PD_ID_3>`
- `servicemag resume|unmark [options] <cage_ID> <magazine>`
- `servicemag status [<cage_ID> <magazine>]`
- `servicemag clearstatus <cage_ID> <magazine>`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

`start`

Specifies that the `servicemag` command informs the system manager to log or relocate disks on a drive magazine so that the drive magazine can be removed for service.

`resume`

Specifies that the `servicemag` command informs the system manager that a drive magazine is replaced and that data services can be resumed.

`unmark`

Specifies that the `servicemag` operation is stopped and its internal state is reset. Since the `servicemag` operation is a multistep process, specifying `unmark` stops the `servicemag` operation at the completion of the current step.

Relocation of chunklets is considered one step and can take from several minutes to several hours (depending on number of chunklets) to complete. If `servicemag unmark` is issued during the relocation phase, the spin-down phase will indicate that the `servicemag start` operation Failed. This can be confirmed by issuing `servicemag status -d`. If the intent was to prevent `servicemag` from completing, issue `servicemag resume` to move the data back onto the drive(s) in this magazine once the `servicemag start` operation reports a failure. The `unmark` option should not be used without consulting with 3PAR engineering.

`status`

Specifies that the status of the `servicemag` command operations on a drive magazine are displayed.

`clearstatus`

Clears the log shown by the `servicemag` command `status` for the given cage and magazine.

## OPTIONS

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-wait`

Specifies that the `servicemag` command must be completed before exiting. If not specified, the `servicemag` command automatically exits before service begins and the remainder of the process proceeds in the background. This option is only valid for the `start` or `resume` operations.

`-log`

Specifies that write operations to chunklets of valid disks are written to a logging area when the magazine is out of service (or removed). When the disks return and the `servicemag resume` option is issued, the data is written from the logging Logical Disks (LDs) to those disks. Chunklets are relocated to free or spare space if their failures would result in a RAID set becoming invalid (for example, if two disks would be missing from a RAID 5 LD). All other used chunklets are placed in the logging mode. This option is only valid for the `start` operation.

`-nodisks`

Specifies that the serviced drive magazine's disk drives are valid and do not need to be replaced. This option can only be used with the `start` subcommand and `-log` option.

**-partial**

This option can only be used with the `resume` subcommand. Specifies that as many chunklets as possible are relocated. Error messages are printed for those chunklets that could not be relocated.

**-pdid <PD\_ID\_0>...<PD\_ID\_3>**

Specifies one to four physical disks (by physical disk ID) that need to be serviced or replaced. If the `-log` option is also specified, logging will only apply to the remaining disks on the magazine and not the ones specified in this option. The disks specified in this option will be vacated out to be replaced and will not be logged. This option can only be used with the `start` subcommand, and cannot be used with the `<cage_ID>` or `<magazine>` specifiers.

**-d**

Displays detailed status of a `servicemag` operation. If the `-d` option is excluded, a summary of the status is displayed. This option is only valid for the `status` subcommand.

**-dryrun**

Estimates the length of time a `servicemag start` or `servicemag resume` command will complete. If the `-dryrun` option is started before an actual `servicemag start` or `resume` command, the estimation will be based on empirical data with no I/O on the system. For a more accurate estimate, which will do runtime calculations based on the system load, use the `servicemag status <cage_ID> <magazine>` command after `servicemag start` or `servicemag resume` command has been initiated. This option is only valid for the `start` and `resume` operations.

**SPECIFIERS****<cage\_ID>**

Specifies the ID of the cage. Use the  `showcage` command to determine the system's drive cage IDs.

**<magazine>**

Specifies the drive magazine within the specified drive cage to be serviced based on the side of the drive magazine and drive magazine slot number.

- ◆ For drive chassis with a single drive cage (type DC2, DC4, and DC3), the valid syntax is `<position>` (the numeric position of the drive magazine). Position values for DC2 drive cages can be from 0 to 9. Position values for DC3 drive cages can be from 0 to 15.



## RESTRICTIONS

- Access to all domains is required to run this command.
- Users with Browse level authority can only issue the `servicemag status` command.

## EXAMPLES

The following example displays the suspension and resumption of data services on drive magazine 0 in drive cage 2:

```

cli% servicemag start -log -wait 2 0.0
Begin servicemag start -log 2 0.0...
... disks in mag : 2 0.0
...   valid disks:   wwn [2000000087043098] id [20]   diskpos [0]
...                   wwn [2000000087008150] id [21]   diskpos [1]
...                   wwn [20000000870042F6] id [22]   diskpos [2]
...                   wwn [2000000087007E6D] id [23]   diskpos [3]
...   not valid disks:
... mark disk wwn [2000000087043098] id [20] as non usable for ld allocation
... mark disk wwn [2000000087008150] id [21] as non usable for ld allocation
... mark disk wwn [20000000870042F6] id [22] as non usable for ld allocation
... mark disk wwn [2000000087007E6D] id [23] as non usable for ld allocation
... relocating chunklets to spare space
... relocating chunklets of fail sets after logging to spare space
... logging chunklets from pd wwn [2000000087043098] id [20]
... logging chunklets from pd wwn [2000000087008150] id [21]
... logging chunklets from pd wwn [20000000870042F6] id [22]
... logging chunklets from pd wwn [2000000087007E6D] id [23]
... spinning down disk wwn [2000000087043098] id [20]
... spinning down disk wwn [2000000087008150] id [21]
... spinning down disk wwn [20000000870042F6] id [22]
... spinning down disk wwn [2000000087007E6D] id [23]
... bypassing mag 2 0.0
... bypassed mag 2 0.0
servicemag start 2 0.0 -- Succeeded
cli%
cli% servicemag resume 2 0.0
Begin servicemag resume 2 0.0...
... onlooping mag 2 0.0
... checking for valid disks...
... disks in mag : 2 0.0
...   valid disks:   wwn [2000000087043098] id [20]   diskpos [0]
...                   wwn [2000000087008150] id [21]   diskpos [1]
...                   wwn [20000000870042F6] id [22]   diskpos [2]
...                   wwn [2000000087007E6D] id [23]   diskpos [3]
...   not valid disks:
... playback chunklets from pd wwn [2000000087043098] id [20]
... playback chunklets from pd wwn [2000000087008150] id [21]
... playback chunklets from pd wwn [20000000870042F6] id [22]
... playback chunklets from pd wwn [2000000087007E6D] id [23]
...   74 chunklets still waiting to be played back or relocating...
.... 18 chunklets still waiting to be played back or relocating..
... All chunklets played back / relocated.
... no chunklets to move
... marking pd wwn [2000000087043098] id [20] as usable for ld allocation
... marking pd wwn [2000000087008150] id [21] as usable for ld allocation
... marking pd wwn [20000000870042F6] id [22] as usable for ld allocation
... marking pd wwn [2000000087007E6D] id [23] as usable for ld allocation
servicemag resume 2 0.0 -- Succeeded

```

## NOTES

- Issuing the `servicemag` command results in chunklet relocation that causes a dip in throughput.
- When a `servicemag` command is issued with the `-log` option, all chunklets on the disks in the drive magazine being serviced are marked as `normal,smag`. This state indicates an active `servicemag` operation on the disks.
  - ◆ Any I/O on the chunklets marked `normal,smag`, changes the states to `logging` and I/O is written to the logging logical disks.
  - ◆ Issuing the `servicemag resume` command causes playback of the log. Any chunklets in the `logging` state enter `playback` state as their data is played back. After all the data is played back, the chunklets return to the `normal` state. Any chunklets in the `normal,smag` state return directly to the `normal` state.
  - ◆ Chunklet states can be checked by issuing either the `showldch` or `showpdch` commands (see [Show Commands](#) on page 22.1).
- By default, the `servicemag` command relocates all chunklets in the magazine to destinations starting first with local (such as on the owning node for the logical disk) spares, then local free space, then remote spare and finally remote free space.
- In the case when a drive needs to be replaced, the `-log` option should always be used in conjunction with the `-pdid` option.
- Replacing disks that have not had data completely relocated can lead to data loss. If the drives need to be replaced, the `-log` option should be used in conjunction with the `-pdid` option.
- For the `servicemag start` command only, instead of using the `<cage_ID>` `<magazine>` specifiers, the disk(s) to be serviced can also be specified using the `-pdid` option. The advantage is that the `servicemag resume` command will dismiss the old PDID when it successfully completes.



# 21

## Set Commands

---

### In this chapter

setalert	21.3
setauthparam	21.5
setbattery	21.12
setcage	21.14
setcim	21.16
setclienv	21.18
setcpgr	21.20
setdate	21.27
setdomain	21.31
setdomainset	21.33
sethost	21.35
sethostset	21.39
setlicense	21.40
setnet	21.42
setnode	21.47

setntp	21.49
setpassword	21.51
setpd	21.53
setrcopygroup	21.55
setrcopytarget	21.62
setsched	21.66
setsnmppw	21.68
setsshkey	21.70
setstatch	21.72
setstatpdch	21.74
setsys	21.75
setsysmgr	21.79
settemplate	21.82
setuser	21.84
setuseracl	21.86
setvv	21.88
setvvset	21.94

---

**COMMAND**

setalert

**DESCRIPTION**

The `setalert` command sets the status of system alerts.

**SYNTAX**

`setalert new|ack|fixed {<alert_ID>...|-a}`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-a`

Specifies that the status of all alerts be set to `new`, `acknowledged (ack)`, or `fixed`. If not specified, the `<alert_ID>` specifier must be specified on the command line.

**SPECIFIERS**

`<alert_ID>...`

Specifies that the status of a specific alert be set. This specifier can be repeated to indicate multiple specific alerts. Up to 99 alerts can be specified in one command. If not specified, the `-a` option must be specified on the command line.

`new|ack|fixed`

Specifies that the alert(s), as indicated with the `<alert_ID>` specifier or with option `-a`, be set as `new`, `acknowledged (ack)`, or `fixed`.

**RESTRICTIONS**

Access to all domains is required to run this command.

## EXAMPLES

The following example sets the status of all system alerts as `new`:

```
cli% setalert -a new
```

## NOTES

Verify the status of alerts by issuing the `showalert` command. See [showalert](#) on page 22.4 for additional information.



---

## COMMAND

setauthparam

## DESCRIPTION

The `setauthparam` command is used to set the authentication and authorization parameters.

## SYNTAX

The syntax of the `setauthparam` command can be one of the following:

```
setauthparam [-f] <param> <value>
setauthparam [-f] <map-param> <map-value>...
setauthparam [-f] -clear <param>...
setauthparam [-f] -clearall
```

## AUTHORITY

Super

## OPTIONS

-f

Does not ask for a confirmation before performing the operation.

-clearall

Clears all the authentication parameters.

-clear

Clears only the specified authentication parameters.

## SPECIFIERS

Specifiers for the `setauthparam` command can be issued as the following:

- `<param> <value>` - See [Table 21-1 on page 21.6](#).
- `<map_param> <map_value>` - See [Table 21-3 on page 21.10](#)
- `<param>` - See [Table 21-1 on page 21.6](#).

**Table 21-1.** Values for Specifiers <param> and <value>

<param>	<value>
ldap-server	Numeric IP address of the LDAP server.
ldap-server-hn	Indicates the host name of the LDAP server. This value must be set when the <code>ldap-reqcert</code> option is set or the <code>sasl-mechanism</code> option is set to GSSAPI. The <i>value</i> is the name of the LDAP server in its certificate or the value of the LDAP principal stored in the Kerberos database, and will usually be a fully-qualified domain name.
ldap-port	Indicates the port of the LDAP server (default: 389 for non-SSL, 636 for SSL).
ldap-ssl	To use SSL when communicating with the LDAP server, set the value to 1. (The default value is 0).
ldap-reqcert	Indicates whether a valid server certificate should be required in order to proceed (The default value is 0).
ldap-ssl-cacert	Indicates the path name of the file containing the certificate of the <i>Certificate Authority</i> that has issued the LDAP server's certificate, or a "–" to prompt you to enter the certificate text.
ldap-StartTLS	Set this parameter to one of the following: <ul style="list-style-type: none"> <li>■ <b>no</b> – to not request the server use StartTLS. Default.</li> <li>■ <b>try</b> – to request the server use StartTLS but does not require it to proceed.</li> <li>■ <b>require</b> – requests that the server uses StartTLS and continues only when it succeeds.</li> </ul>
binding	The LDAP binding type must be one of the following: <ul style="list-style-type: none"> <li>■ <b>simple</b> – use simple binding with the server.</li> <li>■ <b>SASL</b> - use a SASL mechanism that is expected by the server, with the mechanism set by the <code>sasl-mechanism</code> variable.</li> </ul>

**Table 21-1.** Values for Specifiers `<param>` and `<value>` (*continued*)

<code>&lt;param&gt;</code>	<code>&lt;value&gt;</code>
<code>user-dn-base</code>	When using simple binding, the authentication process attempts to bind the user to an entry in the server's Directory Information Tree (DIT). The Distinguished Name (DN) of the entry is a concatenation of the value of <code>user-attr</code> , "=", the username, ",", and the value of <code>user-dn-base</code> . If <code>group-obj</code> is set to <code>posixGroup</code> , the value of <code>user-dn-base</code> is also used as the base for searching for the user's <code>posixAccount</code> entry, regardless of binding type.
<code>user-attr</code>	Indicates the attribute used to form a DN for simple binding. When the attribute ends with a back slash, the DN is the concatenation of the value of the <code>user-attr</code> variable and the username. When the attribute does not end with a back slash, it is as described for the <code>user-dn-base</code> variable.
<code>sasl-mechanism</code>	When the binding is SASL, the SASL mechanism must be one supported by the LDAP server. The InServ allows the mechanisms of PLAIN, DIGEST-MD5, and GSSAPI.
<code>kerberos-server</code>	Indicates the numeric IP address of the Kerberos server if different from the LDAP server.
<code>kerberos-realm</code>	The Kerberos realm.
<code>allow-ssh-key</code>	Set this value to 1 to allow LDAP users to set a public SSH key with the <code>setsshkey</code> command (default 0). Clearing or setting the variable to 0 disables the setting of new keys for LDAP users but any existing keys remain until they are removed with the <code>removesshkey</code> variable. This parameter only affects LDAP users, not local users.
<code>groups-dn</code>	Indicates the base of the subtree in the DIT in which to search for objects that hold group information. It is mutually exclusively with <code>accounts-dn</code> variable.
<code>group-obj</code>	Indicates the <code>objectClass</code> attribute of a group object.
<code>group-name-attr</code>	The attribute in the group object that holds the group's name.

**Table 21-1.** Values for Specifiers <param> and <value> (*continued*)

<param>	<value>
member-attr	The attribute that holds the names of users in the group.
accounts-dn	Indicates the base of the subtree in the DIT in which to search for objects that hold account information. It is mutually exclusively with groups-dn variable.
account-obj	The objectClass attribute of an account object.
account-name-attr	The attribute of an account object that holds the user's username.
memberof-attr	The attribute that holds the name of a group of which the user is a member.

**Table 21-1.** Values for Specifiers `<param>` and `<value>` (*continued*)

<code>&lt;param&gt;</code>	<code>&lt;value&gt;</code>
<code>domain-name-attr</code>	<p>When set, the mapping of groups to domains is enabled. For a user that is a member of a group that maps to a privilege level, the value of <code>domain-name-attr</code> is used to look up an attribute in the group that holds the name of the domain.</p> <p>If the domain is too long or contains characters that are not allowed in a domain name, the name is truncated to the maximum length of a domain name and invalid characters are replaced with an underscore ( <code>_</code> ).</p>
<code>domain-name-prefix</code>	<p>When <code>domain-name-prefix</code> is set, the value of the attribute specified by <code>domain-name-attr</code> is a candidate domain name. The value of <code>domain-name-prefix</code> is a character string used to extract the domain name from the candidate. The value is an optional exclamation point ( <code>!</code> ) followed by a character string called the prefix. The exclamation point is a flag that means the presence of the prefix is required and is described more in the paragraphs that follow.</p> <p>The candidate domain name is searched for the presence of the prefix and if found, the domain name starts after the first occurrence of the prefix and stops before the first space or tab following it or at the end of the candidate domain name.</p> <p>If the prefix is not found, the behavior depends on the flag. If the exclamation point was not used (there is no flag), the candidate domain name becomes the domain name. If the flag is present, the candidate domain name is rejected and there is no domain name.</p> <p>As a last step, and as described for <code>domain-name-attr</code>, domain names can be truncated and have invalid characters replaced.</p>

Some examples of the effects of `domain-name-prefix` are shown in [Table 21-2 on page 21.10](#):

**Table 21-2.** Examples of Domain Name Prefix Effects

candidate	domain-name-prefix	result
dom1	ISDom=	dom1
ISDom=dom2	ISDom=	dom2
ISDom=dom3	!ISDom=	dom3
dom4	!ISDom=	

In the last case there is no resulting domain name because ISDom= does not appear in the candidate.

**Table 21-3.** Values for Specifiers <map\_param> and <map\_value>

<map-param>	<map-value>
super-map	A group name that grants the user the Super privilege level if the user is a member of that group. Multiple group names can be specified using multiple <map-value> arguments. A value of "*" matches any group name.
service-map	Same as super-map, but for the Service level.
edit-map	Same as super-map, but for the Edit level.
browse-map	Same as super-map, but for the Browse level.

## EXAMPLES

For a comprehensive example of the `setauthparam` command used during LDAP setup, see *LDAP Connection* in Chapter 4 of the *InForm OS CLI Administrator's Manual*.

## NOTES

- Users who have been provided with a password that allows successful binding with the LDAP server will nevertheless be denied access if they are not members of any of the groups specified by the map parameters.

- The matching of a user's groups with the mapping rules is done in the order of the mapping parameters provided previously. When there are multiple matches, the first match determines the user's privilege level.
- Domain names found with the use of `domain-name-attr` and `domain-name-prefix` are only potential domains and a user will only have privileges in those if they are actually existing domains. The `showdomain` command will list existing domains.
- The `showauthparam` command displays authentication parameter settings and the `checkpassword` command can be used to see how the parameters are used to bind with an LDAP server and search for data to determine the user's privilege level.
- When 3PAR Domains are enabled, you can only have Super or Service privilege levels for the domain all. Any other domain names are ignored for Super or Service level users. You can only have the Service privilege level when no other domains match for levels other than Super or Service. If other such domains match, the Service level match is ignored.

---

**COMMAND**

setbattery

**DESCRIPTION**

The `setbattery` command sets battery information such as the battery's expiration date, its recharging time, and its serial number. This information gives the system administrator a record or log of the battery age and battery charge status.

**SYNTAX**

```
setbattery [options <arg>] <node_ID> <powersupply_ID> <battery_ID>
```

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-s <serial_number>`

Specifies the serial number of the battery using a limit of eight alphanumeric characters.

`-x <exp_date>`

Specifies the expiration date of the battery (mm/dd/yyyy). The expiration date cannot extend beyond 2037.

`-l`

Specifies that the battery test log is reset and all previous test log entries are cleared.

`-r`

Specifies that the battery recharge time is reset and that 10 hours of charging time are required for the battery to be fully charged.

**SPECIFIERS**

`<node_ID>`

Specifies the node number where the battery is installed.



<powersupply\_ID>

Specifies the power supply number on the node using either 0 (left side from the rear of the node) or 1 (right side from the rear of the node).

<battery\_ID>

Specifies the battery number on the power supply where 0 is the first battery.



**NOTE:** The <powersupply\_ID> and <battery\_ID> specifiers can be obtained from the output of the `showbattery` command.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example resets the battery test log and the recharging time for a newly installed battery on node 2, power supply 1, and battery 0, with an expiration date of July 4, 2010:

```
cli% setbattery -x 07/04/2010 2 1 0
```

## NOTES

To view battery status information, issue the `showbattery` command. See [page 22.8](#) for information about the `showbattery` command.

---

## COMMAND

setcage

## DESCRIPTION

The `setcage` command enables service personnel to set or modify parameters for a drive cage.

## SYNTAX

```
setcage [subcommand <arg>]... <cage_name>
```

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

```
position <position>
```

Sets a description for the position of the cage in the cabinet, where `<position>` is a description to be assigned by service personnel (for example, `left-top`)

```
ps <model>
```

Sets the model of a cage power supply, where `<model>` is a model name to be assigned to the power supply by service personnel. This model name appears in the Model column of the `showcage -d` command output.

## OPTIONS

None.

## SPECIFIERS

```
<cage_name>
```

Indicates the name of the drive cage that is the object of the `setcage` operation.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example demonstrates how to assign `cage1` a position description of Side Left:

```
cli% setcage position Cabinet 0 Bay 5 Side Left cage1
```

The following example demonstrates how to assign model names to the power supplies in `cage1`. In this example, `cage1` has two power supplies (0 and 1). Both power supplies are assigned model name Magnetek.

```
cli% setcage ps 0 Magnetek ps 1 Magnetek cage1
```

## NOTES

- The parameters specified by the `setcage` command appear in the `showcage -d` output (see [page 22.17](#)).
- The power supply model cannot be modified if the information is automatically retrieved from the system.

---

## COMMAND

`setcim`

## DESCRIPTION

The `setcim` command sets the properties of the CIM server, including options to enable or disable the HTTP and HTTPS ports for the CIM server. The command also provides the ability to configure the port numbers associated with these two protocols. In addition, it allows users to enable or disable the SLP port.

## SYNTAX

`setcim [options]`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-f`

Forces the operation of the `setcim` command, bypassing the typical confirmation message.



**NOTE:** At least one of the following options are required when issuing the `setcim` command.

`-slp enable|disable`

Enables or disables the SLP port.

`-http enable|disable`

Enables or disables the HTTP port.

`-httpport <int>`

Sets the HTTP port (49152 - 65535). The default value is 5988.

`-https enable|disable`

Enables or disables the HTTPS port.

`-httpsport <int>`

Sets the HTTPS port (49152 - 65535). The default value is 5989.

## SPECIFIERS

None.

## RESTRICTIONS

- Access to all domains is required to run this command.
- You cannot disable both of the HTTP and HTTPS ports.
- You cannot set the same port number for both of the HTTP and HTTPS ports.

## EXAMPLES

To disable the HTTPS ports:

```
cli% setcim -https disable
Warning: The CIM server is active and will restart.
Are you sure you want to continue (Y/N)? Y
```

To enable the HTTPS port and set the HTTPS port number to 49153:

```
cli% setcim -https enable -httpsport 49153
Warning: The CIM server is active and will restart.
Are you sure you want to continue (Y/N)? Y
```

## NOTES

When the CIM server is active, a warning message appears to inform you of the current status of the CIM server and asks you for confirmation to continue or not. The `-f` option forces the action without a warning message.

---

**COMMAND**

setclienv

**DESCRIPTION**

The `setclienv` command sets the CLI environment parameters.

**SYNTAX**

`setclienv <parameter> <value>`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

The specifiers include the parameters and values to which the parameters should be set. Valid parameters and their values are as follows:

<parameter>	<value>
currentdomain	Enter one of the following: <ul style="list-style-type: none"><li>■ The name of the domain that you wish to set as the working domain for the current CLI session.</li><li>■ <code>-unset</code> to set no current domain.</li></ul>
listdom	Enter one of the following: <ul style="list-style-type: none"><li>■ 0 - (Default) Do not include the domain column in the output.</li><li>■ 1 - Include domain column where relevant.</li></ul>
csvtable	Enter one of the following: <ul style="list-style-type: none"><li>■ 0 - (Default) Normal table printing format.</li><li>■ 1 - Comma Separated Values (CSV) format.</li></ul>
nohdtot	Enter one of the following: <ul style="list-style-type: none"><li>■ 0 - (Default) Show header and total lines.</li><li>■ 1 - Does not show the header and total lines.</li></ul>

<parameter>	<value>
hafter	<nlines> - Specifies the number of lines of data to display before an output header is displayed. If <nlines> is 10 or more, print the header after every <nlines> of data. If <nlines> is less than 10, print only the header at the beginning.
editor	<p>NOTE: This parameter is only supported when connected via SSH.</p> <p>Specifies the command line editing mode. Enter one of the following:</p> <ul style="list-style-type: none"><li>■ emacs - (Default) Use emacs-style line editing.</li><li>■ vi - Use v-style line editing.</li></ul>

## RESTRICTIONS

None.

## EXAMPLES

In the following example, the CLI environment is set to display domains information:

```
cli% setclienv listdom 1
```

## NOTES

This command is only available when you are using a CLI shell or SSH.

---

## COMMAND

setcpg

## DESCRIPTION

The setcpg command modifies existing Common Provisioning Groups (CPG).

## SYNTAX

setcpg [options <arg>] <CPG\_name>

## AUTHORITY

Super, Edit, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

-sa <LD\_name>...

Specifies additional Logical Disks (LDs) that have already been created to be used for snapshot administration space allocation. The <LD\_name> argument can be repeated to specify multiple logical disks. This option is deprecated and will be removed in a subsequent release.

-sd <LD\_name>...

Specifies additional logical disks that have already been created to be used for snapshot data space allocation. The <LD\_name> argument can be repeated to specify multiple logical disks. This option is deprecated and will be removed in a subsequent release.

-aw <percent>

Specifies the percentage of used snapshot data space or snapshot administration space that, when reached, results in a warning alert. To disable the warning, enter 0. This option is deprecated and will be removed in a subsequent release.



**NOTE:** The following options, -sdgs, -sdgl, and -sdgw control the auto logical disk creation for the common provisioning group's snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow (enlarge) size setting options (-sdgs, -sdgl).



`-sdgs <size> [g|G|t|T]`

Specifies the growth increment, the amount of logical disk storage created, on each autogrow operation. The default growth increment is fixed at 32G, but the minimum growth increment varies according to the number of controller nodes in the system. If `<size>` is non-zero it must be 8G or bigger. A size of 0 disables the auto-grow feature. Size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). The following table displays the default and minimum growth increments per number of nodes:

**Table 21-4.** Growth Increment Per Number of Nodes

Number of Nodes	Default	Minimum
1-2	32G	8G
3-4	64G	16G
5-6	96G	24G
7-8	128G	32G

`-sdgl <size> [g|G|t|T]`

Specifies that the autogrow operation is limited to the specified storage amount that sets the growth limit. Storage size can be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). The default size is 0.

`-sdgw <size> [g|G|t|T]`

Specifies that the threshold of used logical disk space, when exceeded, results in a warning alert. This sets the growth warning. Size threshold be specified in megabytes (default) or in gigabytes by using the `g|G` parameter and terabytes by using the `t|T` parameter. To specify the size in gigabytes, enter `g` or `G`, for terabytes, enter `t` or `T`, directly after the specified size (no space). The default size is 0.



**NOTE:** The following options, `-t`, `-ssz`, `-rs`, `-ss`, `-ha`, `-ch`, and `-p` are used to control auto logical disk creation.

`-t <RAID_type>`

Specifies the RAID type of the logical disk: `r0` for RAID-0, `r1` for RAID-1, `r5` for RAID-5, or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklet>`

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row using an integer from 1 through 2147483647. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port|cage|mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the characteristics of the chunklets, either `first` (attempt to use the lowest numbered available chunklets) or `last` (attempt to use the lowest numbered available chunklets). If no argument is specified, the default characteristic is `first`.

`-p <pattern>`

Specifies a pattern for disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to all disks of device type specified at creation time. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The following arguments can be specified as patterns for this option:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0–7). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0–7). The primary path of the disks must be on the specified PCI slot(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–4). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-3). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following options, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid` and `-devtype` are used to select the disks that are used to create common provisioning groups based on the characteristics of the disk.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and connected to node 2 through their primary path.

`-sax <LD_name> [ , <LD_name> . . . ]`

Specifies that the logical disk, as identified with the `<LD_name>` argument, used for snapshot administration space allocation be removed. The `<LD_name>` argument can be repeated to specify multiple logical disks.

`-sdx <LD_name> [ , <LD_name> . . . ]`

Specifies that the logical disk, as identified with the `<LD_name>` argument, used for snapshot data space allocation be removed. The `<LD_name>` argument can be repeated to specify multiple LDs.

`-name <newname>`

Specifies the name of the Common Provisioning Group (CPG) to be modified to. `<newname>` can be up to 31 characters in length.

## SPECIFIERS

<CPG\_name>

Specifies the name of the CPG being modified.

## RESTRICTIONS

- Access to all domains is required to run this command.
- This command sets the parameters that will be used for future grows and it will OVERWRITE all previously set parameters.

## EXAMPLES

The following example displays the modification of the auto-growth parameters for CPG `cpg1`:

```
cli% setcpg -sdgs 16 -sdgl 48 -sdgw 36 cpg1
```

The following example removes the `-cpg` pattern:

```
cli% setcpg -p -cg ""
```

## NOTES

- With this command, you can change the device type of a CPG from logical disks of one device type to logical disks of another device type (device types are Fibre Channel (FC), Nearline (NL), or Solid State Drive (SSD)). This implies that, within a CPG, one can have logical disks of type FC, type NL, and type SSD. However, this is only permitted so that, if a user wants to change the type of a CPG from FC to NL or SSD, they can first change the new logical disk creation characteristics using `setcpg` and then use region moves to change the device types of existing logical disks. For this same reason, users are permitted to add logical disks of a different device type to a CPG.
- The options `-sdgs`, `-sdgl` and `-sdgw` control the auto LD creation for the CPG's snapshot data regions. Auto LD creation occurs when the amount of free LD space falls below the specified grow (enlarge) size setting options (`-sdgs`, `-sdgl`).
- If auto-grow is enabled, new LDs will be created automatically in addition to any LDs specified with `-sa` or `-sd` options.
- Enter `""` after the `<pattern>` option to remove the pattern.
- Use `movetodomain` command to move a CPG from one domain to another.

---

## COMMAND

setdate

## DESCRIPTION

The `setdate` command allows you to set the system time and date on all nodes.

## SYNTAX

The syntax for the `setdate` command can be one of the following:

- `setdate <MMDDhhmm> [ [ <CC> ] <YY> ] [ <.ss> ]`
- `setdate -tzlist [group]`
- `setdate -tz <tzname>`

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-tzlist [group]`

Displays a timezone within a group, if a group is specified. If a group is not specified, displays a list of valid groups.

`-tz <tzname>`

Sets the timezone on all nodes. The option must have a valid `tzname` from the list provided with the `-tzlist` command.

## SPECIFIERS

`<MMDD>`

Specifies the month (`MM`) and day (`DD`).

`<hhmm>`

Specifies the hour (`hh`) and minute (`mm`) on a 24-hour clock.

<CC>

Specifies the century (CC) and cannot be used unless a year is specified (YY). This specifier is not required.

<YY>

Specifies a year (YY). This specifier is not required.

<.ss>

Specifies seconds (ss). This specifier is not required.

Specifiers can only be used in combinations as listed in [Restrictions](#) in the section that follows.

## RESTRICTIONS

- Access to all domains is required to run this command.
- Specifiers must be provided in one of the combinations listed as follows when issuing the `setdate` command:
  - ◆ MMDDhhmm
  - ◆ MMDDhhmmYY
  - ◆ MMDDhhmmCCYY
  - ◆ MMDDhhmm.ss
  - ◆ MMDDhhmmYY.ss
  - ◆ MMDDhhmmCCYY.ss

## EXAMPLES

The following example displays the current date on the node:

```
cli% showdate
Node Date
4    Mon Oct 10 16:14:28 PDT 2005
5    Mon Oct 10 16:14:28 PDT 2005
6    Mon Oct 10 16:14:28 PDT 2005
7    Mon Oct 10 16:14:28 PDT 2005
```



The following example displays the timezones with the `-tzlist` option:

```
cli% setdate -tzlist
Africa
America
Antarctica
Arctic
Asia
...
```

The following example narrows down the list to the required timezone of `Etc`:

```
cli% setdate -tzlist Etc
Etc/GMT
Etc/GMT+0
Etc/GMT+1
...
```

The following example shows the timezone being set:

```
cli% setdate -tz Etc/GMT
Timezone set successfully.
```

The following example verifies the timezone is set to the required setting:

```
cli% showdate
Node Date
4 Mon Oct 10 23:14:52 GMT 2005
5 Mon Oct 10 23:14:52 GMT 2005
6 Mon Oct 10 23:14:52 GMT 2005
7 Mon Oct 10 23:14:52 GMT 2005
```

## NOTES

- Specifies must be provided in one of the combinations listed as follows when issuing the `setdate` command:
  - ◆ `MMDDhhmm`
  - ◆ `MMDDhhmmYY`
  - ◆ `MMDDhhmmCCYY`
  - ◆ `MMDDhhmm.ss`
  - ◆ `MMDDhhmmYY.ss`
  - ◆ `MMDDhhmmCCYY.ss`
- Check node dates by issuing the `showdate` command. See [showdate](#) on page 22.33 for additional information.

---

## COMMAND

setdomain

## DESCRIPTION

The `setdomain` command sets the parameters and modifies the properties of a domain.

## SYNTAX

`setdomain [options <arg>] <domain_name>`

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-name <name>`

Changes the name of the domain.

`-comment <comment>`

Specifies comments or additional information for the domain. The comment can be up to 511 characters long and must be enclosed in quotation marks. Unprintable characters are not allowed within the `<comment>` specifier.

`-vvretentiontimemax <value>[h|H|d|D]`

Specifies the maximum value that can be set for the retention time of a volume in this domain. `<time>` is a positive integer value and in the range of 0 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

To remove the maximum volume retention time for the domain, enter `-vvretentiontimemax " "`. As the result, the maximum volume retention time for the system is used instead.

To disable setting the volume retention time in the domain, enter 0 for `<time>`.

**SPECIFIERS**

<domain\_name>

Indicates the name of the domain.

**RESTRICTIONS**

You need access to all domains to run this command.

**EXAMPLE**

In the following example, the name of a domain named `Domain1` is changed to `DomainX`:

```
cli% setdomain -name DomainX Domain1
```

The following example displays the addition of a comment to the domain `Engineering`:

```
cli% setdomain -comment "This is a comment for engineering." Engineering
```

**NOTES**

For moving objects into a domain, use the `movetodomain` command.

---

**COMMAND**

setdomainset

**DESCRIPTION**

The `setdomainset` command sets the parameters and modifies the properties of a domain set.

**SYNTAX**

```
setdomainset [options <arg>] <setname>
```

**AUTHORITY**

Super, Edit

**OPTIONS**

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters in length. Unprintable characters are not allowed.

`-name <newname>`

Specifies a new name for the domain set. The name can be up to 27 characters in length.

**SPECIFIERS**

`<setname>`

Specifies the name of the domain set to modify.

**RESTRICTIONS**

None.

**EXAMPLE**

To rename a set from foo to bar:

```
cli% setdomainset -name bar foo
```

To change the comment on set bar:

```
cli% setdomainset -comment "This used to be set foo" bar
```

## NOTES

None.

---

## COMMAND

sethost

## DESCRIPTION

The `sethost` command sets properties on existing system hosts, including options to annotate a host with descriptor information such as physical location, IP address, operating system, model, and so on. The command also provides the ability to configure or remove iSCSI CHAP authentication information.

## SYNTAX

The syntax for the `sethost` command can be one of the following:

- `sethost [options <arg>] <host_name>`
- `sethost initchap [-f] [options <arg>] <secret>`  
`{<host_name>|<pattern>}...`
- `sethost targetchap [-f] [options <arg>] <secret>`  
`{<host_name>|<pattern>}...`
- `sethost removechap [-target] [-f] {<host_name>|<pattern>}...`
- `sethost clearagent <WWN|iscsi_name>`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

`initchap`

Sets the initiator CHAP authentication information on one or more hosts.

`targetchap`

Sets the target CHAP authentication information on one or more hosts.

### removechap

Removes CHAP authentication information on one or more hosts. By default, this removes all CHAP information for all specified hosts. Using this subcommand with the `-target` option removes only target CHAP information.

### clearagent

Clears any host agent data associated with the host.

## OPTIONS

`-loc <location>`

Specifies the location of the host.

`-ip <IP_address>`

Specifies the IP address of the host.

`-os <OS>`

Specifies the operating system running on the host.

`-model <model>`

Specifies the model of the host.

`-contact <contact>`

Specifies the contact information for the host.

`-comment <comment>`

Specifies any additional information for the host.

`-name <hostname>`

Specifies the new name of the host up to 31 characters in length.

`-persona <hostpersonaval>`

Sets the host persona that specifies the personality for all ports which are part of the host set. This selects certain variations in SCSI command behavior which certain operating systems expect. The `<hostpersonaval>` is the host persona ID number with the desired capabilities. These can be seen with `showhost -listpersona`.

The following options are for use with the `initchap` and `targetchap` subcommands:

`-f`

Do not ask for confirmation before performing the operation.



`-chapname <chapname>`

Used to specify the initiator or target CHAP name. If this option is not specified, then the initiator CHAP name defaults to the host name and the target CHAP name defaults to the 3PAR System name.

`-hex`

The CHAP secret is treated as a hex number.

The following options are for use with the `removechap` subcommand:

`-target`

Removes only the target CHAP authentication.

## SPECIFIERS

`<host_name>`

Name of the host with a maximum of 31 characters in length.

`<pattern>`

Specifies that the properties are set for all hosts matching the specified pattern.

`<secret>`

The CHAP secret for the host or the target. If `-hex` is specified, it is treated as a hex number. Otherwise it should be a printable ASCII string 12 to 16 characters in length with no spaces, or 16 bytes in HEX.



**NOTE:** The CHAP configuration operations are applied to all hosts whose names match one or more of the specified `<hostname>` or `<pattern>`. The patterns are treated as glob-style (shell-style) patterns (see Help on `sub, globpat`).

`<WWN|iSCSI_name>`

The World Wide Name (WWN) or iSCSI name of the host to clear the host agent data of.

## RESTRICTIONS

Access to all domains is required for the `clearagent` subcommand.

## EXAMPLES

The following example change the settings of a host:

```
cli% sethost -contact "Joe Smith" -model "Sun Ultra 60" queasy10
cli% sethost initchap "MyChapSecret" queasy10
cli% sethost targetchap -hex "30313233343536373839303132333435" queasy10
cli% sethost removechap -target queasy10
cli% sethost clearagent 210100E08B32A58A
cli% sethost -persona 1 queasy10
```

## NOTES

- The CHAP configurations are applied to all hosts whose names match one or more of the specified <host\_name> or <pattern> specifiers. The patterns are treated as glob-style (shell-style) patterns (see help on sub,globpat).
- The options that allow for adding descriptive information are for annotation purposes only; the storage server does not actively use the information provided here.
- Remove descriptors by passing an empty string to the command.
- Verify modification of host properties by issuing the `showhost` command.

---

**COMMAND**

sethostset

**DESCRIPTION**

The `sethostset` command sets the parameters and modifies the properties of a host set.

**SYNTAX**

`sethostset [options <arg>] <setname>`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-comment <comment>`

Specifies a comment relating to the set.

`-name <newname>`

Specifies a new name for the host set.

**SPECIFIERS**

`<setname>`

Specifies the name of the host set to modify.

**RESTRICTIONS**

None.

**EXAMPLES**

To rename a set from foo to bar:

```
cli% sethostset -name bar foo
```

To change the comment on a set:

```
cli% sethostset -comment "This used to be set foo" bar
```

**NOTES**

None.

---

**COMMAND**

setlicense

**DESCRIPTION**

The setlicense command sets the license key information.

**SYNTAX**

setlicense [options <arg>]

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

-f <filename>

Specifies the file from which the license key is read.

-noconfirm

Specifies that the system does not prompt for confirmation for the new license key.

**SPECIFIERS**

None.

**RESTRICTIONS**

- Access to all domains is required to run this command.
- The entered license key is accepted only if it is recognized as a valid key. A valid license key includes an appropriate serial number and is associated with the number of nodes in the system for which the license key is being entered.
- 3PAR's license terms and conditions must be accepted before proceeding with this command.

## EXAMPLES

The following example displays the setting of a license key:

```
cli% setlicense
```

If this software is being provided to you for a limited evaluation period, then your license shall be governed by the terms of 3PAR's standard Master Purchase and License Agreement which can be found at <http://www.3par.com/mppla> (or as otherwise agreed between us) with the exception that the term of the license shall expire upon the earlier of the evaluation period notified to you or 120 days. Upon expiration of the license, you must cease using the software and 3PAR reserves the right to disable the software without notice. By using or activating the software you are agreeing to these terms.

Do you agree to these terms and conditions? y=yes n=no: yes

Please enter the new license key below. When finished, press enter twice. If the key is entered by hand, note that characters other than letters and numbers are ignored, and the key is not case-sensitive.

```
60R3-0C1G...
```

## NOTES

- This command prompts for a new license key. To finish entering the license key, press ENTER on a blank line.
- When the license key is being interpreted, all characters other than letters (without case-sensitivity) and numbers are ignored, and the letters are not case-sensitive.
- After the new license key has been entered, the changes between the existing license key and the new license key are displayed. There is a prompt to confirm the changes unless the `-noconfirm` option is given, in which case the information is not displayed, and the new license key is entered immediately.

---

**COMMAND**

setnet

**DESCRIPTION**

The setnet command sets the administration network interface configuration.

**AUTHORITY**

Service, Super



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

The syntax for the setnet command can be one of the following:

- `setnet startaddr <old_IP> <new_IP> <new_netmask>`
- `setnet startgateway <new_gateway>`
- `setnet finish [-f]`
- `setnet abort`
- `setnet cleargateway [-f]`
- `setnet speed <IP_addr> auto|<mbps> <duplex>`
- `setnet failoverping <IP_addr> <ping_addr>|none`
- `setnet ntp none|<server_addr>`
- `setnet changenode [<node_ID>]`

**SUBCOMMANDS**

startaddr

Specifies that the system start switching the old IP address (<old\_IP> specifier) to the new IP address (<new\_IP> specifier) with the specified netmask (<new\_netmask> specifier).

### `startgateway`

Specifies that the gateway be immediately set to the specified IP address if no gateway is currently defined, or switched to the specified IP address if currently defined, in both cases indicated with the (`<new_gateway>` specifier).

### `finish`

Specifies that outstanding changes from the `startaddr` and `startgateway` subcommands be completed.

### `abort`

Specifies that any attempt to configure a new IP address or gateway fails. The system returns to its previous state.

### `cleargateway`

Specifies that the existing system gateway is removed.

### `speed`

Specifies that the network interface is set to the specified speed and duplex as indicated with the `<mbps>` and `<duplex>` specifiers.

### `failoverping`

Specifies that on IP failover, a ping is sent to the specified IP address as indicated with the `<ping_addr>` specifier.

### `ntp none|ntp <serveraddr>`

Specifies the NTP server the system should use to synchronize its clocks. The server must be specified as an IP address.

### `changenode`

Forces the system to change which node has an active Ethernet interface. If a node ID is specified, it switches to that node. Otherwise, any node with a connected Ethernet interface is chosen.



**WARNING:** If successful, this command causes any 3PAR management application clients currently connected to lose their connection.

## OPTIONS

`-f`

Specifies that the operation is forced even if verification has not occurred. This option can only be used with the `finish` and `cleargateway` subcommands.

## SPECIFIERS

`<old_IP>`

Specifies an existing IP address that is to be changed. This specifier is used in conjunction with the `<new_IP>` and `<new_netmask>` specifiers and can only be used with the `startaddr` subcommand.

`<new_IP>`

Specifies a new IP address to which the system is configured. This specifier is used in conjunction with the `<old_IP>` and `<new_netmask>` specifiers and can only be used with the `startaddr` subcommand.

`<new_netmask>`

Specifies a new netmask to which the system is configured. This specifier is used in conjunction with the `<old_IP>` and `<new_IP>` specifiers and can only be used with the `startaddr` subcommand.

`<new_gateway>`

Specifies the IP address of the new gateway for the system. This specifier can only be used with the `startgateway` subcommand.

`auto|<mbps> <duplex>`

Specifies that the speed of the network interface is either auto negotiated (`auto`), or specified manually using the `<duplex>` and `<mbps>` specifiers.

`<mbps>`

Specifies the speed of the network interface. Valid values are either 10 or 100. This specifier can only be used with the `<duplex>` specifier and with the `speed` subcommand.

`<duplex>`

Specifies the duplex of the network interface. Valid values are either `half` or `full`. This specifier can only be used with the `<mbps>` specifier and with the `speed` subcommand.



<IP\_addr>

Specifies the IP address of the node. This specifier can only be used with the `failoverping` subcommand.

<ping\_addr> | none

Specifies that during an IP failover, a ping either be sent to the specified IP address (<ping\_addr>) or not sent at all (none). This specifier can only be used with the `failoverping` subcommand.

[ <node\_ID> ]

Specifies the node, by ID, that has an active Ethernet interface. This specifier can only be issued with the `changenode` subcommand. This specifier is not required.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the switching of the old IP address with a new IP address and netmask:

```
cli% setnet startaddr 10.0.23.42 192.168.5.218 255.255.252.0
Change of IP address successfully started.
```

The following example displays the gateway being cleared:

```
cli% setnet cleargateway

If the machine that the 3Par CLI is running on is not on the same subnet as
the 3Par InServ it is connected to, clearing the gateway will render the
InServ unreachable from that machine, and any future connections will need
to be made from a system which is on the same subnet as the InServ system.

Are you sure you want to clear the gateway (y/n)?
y
Gateway modified successfully.
```

## NOTES

- If successful, this command causes any CLI or GUI clients currently connected to lose their connection.

- To make it possible to change the network configuration without running the risk of losing contact with the system because of misconfiguration, the `setnet` command uses a two step process. When a new IP address is specified with the `setnet startaddr` command, the system is configured to listen as both the old and new IP addresses. When a new gateway is specified with the `setnet gateway` command, the system switches between the old and new gateways when it sees packets addressed to it being routed through those gateways. After a connection has been made with the new configuration, the `setnet finish` command can be used to remove the old configuration. While in the middle of this process, additional work must be done by the system. It is preferable to run the `setnet finish` command after the new configuration has been verified.
- When changing gateways, starting a CLI connection takes longer than usual, as the first reply packet is typically routed through the previously used gateway address.

---

**COMMAND**

setnode

**DESCRIPTION**

The `setnode` command sets the properties of the node components such as the serial number of the power supply.

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

```
setnode ps <PS_ID> [options] <node_ID>
```

**SUBCOMMAND**

ps

Sets the power supply properties.

**OPTIONS**

`-s <serial_number>`

Specifies the serial number up to eight characters in length.

**SPECIFIERS**

`<PS_ID>`

Specifies the power supply ID.

`<node_ID>`

Specifies the node ID.

**RESTRICTIONS**

- Access to all domains is required to run this command.
- At least one option must be specified.

**EXAMPLE**

The following example displays the node power supply:

```
cli% shownode -ps
Node PS -Serial- -PSState-- FanState ACState DCState -BatState- ChrgLvl(%)
  0  0  --          NotPresent --          --          --          NotPresent      0
  0  1  FFFFFFFF OK           OK           OK           OK           OK           0
  1  0  FFFFFFFF OK           OK           OK           OK           OK           0
  1  1  --          NotPresent --          --          --          NotPresent      0
cli%
cli% setnode ps 1 -s 12345678 0
cli% shownode -ps
Node PS -Serial- -PSState-- FanState ACState DCState -BatState- ChrgLvl(%)
  0  0  --          NotPresent --          --          --          NotPresent      0
  0  1  12345678 OK           OK           OK           OK           OK           0
  1  0  FFFFFFFF OK           OK           OK           OK           OK           0
  1  1  --          NotPresent --          --          --          NotPresent      0
cli%
cli% setnode ps 0 -s aabbccdd 1
cli% shownode -ps
Node PS -Serial- -PSState-- FanState ACState DCState -BatState- ChrgLvl(%)
  0  0  --          NotPresent --          --          --          NotPresent      0
  0  1  12345678 OK           OK           OK           OK           OK           100
  1  0  AABCCDD OK           OK           OK           OK           OK           100
  1  1  --          NotPresent --          --          --          NotPresent      0
cli%
cli% setnode ps 0 -s aabbccdde 1
Error: The serial number aabbccdde is too long, should be less than 9
characters.
```

**NOTE**

None.

---

**COMMAND**

setntp

**DESCRIPTION**

The `setntp` command sets the system Network Time Protocol (NTP) server. This command is deprecated and will be removed in a subsequent release.

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**SYNTAX**

`setntp <server_IP_adress> | none`

**OPTIONS**

None.

**SPECIFIERS**

`<server_IP_address>`

Specifies the IP address of the NTP server to which the InServ system synchronizes its internal clocks.

`none`

Specifies that the InServ system should not synchronize its internal clocks with an external NTP server. Instead, the InServ system synchronizes its clocks internally.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example displays the setting of the system NTP server:

```
cli% setntp 192.168.1.1
NTP server successfully updated.
```

## NOTES

None.

---

## COMMAND

setpassword

## DESCRIPTION

The `setpassword` command allows a user with Super level privileges to change the password for any user and create a password file on a client. Edit-, browse-, or service-level users can use the `setpassword` command to change their own passwords or save their password files on a client.

## SYNTAX

`setpassword [options <arg>]`

## AUTHORITY

Super, Edit, Browse, Service

Only the Super user can edit the password of a different user.

## OPTIONS

`-save|-saveonly`

This option cannot be used independently of the `-file` option.

`-save`

Specifies that the password, as specified with the `-file` option, is saved on both the InServ storage system and on a client of the system.

`-saveonly`

Specifies that the password, as specified with the `-file` option, is only saved on the client.

`-file <pwfile>`

Specifies the password file to be saved. The `<pwfile>` option can be any valid file name in the client system. This option cannot be used independently of the `-save` or `-saveonly` options.

`-u <username>`

Specifies the login name of the user whose password is being changed. If a login name is not specified, the command defaults to the current user.

## SPECIFIERS

None.

## RESTRICTIONS

- Only a user with Super level privileges can set another user's password.
- Passwords can be no longer than eight non-null characters, and must be at least six non-null characters long.

## EXAMPLES

The following example shows how to set a user's (user1) password on an InServ storage system and on a client:

```
cli% setpassword -save -file <insertfile> user1
```

The following example displays the prompts encountered when changing a user's (user1) password:

```
cli% setpassword -u user1
password:
Old password:
```

## NOTES

- The format of the entry in the file is <username> <encrypted\_password>. This file may be referenced by the TPDPWFILE environment variable or -pwf command line option for subsequent commands.
- The -save or -saveonly option requires the -file option.
- Without any options, the command will prompt to change the invoking user's password on the InServ system.
- Changing a user's password has no effect on SSH access if the user has set a valid key with the setsshkey command. Until the SSH key is removed the user will not have to provide the new password.
- For additional information about password files, see the *InForm OS Concepts Guide* and the *InForm OS CLI Administrator's Manual*.



---

**COMMAND**

setpd

**DESCRIPTION**

The `setpd` command marks a Physical Disk (PD) as allocatable or non allocatable for Logical Disks (LDs).

**SYNTAX**

```
setpd ldalloc on|off <PD_ID>...
```

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**SUBCOMMANDS**

```
ldalloc on|off
```

Specifies that the PD, as indicated with the `PD_ID` specifier, is either allocatable (`on`) or nonallocatable for LDs (`off`).

**OPTIONS**

None.

**SPECIFIERS**

```
<PD_ID>...
```

Specifies the PD identification using an integer.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example displays PD 0 marked as non allocatable for LDs.

```
cli% setpd ldalloc off 0
```

**NOTES**

- This command can be used when the system has disks that are not to be used until a later time.
- Verify the status of PDs by issuing the `showpd -s` command. See [showpd](#) on page 22.91.

---

## COMMAND

setrcopygroup

## DESCRIPTION

The `setrcopygroup` command performs the following actions:

- Sets the policy of the Remote Copy volume group for dealing with I/O failure and error handling.
- Switches the direction of transfer between volume groups.
- Sets a resynchronization period for volume groups in asynchronous periodic mode.
- Sets the group's mode.

## SYNTAX

The syntax for the `setrcopygroup` command can be one of the following:

- `setrcopygroup pol [option] <policy> <group_name>`
- `setrcopygroup period <period_value> <target_name> <group_name>`
- `setrcopygroup mode <mode_value> <target_name> <group_name>`
- `setrcopygroup <dr_operation> [options] <target_name|group_name>...`

## AUTHORITY

Super, Edit

## SUBCOMMANDS

`pol`

Sets the policy of the Remote Copy volume group for dealing with I/O failure and error handling.

`period`

Specifies that groups that are in asynchronous periodic mode should be periodically synchronized in accordance with the specified `<period_value>`.

`mode`

Specifies the mode to which the volume group is set.

<dr\_operation>

Specifies the operation of the group(s). Valid operations are:

`reverse`

Changes the natural and current direction of all specified groups. The operation is mirrored resulting in a direction change on both systems. This option is very flexible depending on the options that are provided, however it should not be used as a part of the normal disaster recovery process.

`failover`

Changes secondary volume groups to primary volume groups on the active system in the event of a server failure. If the group has multiple targets it will also attempt to pull more recent data from other targets and start remote copy to those targets when complete.

`recover`

Used for groups on which the failover operation has already been run. Changes matching primary volume groups on the backup system to secondary volume groups and then starts and synchronizes all groups.

`restore`

Used on groups on which the recover operation has already been run. Returns all groups to their natural direction and starts them.

## OPTIONS

`-t`

Specifies that the `setrcopygroup <dr_operations>` command be applied to all relevant groups of the indicated target. If no group names are specified then all groups in that target that are in the correct state for the specified subcommand will have it applied to them. For groups with multiple targets this option must be specified for the `recover`, `restore`, `reverse` and `pol` subcommands.

`-f`

Does not ask for confirmation for disaster recovery commands.

`-nostart`

Specifies that groups are not started after role reversal is completed. This option can be used for `failover`, `recover` and `restore` subcommands.

**-nosync**

Specifies that groups are not synchronized after role reversal is completed through the `recover`, `restore` and `failover` specifiers.

**-discard**

Specifies not to check a group's other targets to see if newer data should be pushed from them if the group has multiple targets. The use of this option can result in the loss of the most recent changes to the group's volumes and thus should be used carefully. This option is only valid for the `failover` subcommand.

**-nopromote**

This option is only valid for the `failover` and `reverse` specifiers. When used with the `reverse` specifier it indicates that the synchronized snapshots of groups that are switched from primary to secondary not be promoted to the base volume. When used with the `failover` specifier it indicates that snapshots of groups that are switched from secondary to primary should not be promoted to the base volume in the case where all volumes of the group were synced to the same time point. The incorrect use of this option can lead to the primary and secondary volumes not being consistent.

**-nosnap**

Specifies that snapshots are not taken of groups that are switched from secondary to primary. Additionally, existing snapshots are not deleted if groups are switched from primary to secondary. The use of this option may result in a full synchronization of the secondary volumes. This option can be used for `failover` and `reverse` subcommand.

**-stopgroups**

Specifies that groups are stopped before running the `reverse` subcommand.

**-local**

When issuing the command with the `reverse` specifier, only the group's direction is changed on the system where the command is issued.

**-natural**

When issuing the command with the `reverse` specifier, only the natural direction of the groups is reversed, leaving the current direction unchanged.

**-current**

When issuing the `setrcopygroup` command with the `reverse` subcommand, only the current direction of the groups is reversed.

**-waittask**

Wait for all tasks created by this command to complete before returning. This option applies to the `failover`, `recover`, `restore`, and `reverse` subcommands.

## SPECIFIERS

`<policy>`

Specifies the policy to assign to the group. This specifier can only be used with the `pol` subcommand. Valid policies are:

`fail_wrt_on_err`

Specifies that if Remote Copy is started for the volume group and a write to the secondary system fails, then an I/O error is returned to the host.

`no_fail_wrt_on_err`

Specifies that if Remote Copy is started for the volume group and a write to the secondary system fails, then the Remote Copy operation is stopped and an I/O error is not returned to the host (default). This allows the application writing the data to continue, but makes the secondary volumes out of date with the primary volumes.

`auto_recover`

Specifies that if the Remote Copy is stopped as a result of the Remote Copy links going down, the group is restarted automatically after the links come back up. If this policy is enabled for a group while the group is stopped after link failures it will be only be started when the links come up for the failed target. If the links are already up at the time the policy is set then the group will not be restarted at that time.

`no_auto_recover`

Specifies that if the Remote Copy is stopped as a result of the Remote Copy links going down, the group must be restarted manually after the links come back up (default).

`over_per_alert`

If a synchronization of a periodic Remote Copy group takes longer to complete than its synchronization period then an alert will be generated. This is the default behavior.

`no_over_per_alert`

If a synchronization of a periodic Remote Copy group takes longer to complete than its synchronization period then an alert will not be generated.



**NOTE:** When issuing the `setrcopygroup <dr_operation>` command, either the `<group_name>` specifier or the `<target_name>` specifier must be specified.

<group\_name>

Specifies the name of the volume group whose policy is set, or whose target direction is switched.

<target\_name>

Specifies the target name for the target definition created with the `creatercopytarget` command.

<mode\_value>

Specifies the mode, `sync` or `periodic`, to which the group is set. This specifier can only be used with the `mode` subcommand.

<period\_value>`s|m|h|d`

Specifies the time period in units of seconds (`s`), minutes (`m`), hours (`h`), or days (`d`), for automatic resynchronization (e.g. `14h` for 14 hours). The time must be longer than or equal to five minutes and not more than one year in duration, or set to zero to indicate that no period should be used. This specifier can only be used with the `period` subcommand.

## RESTRICTIONS

- This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- When issuing the `setrcopygroup <dr_operation>` command, either the `<group_name>` specifier or the `<target_name>` specifier must be specified.
- Reversing the direction of primary volumes will result in the loss of any data changed after the group was stopped.
- The `reverse` specifier is not to be used as part of the normal disaster recovery process.

## EXAMPLES

The following example sets the group policy for `Group1`:

```
cli% setrcopygroup pol fail_wrt_on_err Group1
```

The following example reverses the current direction of secondary group `Group1` so that I/O might be applied to the group after disaster recovery:

```
cli% setrcopygroup failover Group1.r121
```



The following example sets volume group `Group1` to synchronize to its asynchronous periodic mode target `InServ2` every 30 minutes:

```
cli% setcopygroup period 30m InServ2 Group1
```

## NOTES

- There is no default resynchronization period. For groups whose mode is asynchronous periodic, you must specify a resynchronization period using `setcopygroup period <period_value>` or resynchronizations does not automatically take place.
- The minimum interval for periodic resynchronizations is five minutes.
- Use the `-nosnap` option when the primary server has failed or where the disks are ruined or in an unknown state. For example, an uncontrolled shutdown can result in loss of data. If you suspect that the primary volumes are not in a known good state, you should use this option to force a FULL RESYNC when the primary system is restored.
- The `-nosnap` option can be used when making a secondary group take over as the primary after a disaster takes down the primary (`setcopygroup failover -nosnap...`). This option indicates that no incremental resynchronization of the primary group is possible while the primary system is coming back online. Without this option, a snapshot is taken when the secondary server takes over as the primary. That snapshot is used to do an INCREMENTAL synchronization of the primary after it is restored. This assumes that there was no loss of data in the primary volumes when the primary server went down.

---

## COMMAND

setrcopytarget

## DESCRIPTION

The `setrcopytarget` command sets the name, policies, and throughput of a target definition.

## SYNTAX

The syntax for the `setrcopytarget` command can be one of the following:

- `setrcopytarget pol <policy> <target_name>`
- `setrcopytarget name <new_name> <target_name>`
- `setrcopytarget tput <tput_value> <target_name>`
- `setrcopytarget tunelinks <bandwidth> <latency> <target_name>`
- `setrcopytarget enable|disable <target_name>`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

`pol`

Sets the policy for the specified target using the `<policy>` specifier.

`name`

Changes the name of the indicated target using the `<new_name>` specifier.

`tput`

Sets the maximum throughput value for the target's links.

`tunelinks`

Adjust performance values for the target's links using the `<bandwidth>` and `<latency>` specifiers.

`enable|disable`

Enables or disables the target.

## OPTIONS

None.

## SPECIFIERS

`<target_name>`

Specifies the target name for the target definition previously created with the `creatercopytarget` command.

`<policy>`

This specifier can only be used with the `pol` subcommand. The policy can be one of the following:

`mirror_config|no_mirror_config`

Specifies that all configuration commands (`creatercopygroup`, `removercopygroup`, `admitrcopyvv`, `dismissrcopyvv`, `setrcopygroup pol/period`, `startrcopygroup`, and `stoprcopygroup`) involving the specified target are duplicated (`mirror_config`) or not duplicated (`no_mirror_config`). If not specified, all configuration commands are duplicated.

`<new_name>`

The new name for the indicated target. This specifier can only be used with the `name` subcommand.

`<tput_vlaue>`

Specifies the maximum throughput for this target's links, and is used to limit the total throughput of the link. You can optionally specify `g` or `G` (gigabytes), `m` or `M` (megabytes), or `k` or `K` (kilobytes) following the throughput value to indicate size (with no space between the specified value and size type). The default is kilobytes. This specifier can only be used with the `tput` subcommand. A value of 0 will remove the throughput.

`<bandwidth>`

The measured bandwidth of the connection to the target, specified in Kilobytes (KB) per second. This specifier can only be used with the `tunelinks` subcommand.

<latency>

The measured round-trip latency of the connection to the target, specified in milliseconds (ms). This specifier can only be used with the `tunelinks` subcommand.

## RESTRICTIONS

- Access to all domains is required to run this command.
- This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- If the `mirror_config` policy is set and the `setrcopytarget` command is issued with the `pol` subcommand, the duplicated configuration commands cannot be issued on the secondary. Doing so results in an error.
- There must be an active connection between the systems in the Remote Copy pair to issue commands on the primary to be mirrored to the secondary. If there is no connection, the commands return an error.

## EXAMPLES

The following example will set the throughput of each link of target `InServB` to 5 megabytes per second:

```
cli% setrcopytarget tput 5M InServB
```

## NOTES

- The `setrcopytarget` command requires the groups associated with it be stopped prior to using the following options:

- ◆ `setrcopytarget name <new_name> <target_name>`
- ◆ `setrcopytarget tput <throughput> <target_name>`

The `setrcopytarget` command with the following arguments can be run without bringing down its Remote Copy groups:

- ◆ `setrcopytarget pol <policy> <target_name>`
- ◆ `setrcopytarget tunelinks <bandwidth> <latency> <target_name>`

- Under normal operating conditions the `mirror_config` policy should never be changed to `no_mirror_config`. This policy option is included only as a method to correct several unusual error conditions that might occur in the course of operation which result in a mismatch in configuration between the two sides of a Remote Copy pair. For instance, it is possible for a group to be created, or a volume to be added to a group, only on one side of the pair if the operation is interrupted by a network failure. In such cases it might be necessary to temporarily change the policy to `no_mirror_config` in order to bring the configurations into alignment. After being corrected the `mirror_config` policy should be immediately restored. Such operations should generally only be undertaken on the advice of a 3PAR representative.

---

**COMMAND**

setsched

**DESCRIPTION**

The `setsched` command allows users to suspend, pause, change the schedule, change the parameters, and change the name of currently scheduled tasks.

**SYNTAX**

`setsched [options <arg>] <schedname>`

`setsched [-suspend | -resume] <schedname>`

`setsched [-suspend_all | -resume_all]`

**AUTHORITY**

Super, Service

**OPTIONS**

`-suspend <schedname>`

Suspends execution of a specified task. This option will not halt execution of a task that is already running. This option cannot be used with other options.

`-suspend_all`

Suspends execution of all scheduled tasks. This option will not halt execution of a task that is already running. This option cannot be used with other options.

`-resume`

Resumes scheduling of a specified task. This option cannot be used with other options.

`-resume_all`

Resume scheduling of all specified tasks. This option cannot be used with other options.

`-s <newschedule>`

Sets a new schedule for a given task. The `newschedule` is a cron-style schedule for the task to be run on. Each field can be up to 127 characters in length.

`-name <newname>`

Sets a new name for a given task. The name can be up to 31 characters in length.

`-no_alert`

Failures of tasks will no longer generate an alert.

`-alert`

Failures of tasks will generate an alert. This is set by default.

## **SPECIFIERS**

`<schedname>`

Specifies the name of the scheduled task to modify.

## **RESTRICTIONS**

None.

## **EXAMPLES**

None.

## **NOTES**

None.

---

**COMMAND**

setsnmppw

**DESCRIPTION**

The setsnmppw command allows a user to update SNMP access community strings passwords. The SNMP password is required for the system manager to send requests to the SNMP agent.

**SYNTAX**

setsnmppw [options] <password>

**AUTHORITY**

Super



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

-rw | -r | -w

Specifies that the read-write (-rw), read-only (-r), or write-only (-w) community password is changed. If not specified, the read/write password is changed.

**SPECIFIERS**

<password>

Specifies the new user-defined password using up to 50 alphanumeric characters.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example changes the read/write SNMP community string password to newpassword1:

```
cli% setsnmppw newpassword1
```



The following example changes the read-only SNMP password to `newpassword2` by specifying the `-r` option on the command line:

```
cli% setsnmppw -r newpassword2
```

The following example changes the write-only SNMP password to `newpassword3` by specifying the `-w` option on the command line:

```
cli% setsnmppw -w newpassword3
```

## NOTES

- The default initial read/write password is `public`.
- If the read-only or write-only passwords do not exist, they are created.
- Verify SNMP passwords by issuing the `showsnmppw` command. See [showsnmppw](#) on page 22.159 for more information.

---

**COMMAND**

setsshkey

**DESCRIPTION**

The `setsshkey` command sets the SSH public key for a user.

The user will be prompted to provide the SSH public key. To finish entering the public key, press enter on a blank line. The key must have been generated using the `ssh-keygen` utility. The public key is contained in the user-defined file named with `.pub` extension. The user can open this file with an ASCII editor to copy the key and paste it. After setting the SSH public key on the InServ Storage Server, the user can use the corresponding private key to log on without a password. This new key replaces the existing key if any.

**SYNTAX**

setsshkey [options]

**AUTHORITY**

Super, Edit, Browse, Service

**OPTIONS**

-add

Specifies that the given key should be added to the list of authorized keys instead of replacing existing ones.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays the setting of a user's SSH public key:

```
cli% setsshkey
```

```
Please enter the SSH public key below. When finished, press enter
twice. The key is usually long. It's better to copy it from inside an
editor and paste it here. (Please make sure there are no extra blanks.)
The maximum length of the key is 4096 bits.
```

```
ssh-rsa AAAAB3NzaClyc2EAAAABIwAAAQEA+4Z3zT8Sq6t3s08q/MMd2ZnMpRc/3tyHMk63dH
R8b/VgV6ewXNfYhhfTGWRdZldtyLFXnuEPuf+z8EtinFStinTzA3FiJ0agK7rLoNtn/F0jBaGwM
SWukqzAQA2VJvq/keaLVMT3+J3nvXEucjS4fApeeLwFgKczOX51loaHDtn2ys2C5l+mLwlVDUJL
wIKJlJOWqr68ToeRgTDMoppOI3cG14ryF0re4xKANHFQLnSct5ANjqD2jpnYEABuOvBX7G6vI5g
zQFfcAda/a2bq563/AHr8Ehhi5EVy+GSCqjj8cr0/zHSZyX6llQUfe0YxmayPEKlods6bDi5oxG
COw== user1@server1"
```

## NOTES

- Issuing the `setsshkey` command results in the system prompting for your SSH public key. Copy and paste the key using the aforementioned ASCII editor on the command line and then press ENTER.
- After setting the SSH public key on the system, use the corresponding private key to log on without a password. A single ssh key may be used for multiple hosts with a single use of `setsshkey`, or different keys may be used for different hosts. To restrict the usage of a key to a particular host, a `-from` option can be added to the beginning of a key. For example:  
`from="192.168.1.1" ssh-rsa AAAAB3Nza...`
- The maximum number of characters used to represent the SSH key (including the `-from` option, key type, and additional comments) is 4095.
- LDAP users are only allowed to set an SSH key if the `setauthparam` command has been used to set the `allow-ssh-key` parameter to 1. When an LDAP user runs the `setsshkey` command, the user's privilege level is recorded and is assigned when the user logs in using the key. Changes in the group-to-privilege mappings set with the `setauthparam` command or changes in the user's data in the LDAP server have no effect as long as the user has an SSH key.
- Removing the user's SSH key forces a new privilege to be determined at the user's next login.
- Only one key may be entered at a time; to enter multiple keys, run the `setsshkey` command again with the `-add` option.

---

**COMMAND**

setstatch

**DESCRIPTION**

The `setstatch` command starts and stops the statistics collection mode for chunklets.

**SYNTAX**

`setstatch start|stop <LD_name> <chunklet_num>`

**AUTHORITY**

Super, Edit

**SUBCOMMANDS**

`start|stop`

Specifies that the collection of statistics is either started or stopped for the specified Logical Disk (LD) and chunklet.

**OPTIONS**

None.

**SPECIFIERS**

`<LD_name>`

Specifies the name of the LD in which the chunklet to be configured resides.

`<chunklet_num>`

Specifies the chunklet that is configured using the `setstatch` command.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the start of statistics collection on chunklet 0 of LD test:

```
cli% setstatch start test 0
```

**NOTES**

After the statistic collection mode for the chunklet is set, you can then use either the `histch` command ([page 15.2](#)) or the `statch` command ([page 25.2](#)) to view the chunklet's statistics.

---

**COMMAND**

setstatpdch

**DESCRIPTION**

The `setstatpdch` command sets the statistics collection mode for all in-use chunklets on a Physical Disk (PD).

**SYNTAX**

`setstatpdch start|stop <PD_ID>`

**AUTHORITY**

Super, Edit

**SUBCOMMANDS**

`start|stop`

Specifies that the collection of statistics is either started or stopped for chunklets on the specified PD used by Logical Disks (LDs).

**OPTIONS**

None.

**SPECIFIERS**

`<PD_ID>`

Specifies the PD ID.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the start of statistics collection on all PD chunklets of PD 0:

```
cli% setstatpdch start 0
```

**NOTES**

After the statistic collection mode for the chunklet is set, you can then use either the `histch` command ([page 15.2](#)) or the `statch` command ([page 25.2](#)) to view the chunklet's statistics.

---

## COMMAND

setsys

## DESCRIPTION

The `setsys` command sets the properties of the system, and includes options to annotate a system with descriptor information such as physical location, owner, contact information, and so on. The command also enables you to set system-wide parameters such as the raw space alert.

## SYNTAX

```
setsys [options]  
setsys <parameter>
```

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

The following option is designed for changing the name of the system:

`-name <systemname>`

Specifies the new name of the system up to 31 characters in length.

The following options allow the annotation of the system with descriptor information:

`-loc <location>`

Specifies the location of the system.

`-owner <owner>`

Specifies the owner of the system.

`-contact <contact>`

Specifies the contact information for the system.

`-comment <comment>`

Specifies any additional information for the system.

## SPECIFIERS

The following parameters can be configured on the system and are issued for the <parameter> specifier:

`RawSpaceAlertFC <value>`

Sets the user configurable space alert threshold (100 to 100000 GB) for Fibre Channel type drives. When the total space on the available chunklets (both clean and unclean) for the specified drive type falls below the specified value, the alert is posted. A value of 0 will disable the alert.

`RawSpaceAlertNL <value>`

Performs the same function as `RawSpaceAlertFC`, but should be used for Nearline type drives.

`RawSpaceAlertSSD <value>`

Performs the same function as `RawSpaceAlertFC`, but should be used for Solid State Drive type drives.

`RemoteSyslog <value>`

Enables or disables sending events as `syslog` messages to a remote system. A value of 0 disables the message, and a value of 1 enables the message. `syslog` messages are sent with a facility user and with event severities mapped to `syslog` levels such as:

**Table 21-5.**

Event Severity	syslog Level
fatal	alert
critical	alert
major	crit
minor	err
degraded	warning
info	info



RemoteSyslogHost <value>

Sets the IP address of the system to which events will be sent as syslog messages. The value must be a valid IP address.

SparingAlgorithm <value>

Sets the sparing algorithm used by the `admithw` command. Valid values are `Default`, `Minimal`, `Maximal`, and `Custom`.

MgmtOldPorts <value>

Enables or disables listening on the old management ports 2540 and 2550 (for SSL). The server now listens on ports 5782 and 5783 (for SSL). By default, the server also listens on the old ports unless they are disabled using this parameter. The value must be either `enable` or `disable`. Default ports 2540 and 2550 were deprecated in the 2.2.4 release and support for them will be removed in a future release.



**NOTE:** The server must be restarted for changes made with the `MgmtOldPorts` parameter to take effect.

RemoteCopyOldPort <value>

Enables or disables listening on the old TCP port 3491. The value must be either `enable` or `disable`. By default, remote copy links listen on both IANA registered port 5785 AND port 3491.



**WARNING:** Changing this parameter will reset all remote copy links. This may cause remote copy targets to lose their connections temporarily.

CopySpaceReclaim <value>

Enable the system to reclaim unused snapshot copy space if the <value> argument is set to 1. The default is 1 for a newly installed system with InForm 2.3.1 or later, and 0 if the system had been upgraded from any version prior to 2.3.1. Once the <value> argument is set to 1, setting it back to 0 is not allowed. Downgrading to a software release prior to 2.3.1 if <value> is 1 will fail if there is any VVs in the system with reclaimed snapshot copy space. If the downgrade fails, refer to the CLI Administrator's Manual for instructions to enable the downgrade.

VVRetentionTimeMax <value>[h|H|d|D]

Specifies the maximum value that can be set for the retention time of a volume. <time> is a positive integer value and in the range of 0 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either d or D for day and h or H for hours following the entered time value. The default value for <time> is 14 days.

To disable setting the volume retention time in the system, enter 0 for <time>.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays setting a raw space alert of 800 gigabytes:

```
cli% setsys RawSpaceAlertFC 800
cli% showsys -param
System parameters from configured settings

----Parameter-----      ---Value---
RawSpaceAlertFC           :              800
RawSpaceAlertNL           :                0
RemoteSyslog              :                1
RemoteSyslogHost          :      192.168.6.15
SparingAlgorithm          :           Minimal
CopySpaceReclaim          :                0
EventLogSize              :               3M
VVRetentionTimeMax        :      336 Hours
```

## NOTES

Use the `showsys -param` command to see the current raw space alert setting (see [showsys](#) on page 22.171).

---

## COMMAND

setsysmgr

## DESCRIPTION

The `setsysmgr` command sets the system manager startup state.



**CAUTION:** Issuing the `setsysmgr` command can potentially remove the entire state of the system causing data loss.

## SYNTAX

The syntax of the `setsysmgr` command can be one of the following:

- `setsysmgr wipe [-f] <system_name> [<system_ID>]`
- `setsysmgr tocgen [-f] [<toc_gen_number> [<disk_quorum>]]`
- `setsysmgr force_iderecovery [-f]`
- `setsysmgr force_idewipe [-f]`

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

`wipe`

Requests that the specified system be started in the new system state. If this subcommand is not used, then the `tocgen` subcommand must be used.

`tocgen`

Specifies that the system is to be started with the specified table of contents generation number. If this subcommand is not used, then the `wipe` subcommand must be used.

`force_iderecovery`

Specifies that the system starts the recovery process from the IDE disk even if all Virtual Volumes (VV) have not been started.



**CAUTION:** Issuing the `setsysmgr force_idewipe` command can result in data loss.

`force_idewipe`

Specifies that the system wipes the IDE powerfail partition. The system is shutdown and restarted, during which time all logical disks and VVs are checked.

## OPTIONS

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

`<system_name> [ <system_ID> ]`

Specifies the name of the system to be started, using up to 31 characters in length.

`<system_ID>`

Specifies the ID of the system to be started.

`<toc_gen_number>`

Specifies the table of contents generation number for the system to start with.

`<disk_quorum>`

Specifies the disk quorum for the system to start with.

## RESTRICTIONS

- Access to all domains is required to run this command.
- Use this command only when the system cannot start up normally.

## EXAMPLES

The following example starts system `mysystem` in the new system state:

```
cli% setsysmgr wipe mysystem systemid
```

The following example starts the system with the table of contents generation number of 42956, where 91 is the highest disk quorum displayed:

```
cli% setsysmgr tocgen 42956 91
```

The following example displays the start of a system's recovery process from its IDE disk:

```
cli% setsysmgr force_iderecovery
```

## NOTES

- If the `wipe` subcommand is specified, all system data and configuration information, including customer data and VV layout, are destroyed.
- If the `force_iderecovery` subcommand is specified, the system can delete data for some of the unstarted VVs. The system can run the `checkvv` and `checkld` commands to make the VVs and logical disks consistent, thereby resulting in a possible data loss.

---

## COMMAND

settemplate

## DESCRIPTION

The `settemplate` command modifies the properties of existing templates.

## SYNTAX

```
settemplate <option_value>... [option <arg>] <template_name>
```

## AUTHORITY

Super



**NOTE:** You need access to all domains in order to run this command.

## OPTION

`-remove <option>...`

Indicates that the option(s) that follow `-remove` are removed from the existing template. When specifying an option for removal, do not specify the option's value. For valid options, refer to [createtemplate](#) on page 11.61.

## SPECIFIERS

`<option_value>...`

Indicates the specified options and their values (if any) are added to an existing template. The specified option replaces the existing option in the template. For valid options, refer to [createtemplate](#) on page 11.61.

`<template_name>`

Specifies the name of the template to be modified, using up to 31 characters.

## RESTRICTIONS

- Access to all domains is required to run this command.
- The `-desc` option cannot be removed from a template (it can be changed to an empty string).
- The `-nrw` and `-nro` options cannot be removed from a template. These options can only be replaced by specifying either `-nrw` or `-nro` before the `-remove` option.

## EXAMPLES

- The following example displays template `vvtemp1` modified to support the availability of data should a drive magazine fail (`mag`) using the `stale_ss` policy:

```
cli% settemplate -ha mag -pol stale_ss vvtemp1
```

- In the following example, the `-nrw` and `-ha mag` options are added to the template `templatel`, and the `-t` option is removed:

```
cli% settemplate -nrw -ha mag -remove -t templatel
```

## NOTES

- All options available for the `createald`, `createaldvv`, `createcpg`, and `createtpvv` commands can be used with the `settemplate` command to modify an existing template.
- Options preceded with `-remove` are removed from an existing template.
- When adding options to a template, specify the option flag and its value (if any).
- When removing options from a template, specify the option flag only.

---

**COMMAND**

setuser

**DESCRIPTION**

The `setuser` command sets your user properties.

**SYNTAX**

setuser [options <arg>] <user>

**AUTHORITY**

Super, Edit, Browse

**OPTIONS**

-f

Specifies that the command is forced. No confirmation is requested before executing the command.

-adddomain <domain>:<privilege>[,<domain>:<privilege>...]

Adds a specified user (<user>) to the specified domain (<domain>) at the specified privilege level (<privilege>). Permitted values for <priv> are Super, Edit, Browse and Service.

-rmdomain <domain\_name\_or\_pattern>[,<domain\_name\_or\_pattern>...]

Remove the user from each domain with a name that matches one or more of the <domain\_name\_or\_pattern> options.

-defaultdomain <domain>

Changes the default domain of the user to a specified domain. You must already have permission set in the domain. Specify the `-unset` option to specify no default domain. The Browse and Edit authorities for this command and option can be used by any user with their own user name.

**SPECIFIERS**

<user>

Specifies the name of the user.



## RESTRICTIONS

The `-adddomain` and `-rmdomain` options cannot be used on a user that is currently logged in. Existing sessions can be terminated with the `removeuserconn` command.

## EXAMPLES

In the following example, user `3paruser` is permitted `edit` level privileges in the domain `Engineering`:

```
cli% setuser -adddomain Engineering:edit 3paruser
```

## NOTES

None.

---

**COMMAND**

setuseracl

**DESCRIPTION**

The `setuseracl` command sets the Access Control List (ACL).

**SYNTAX**

`setuseracl [options] <user_name> <operation> [<name_or_pattern>]...`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-add`

Adds names or patterns of objects at the end of an existing ACL.

`-remove`

Removes names or patterns of objects from an ACL.

**SPECIFIERS**

`<user_name>`

Specifies the name of the user whose ACL is being set.

`<operation>`

Specifies the operation for which the ACL is being defined. The only value currently accepted is the `updatevv` command, which updates a snapshot with a new snapshot.

Refer to [updatevv](#) on page 29.4 for details.

`[<name_or_pattern>]...`

Specifies a list of names or patterns of objects on which the operation (as specified by the `<operation>` argument) is performed. The object type is dependent on the specified operation. For example, the objects or names specified for the `updatevv` operation are Virtual Volume (VV) names. If an empty string is specified without any option, then any existing ACL for the `<user_name>`, `<operation>` is removed.

**RESTRICTIONS**

None.

## EXAMPLES

The following example sets the ACL for user `testuser1` and allows `testuser1` to update the snapshot VVs of VVs `vv1` and `vv2`.

```
cli% setuseracl testuser1 updatevv vv1 vv2
```

The following example shows how to add `vv3` and `vv4` to the ACL for user `testuser1`:

```
cli% setuseracl -add testuser1 updatevv vv3 vv4
```

The following example shows how to remove `vv3` from the ACL for user `testuser1`:

```
cli% setuseracl -remove testuser1 updatevv vv3
```

The following example shows how to clear the ACL for user `testuser1`:

```
cli% setuseracl testuser1 updatevv ""
```

## NOTES

- The `-add` and `-remove` options are mutually exclusive.
- If the names or pattern of objects to add already exist, the request is ignored.
- If the names or pattern of objects to remove do not exist, the request is also ignored.

---

**COMMAND**

setvv

**DESCRIPTION**

The `setvv` command changes the properties associated with Virtual Volumes (VVs). It can be used to modify volume names, volume policies, allocation warning, and limit levels, and the volume's controlling Common Provisioning Group (CPG).

**SYNTAX**

```
setvv [options <arg>] <VV_name|pattern>...
```

**AUTHORITY**

Super, Edit

**OPTIONS**

At least one of the following options must be specified:

`-name <new_name>`

Specifies that the name of the VV be changed to a new name (as indicated by the `<new_name>` specifier) that uses up to 31 characters.

`-clrrsv`

Specifies that all reservation keys (i.e. registrations) and all persistent reservations on the VV are cleared.

`-exp <time>[d|D|h|H]`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 0 - 43,800 hours (1825 days). Time can be specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

To remove the expiration time for the volume, enter 0 for `<time>`.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.  
Use `-comment ""` to remove the comments.

`-f`

Do not ask for confirmation before setting or modifying volumes with retention time (`-retain`).

`-retain <time>[d|D|h|H]`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.



**NOTE:** If the volume is not in any domain, then its retention time cannot exceed the value of the system's `VVRetentionTimeMax`. The default value for the system's `VVRetentionTimeMax` is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's `VVRetentionTimeMax`, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the 3PAR Virtual Lock License. Contact your local 3PAR Authorized Service Provider for more information.

`-pol <policy>[, <pol>...]`

Specifies virtual volume policies. If an argument is not specified, the option defaults to `stale_ss`. Valid policies are as follows:

`stale_ss`

Specifies that invalid (stale) snapshot volumes are permitted. Failures to update snapshot data does not affect the write to the base volume. However, the snapshot is considered invalid.

`no_stale_ss`

Specifies that stale snapshot volumes are not permitted. A failure to update a snapshot is considered a failure to write to the base volume.

`one_host`

This constrains the export of a volume to one host or one host cluster (when cluster names may be used as a host name).

`no_one_host`

This policy should only be used when exporting a virtual volume to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.

### `tp_bzero`

Specifies that if a host write results in the allocation of a new data page that is only partially filled by the host write, then a zero-fill is performed on the unwritten portion of the data page to ensure that the host cannot read data from deleted volumes or snapshot. The default allocation page size is 16 KB. This is the default setting.

### `no_tp_bzero`

Specifies that the zero-fill operation is bypassed on the allocation of partially written data pages.

### `zero_detect`

This policy enables the InServ to scan for zeros in the incoming write data. This feature when used during physical copy to a TPVV will avoid allocating space for blocks containing zero. This feature when used with a Thin Persistence license will reclaim allocated space when zero blocks are written to the TPVV. This policy is only applicable for the base TPVV.



**NOTE:** There can be some performance implication under extreme busy systems so it is recommended for this policy to be turned on only during Fat to Thin and re-thinning process and be turned off during normal operation.

### `no_zero_detect`

This policy disables the InServ to scan for zeros in the incoming write data to reclaim allocated space on the volume. This is the default policy setting.

Multiple policies can be specified and are separated with commas. If a policy is not specified, the policy defaults to `stale_ss`.

The following options can only be used on Thinly Provisioned Virtual Volumes (TPVVs):

`-usr_aw <percent>`

This option returns a user space allocation warning. It generates a warning alert when the user space of the TPVV exceeds the specified percentage of the VV size.

`-usr_al <percent>`

This option returns the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the virtual volume size. After this size is reached, any new writes to the VV will fail.

The following option can only be used on fully provisioned volumes:

`-usr_cpg <usr_CPG>`

Specifies that the volume user space that is to be provisioned from the specified CPG. This option moves all the logical disks currently contained in the user space for these volumes into the CPG. This is permitted only when none of the logical disks are shared with other volumes that are not specified with this option. If the `<usr_CPG>` specifier is specified as `" "`, the volume user space is no longer provisioned from a CPG and the existing user space logical disks will be removed from the CPG. This is permitted only when the existing user's logical disks are exclusively being used by the volumes specified with this option.

The following options can only be used on thinly provisioned volumes:

`-snp_cpg <snp_cpg>`

Specifies that the volume snapshot space is to be provisioned from the specified CPG. This option moves all the logical disks currently contained in the snapshot space for these volumes into the CPG. This is permitted only when none of the logical disks are shared with other volumes that are not specified in this option. If the `<snp_CPG>` specifier is specified as `" "`, the volume snapshot space is longer provisioned from a CPG and the existing snapshot space logical disks are removed from the CPG. This is permitted only when the existing snapshot logical disks are exclusively used by the volumes specified in this option.

`-snp_aw <percent>`

Indicates a snapshot space allocation warning. Through this option you can generate a warning alert when the snapshot space of the VV exceeds the indicated percentage of the VV size.

`-snp_al <percent>`

Sets a snapshot space allocation limit. The snapshot space of the VV is prevented from growing beyond the indicated percentage of the VV size.

`-spt <sectors_per_track>`

Defines the VV geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is from 4 to 8192 and the default value is 304.

`-hpc <heads_per_cylinder>`

Defines the VV geometry heads per cylinder value that is reported to the hosts through the SCSI mode pages. The valid range is from 1 to 1024 and the default value is 8.

## SPECIFIERS

`<VV_name|pattern>...`

Specifies the VV name or all VVs that match the pattern specified, using up to 31 characters. The patterns are glob-style patterns (see Help on sub, or globpat). Valid characters include alphanumeric characters, periods, dashes, and underscores.

## RESTRICTIONS

At least one option must be specified.

## EXAMPLES

The following example displays space usage information for thin provisioned VVs that are exported to host `hname`.

```
cli% showvv -s -p -prov tp* -host hname
```



**NOTE:** The `tp*` matches `tpvv` as well as `tpsd`.

The following example displays raw space usage information for thin provisioned VVs exported to host `hname` and all VVs in the tree with the base VV of Id 50:

```
cli% showvv -r -p -prov tp* -host hname -p -baseid 50
```

The following example lists the Id, Name and VSize\_MB columns only for thin provisioned VVs exported to host `hname`:

```
cli% showvv -showcols Id,Name,VSize_MB -p -prov tp* -host hname
```



**NOTES**

- To view policies assigned to the system's VVs, issue the `showvv -p` command. See [showvv](#) on page 22.201 for more information.
- It is not possible to rename a VV that is already associated with a Remote Copy group.
- Changing the CPG for a TPVV is not allowed.

---

**COMMAND**

setvvset

**DESCRIPTION**

The setvvset command sets the parameters and modifies the properties of a Virtual Volume (VV) set.

**SYNTAX**

```
setvvset [options <arg>] <setname>
```

**AUTHORITY**

Super, Edit

**OPTIONS**

-comment <comment>

Specifies any comment or additional information for the set. The comment can be up to 255 characters in length. Unprintable characters are not allowed.

-name <newname>

Specifies a new name for the VV set.

**SPECIFIERS**

<setname>

Specifies the name of the vv set to modify.

**RESTRICTIONS**

None.

**EXAMPLES**

To rename a set from foo to bar:

```
cli% setvvset -name bar foo
```

To change the comment on a set bar:

```
cli% setvvset -comment "This used to be set foo" bar
```

## NOTES

None.



# 22

## Show Commands

---

### In this chapter

<code>showalert</code>	<a href="#">22.4</a>
<code>showauthparam</code>	<a href="#">22.6</a>
<code>showbattery</code>	<a href="#">22.8</a>
<code>showblock</code>	<a href="#">22.14</a>
<code>showcage</code>	<a href="#">22.17</a>
<code>showcim</code>	<a href="#">22.23</a>
<code>showclienv</code>	<a href="#">22.25</a>
<code>showcpg</code>	<a href="#">22.26</a>
<code>showdate</code>	<a href="#">22.33</a>
<code>showdomain</code>	<a href="#">22.34</a>
<code>showdomainset</code>	<a href="#">22.36</a>
<code>showeeprom</code>	<a href="#">22.38</a>
<code>showeventlog</code>	<a href="#">22.41</a>
<code>showfirmwaredb</code>	<a href="#">22.45</a>
<code>showhost</code>	<a href="#">22.47</a>

showhostset	22.52
showinventory	22.54
showiscsisession	22.56
showld	22.58
showldch	22.66
showldmap	22.71
showlicense	22.73
shownet	22.75
shownode	22.77
shownodeenv	22.87
showpatch	22.89
showpd	22.91
showpdata	22.108
showpdch	22.109
showpdvv	22.116
showport	22.121
showportarp	22.132
showportdev	22.134
showportisns	22.137
showportlesb	22.139
showrcopy	22.145
showrctransport	22.150
showrsv	22.153
showsched	22.155
showsnmppw	22.159
showsnmpmgr	22.157

showspace	<b>22.161</b>
showspare [-used]	<b>22.166</b>
showsshkey	<b>22.169</b>
showsys	<b>22.171</b>
showsysmgr	<b>22.176</b>
showtarget	<b>22.179</b>
showtask	<b>22.180</b>
showtemplate	<b>22.184</b>
showtoc	<b>22.185</b>
showtocgen	<b>22.187</b>
showuser	<b>22.188</b>
showuseracl	<b>22.190</b>
showuserconn	<b>22.191</b>
showversion	<b>22.193</b>
showvlun	<b>22.195</b>
showvv	<b>22.201</b>
showvvmap	<b>22.217</b>
showvvpd	<b>22.219</b>
showvvset	<b>22.224</b>

---

**COMMAND**

`showalert`

**DESCRIPTION**

The `showalert` command displays the status of system alerts. When issued without options, all alerts are displayed.

**SYNTAX**

`showalert [options]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

The selection of alerts to be displayed is controlled by selecting one of the following options:

`-n`

Specifies that only new alerts are displayed.

`-a`

Specifies that only acknowledged alerts are displayed.

`-f`

Specifies that only fixed alerts are displayed.

`-all`

Specifies that all alerts are displayed.

The format of the alert display is controlled by the following options:

`-d`

Specifies that detailed information is displayed. Cannot be specified with the `-online` option.

`-online`

Specifies that summary information is displayed in a tabular form with one line per alert. The message text will be truncated if it is too long unless the `-wide` option is also specified.



`-wide`

Do not truncate the message text. Only valid if the `-online` option is also specified.

## SPECIFIERS

None.

## RESTRICTIONS

Without any options, the `showalert` command displays all alerts in the `New` state.

## EXAMPLES

The following example displays new alerts on a system:

```
cli% showalert -n

Id           : 1
State        : New
Message Code: 0x2200de
Time         : 2008-07-17 20:14:29 PDT
Severity     : Degraded
Type         : Component state change
Message      : Node 0, Power Supply 1, Battery 0 Degraded (Unknown)

Id           : 2
State        : New
Message Code: 0xe000c
Time         : 2008-07-17 20:15:31 PDT
Severity     : Informational
Type         : Cluster shutdown after system recovery completion
Message      : System recovery completed for node 0 (reason Power Loss)

2 alerts
```

## NOTES

- Set the status of alerts by issuing the `setalert` command. See [setalert](#) on page 21.3 for more information.
- Alerts can be removed by issuing the `removealert` command. See [removealert](#) on page 19.3 for more information.

---

**COMMAND**

showauthparam

**DESCRIPTION**

The showauthparam command shows authentication parameters.

**SYNTAX**

showauthparam

**AUTHORITY**

Super

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLE**

The following is example output from the showauthparam command:

```
cli% showauthparam
Param                -----Value-----
ldap-server          xxx.xxx.xx.xx
ldap-server-hn       domaincontroller.work.com
binding              sasl
sasl-mechanism        GSSAPI
kerberos-realm       NTDOM1.work.COM
accounts-dn          OU=Users,DC=work,DC=com
account-obj          user
account-name-attr     sAMAccountName
memberof-attr        memberOf
edit-map              CN=Software,CN=Users,DC=work,DC=com
browse-map            CN=Eng,CN=Users,DC=work,DC=com
domain-name-attr     description
group-obj             group
domain-name-prefix   !InServDomain=
```

## NOTES

None.

---

**COMMAND**

`showbattery`

**DESCRIPTION**

The `showbattery` command displays battery status information including serial number, expiration date, battery life, and so on, which could be helpful when determining battery maintenance schedules.

**SYNTAX**

`showbattery [options]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d`

Specifies that detailed battery information, including battery test information, serial numbers, and expiration dates, is displayed.

`-log`

Specifies the battery test log information.

`-i`

Specifies the battery inventory information.

`-s`

Specifies the detailed states of the battery.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays battery status information:

```
cli% showbattery
```

Node	PS	Bat	Serial	--State---	ChrgLvl(%)	-ExpDate--	Expired	Testing
0	0	0	FFFFFFFF	OK	100	05/15/2007	No	No
0	1	0	FFFFFFFF	OK	100	05/15/2007	No	No
1	0	0	--	NotPresent	--	--	No	No
1	1	0	--	NotPresent	--	--	No	No

The columns in the previous example are identified as follows:

- Node. The controller node number.
- PS. The power supply number.
- Bat. The battery ID.
- Serial. The battery serial number.
- State. The current status of the battery. Battery states can be as follows:
  - ◆ --. Cannot determine the battery state.
  - ◆ NotPresent. Battery is missing.
  - ◆ OK. Battery is operating normally.
  - ◆ Failed. Battery is operating abnormally.
  - ◆ MaxLifeLow. Maximum battery life is low (less than 12 minutes).
- ChgLvl. Percentage of battery charge status for the node.
- ExpDate. The expiration date of the battery.
- Expired. Indicates whether expired batteries are connected to the power supply.
- Testing. Indicates whether a battery test is in progress. When a battery test is in progress, the batteries cannot be counted so the number of batteries shown is a cached value.

The following example displays detailed battery information:

```
cli% showbattery -d
-----Node 0 PS 0 Battery 0-----
Node ID           :    0
Power Supply ID   :    0
Battery ID        :    0
Manufacturer      :   MAG
Model             :  0800-0016-50.0B
Serial Number     :  70315366
State             :   OK
Charge State      :  FullyCharged
Charge Level(%)   :   100
Max Battery Life(mins) :  25
Expired           :   No
Test in Progress  :   No
Expiration Date   :  2010-07-17 13:00:00 PD
...
```

The following example displays the inventory information for the battery using the `showbattery -i` command:

```
cli% showbattery -i
Node PS Bat -Manufacturer- -Model- -Serial-
  0  0  0 --          --      FFFFFFFF
  0  1  0 --          --      FFFFFFFF
  1  0  0 --          --      --
  1  1  0 --          --      --
```

- Node. The controller node number.
- PS. The power supply number.
- Manufacturer. The manufacturer of the battery.
- Model. The battery model.
- Serial. Indicates the serial number of the battery.

The following example displays the battery test log information using the `showbattery -log` command:

```
cli% showbattery -log
Node PS Bat Test Result Dur(mins) -----Time-----
  0  0  0    0 Passed      1 Fri Jan 19 13:16:51 PST 2007
  0  0  0    1 Passed      1 Fri Feb 02 13:18:51 PST 2007
  0  0  0    2 Passed      1 Tue Feb 20 12:41:06 PST 2007
  0  0  0    3 Passed      1 Tue Mar 06 12:42:07 PST 2007
  0  0  0    4 Passed      1 Tue Mar 20 13:43:58 PDT 2007
  0  0  0    5 Passed      1 Tue Apr 03 13:45:35 PDT 2007
  0  0  0    6 Passed      1 Tue Apr 17 13:47:07 PDT 2007
  0  0  0    7 Passed      1 Tue May 01 13:49:05 PDT 2007
  0  0  0    8 Passed      1 Tue May 15 13:50:10 PDT 2007
  0  0  0    9 Passed      1 Tue May 29 13:51:36 PDT 2007
  1  1  0    0 Passed      1 Fri Jan 19 14:17:44 PST 2007
  1  1  0    1 Passed      1 Fri Feb 02 14:19:45 PST 2007
  1  1  0    2 Passed      1 Tue Feb 20 13:42:02 PST 2007
  1  1  0    3 Passed      1 Tue Mar 06 13:44:02 PST 2007
  1  1  0    4 Passed      1 Tue Mar 20 14:45:53 PDT 2007
  1  1  0    5 Passed      1 Tue Apr 03 14:47:24 PDT 2007
  1  1  0    6 Passed      1 Tue Apr 17 14:48:56 PDT 2007
  1  1  0    7 Passed      1 Tue May 01 14:50:54 PDT 2007
  1  1  0    8 Passed      1 Tue May 15 14:51:59 PDT 2007
  1  1  0    9 Passed      1 Tue May 29 14:53:25 PDT 2007
```

- Node. The controller node number.
- PS. The power supply number.
- Bat. The battery ID.
- Test. The current status of the battery.
- Result. The battery status, shows as passed or failed.
- Dur(mins). Indicates the length of time for the results to return.
- Time. Indicates the time of the last battery test.

The following example displays the detailed status of the battery using the `showbattery -s` command:

```
cli% showbattery -s
Node PS Bat -State- -Detailed_State-
  0  0  0 OK      Normal
  0  1  0 OK      Normal
  1  0  0 OK      Normal
  1  1  0 OK      Normal
```

- State
- Node. The controller node number.
- PS. The power supply number.
- Bat. The battery ID.
- State. The current status of the battery and can be one of the following:
  - ◆ OK. The battery is operating normally.
  - ◆ Failed. The battery is operating abnormally.
  - ◆ Degraded. The battery is in degraded state.
  - ◆ Not Present. The battery is missing.
- Detailed\_State. The detailed state of the battery and can be one of the following:
  - ◆ Not Present. The battery is missing.
  - ◆ Expired. The battery is already expired.
  - ◆ Failed. The battery is operating abnormally.
  - ◆ InvalidFirmware. The battery has invalid firmware.
  - ◆ Unknown. The battery state is unknown.
  - ◆ Undefined. The battery state can not be determined.

## NOTES

- Battery information is set by issuing the `setbattery` command. See [setbattery](#) on page 21.12 for more information.



- For the E-Class and F-Class Storage Servers, the Power Supply 0's primary node is 0 and its secondary node is 1. The Power Supply 1's primary node is 1 and its secondary node is 0. The primary node can read and modify the power supply's properties. The secondary node can only read the power supply's properties. Therefore, for the E-Class and F-Class, with the `-old` option, the battery information for Node 0 Power Supply 0 is the same as the battery information for Node 1 Power Supply 0. The battery information for Node 1 Power Supply 1 is the same as the battery information for Node 0 Power Supply 1.

---

**COMMAND**

showblock

**DESCRIPTION**

The `showblock` command displays block mapping information for Virtual Volumes (VVs), Logical Disks (LDs), and Physical Disks (PDs).

**SYNTAX**

`showblock [options <arg>] <dev> <block> [<eblock>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d`

Specifies that detailed information is displayed for the specified device and block.

**SPECIFIERS**

`<dev>`

The `<dev>` specifier is specified as one of the following arguments:

`vv <VV_name> usr|snp|adm`

Specifies the VV name and the area (`usr`, `snp`, or `adm`) of that volume for the block mapping information to be displayed.

`ld <LD_name>`

Specifies the LD name.

`pd <PD_ID>`

Specifies the ID of the PD.

`<block>`

Specifies the 512 byte block number on the specified device.

`[<eblock>]`

Specifies an end range when used with the `<block>` specifier. Additional mapping for blocks at the start of each device mapping boundary for the range indicated by `<block>` and `<eblock>` is displayed. This specifier is optional.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays block mapping information for block 0 in the user space of VV admin:

```
cli% showblock vv admin usr 0
      VVname Spc      VVBlock      LDname      LDBlock PDId Chnk      PDblock
      admin usr 0x00000000      admin.usr.0 0x00000000      88      0 0x00080000
      =      =      =      =      0x00000000      58      0 0x00080000
```

The following example displays detailed block mapping information for block 0 in the user space of VV admin:

```
cli% showblock -d vv admin usr 0
VV/LD      Name      Block      Region
VV usr:      admin 0x00000000 0x00000000-0x0007FFFF (0-256MB)
LD      :      admin.usr.0 0x00000000 0x00000000-0x0007FFFF (0-256MB)

PD Chnk      PDStep      Block      LDStep
88      0 0x00080000-0x000801FF 0x00080000 0x00000000-0x000001FF
58      0 0x00080000-0x000801FF 0x00080000 0x00000000-0x000001FF
```

The following example displays block mapping information from blocks 0x100 through 0x400 in the user space of VV admin:

```
cli% showblock vv admin usr 0x100 0x400
      VVname Spc      VVBlock      LDname      LDBlock PDId Chnk      PDblock
      admin usr 0x00000100      admin.usr.0 0x00000100      88      0 0x00080100
      =      =      =      =      0x00000100      58      0 0x00080100
      admin usr 0x00000200      admin.usr.0 0x00000200      87      0 0x00080000
      =      =      =      =      0x00000200      56      0 0x00080000
      admin usr 0x00000400      admin.usr.0 0x00000400      70      1 0x00100000
      =      =      =      =      0x00000400      54      0 0x00080000
```

The following example displays detailed block mapping information from blocks 0x100 through 0x400 in the user space of VV admin:

```
cli% showblock -d vv admin usr 0x100 0x400
```

VV/LD	Name	Block	Region
VV usr:	admin	0x00000100	0x00000000-0x0007FFFF (0-256MB)
LD :	admin.usr.0	0x00000100	0x00000000-0x0007FFFF (0-256MB)

  

PD Chnk	PDStep	Block	LDStep
88	0	0x00080000-0x000801FF	0x00000000-0x000001FF
58	0	0x00080000-0x000801FF	0x00000000-0x000001FF

---

VV/LD	Name	Block	Region
VV usr:	admin	0x00000200	0x00000000-0x0007FFFF (0-256MB)
LD :	admin.usr.0	0x00000200	0x00000000-0x0007FFFF (0-256MB)

  

PD Chnk	PDStep	Block	LDStep
87	0	0x00080000-0x000801FF	0x00000200-0x000003FF
56	0	0x00080000-0x000801FF	0x00000200-0x000003FF

---

VV/LD	Name	Block	Region
VV usr:	admin	0x00000400	0x00000000-0x0007FFFF (0-256MB)
LD :	admin.usr.0	0x00000400	0x00000000-0x0007FFFF (0-256MB)

  

PD Chnk	PDStep	Block	LDStep
70	1	0x00100000-0x001001FF	0x00000400-0x000005FF
54	0	0x00080000-0x000801FF	0x00000400-0x000005FF

  

```
cli% showblock -d ld admin.usr.1 0x100
```

VV/LD	Name	Block	Region
VV usr:	admin	0x00080100	0x00080000-0x000FFFFF (256-512MB)
LD :	admin.usr.1	0x00000100	0x00000000-0x0007FFFF (0-256MB)

  

PD Chnk	PDStep	Block	LDStep
65	1	0x00100000-0x001001FF	0x00000000-0x000001FF
59	0	0x00080000-0x000801FF	0x00000000-0x000001FF

## NOTES

None.

---

## COMMAND

showcage

## DESCRIPTION

The `showcage` command displays information about drive cages.

## SYNTAX

The syntax of the `showcage` command can be one of the following:

- `showcage [options] [-d] [<cagename>...]`
- `showcage [-sfp [-d|-ddm]] [<cagename>...]`
- `showcage -i [<cagename>...]`

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

`-d`

Specifies that more detailed information about the drive cage is displayed. If this option is not used, then only summary information about the drive cages is displayed.

`-e`

Displays error information.

`-c`

Specifies to use cached information. This option displays information faster because the cage does not need to be probed, however, some information might not be up-to-date without that probe.

`-sfp`

Specifies information about the SFP(s) attached to a cage. Currently, additional SFP information can only be displayed for DC2 and DC4 cages.

`-ddm`

Specifies the SFP DDM information. This option can only be used with the `-sfp` option and cannot be used with the `-d` option.

-i

Specifies that inventory information about the drive cage is displayed. If this option is not used, then only summary information about the drive cages is displayed.

## SPECIFIERS

<cagename> . . .

Specifies a drive cage name for which information is displayed. This specifier can be repeated to display information for multiple cages. If no specifiers are used, the command defaults to displaying information about all cages in the system.

## RESTRICTIONS

None.

## EXAMPLES

The following examples display information for a single system's drive cages:

```
cli% showcage
Id  Name  LoopA  Pos.A  LoopB  Pos.B  Drives  Temp  RevA  RevB  Model  Side
0  cage0  0:1:1    0  1:1:1    0      40  24-33  2.05  2.05   DC4   n/a
1  cage1  0:1:2    0  1:1:2    0      24  26-32  2.05  2.05   DC2    0
1  cage1  0:1:3    0  1:1:3    0      16  28-32  2.05  2.05   DC2    1
2  cage2  0:1:4    0  1:1:4    0      16  33-36   04   04   DC3   n/a
```

The columns in the previous example are identified as follows:

- Id. The cage number.
- Name. The name of the cage.
- LoopA. The position of the controller node port to which the cage's A port (the top port) is connected.
- Pos . A. Indicates how closely connected the cage's A port is to the controller node port. The value is 0 if the A port is non daisy chained to the node or 1 if daisy chained.
- LoopB. The position of the controller node port that is connected to the cage's B port.
- Pos . B. Indicates how closely connected the cage's B port is to the controller node port. The value is 0 if the B port is non daisy chained to the node or 1 if daisy chained.
- Drives. The number of physical disks in the drive cage.
- Temp. The range of current temperatures for the drives in the drive cage, in Celsius.

- RevA. The cage firmware version for side A of the drive cage.
- RevB. The cage firmware version for side B of the drive cage.
- Model. The model can be DC1, DC2, DC3, or DC4.
- Side. Specifies the right (1) or left (0) side of the drive cage.

Specifying the `-d` option provides more detailed information about the drive cages. See the following E200 Storage Server example for cage0:

```
cli% showcage -d
Id  Name LoopA Pos.A LoopB Pos.B Drives  Temp RevA RevB Model Side
0  cage0 0:0:1    0 1:0:1    0      8 29-32  03  03  DC3  n/a

-----Cage detail info for cage0 -----
Position: ---
-----Midplane Info-----
VendorId,ProductId    3PARdata,DC3
Serial_Num    OPS45811C010719
Node_WWN    20000050CC010719
TempSensor_State    OK
TempSensor_Value    35
OpsPanel_State    OK
Audible_Alarm_State    Muted
ID_Switch    4
Cage_State    OK
Interface Board Info    LoopA    LoopB
Firmware_status    Current    Current
Product_Rev    03    03
IFC_State    OK    OK
ESH_State    OK    OK
Master_CPU    Yes    No
Loop_Map    valid    valid
Link_Speed    2Gbps    2Gbps
Port0_State    OK    OK
Port1_State    No_SFP    No_SFP
Port2_State    No_SFP    No_SFP
Port3_State    No_SFP    No_SFP

Power Supply Info State Fan State AC Model
ps0    OK    MedSpeed OK    --
ps1    OK    MedSpeed OK    --

-----Drive Info----- ----LoopA----- ----LoopB-----
Drive          NodeWWN State Temp(C) ALPA LoopState ALPA LoopState
0:0 20000011c60b0379    OK    32 0x88    OK 0x88    OK
3:0 20000011c60b18b4    OK    29 0x81    OK 0x81    OK
4:0 2000000c50c3560d    OK    31 0x80    OK 0x80    OK
7:0 2000000c501fd935    OK    30 0x79    OK 0x79    OK
8:0 2000000c501fd804    OK    32 0x76    OK 0x76    OK
11:0 20000011c60b0aeb    OK    29 0x73    OK 0x73    OK
12:0 2000000c501fd7ad    OK    32 0x72    OK 0x72    OK
15:0 2000000c501fcb9    OK    31 0x6d    OK 0x6d    OK
```



Specifying the `-d` option provides more detailed information about the drive cages. See the following DC4 example from a T-Class Storage Server example for `cage0`:

```
cli% showcage -d cage0
Id  Name LoopA Pos.A LoopB Pos.B Drives  Temp RevA RevB Model Side
  0  cage0 0:1:1      0 1:1:1      0   40 24-33 2.05 2.05  DC4  n/a

-----Cage detail info for cage0 -----

Position: ---

Fibre Channel Info PortA0 PortB0 PortA1 PortB1
Link_Speed 4Gbps 0Gbps 0Gbps 4Gbps

-----SFP Info-----
FCAL SFP -State-- --Manufacturer-- MaxSpeed(Gbps) TXDisable TXFault RXLoss DDM
  0  0 OK      FINISAR CORP.          4.20 No          No      No      Yes
  0  1 OK      FINISAR CORP.          4.20 No          No      Yes      Yes
  1  0 OK      FINISAR CORP.          4.20 No          No      Yes      Yes
  1  1 OK      FINISAR CORP.          4.20 No          No      No       Yes

Interface Board Info      FCAL0      FCAL1
Link A RXLEDs             Green    Off
Link A TXLEDs             Green    Green
Link B RXLEDs             Off      Green
Link B TXLEDs             Green    Green
LED(Loop_Split)           Off      Off
LEDS(system,hotplug) Green,Off Green,Off

-----Midplane Info-----
Firmware_status           Current
Product_Rev               2.05
State                     Normal Op
Loop_Split                0
VendorId,ProductId        3PARdata,DC4
Unique_ID 1062010000001C00

Power Supply Info State Fan State AC Model
ps0      OK          OK OK   POI
ps1      OK          OK OK   POI
ps2      OK          OK OK   POI
ps3      OK          OK OK   POI

-----Magazine Info----- ---State---
Mag SysLED HplLED Disks LoopA LoopB
  0  Green  Off    4 Ready Ready
  1  Green  Off    4 Ready Ready
  2  Green  Off    4 Ready Ready
  ...
  9  Green  Off    4 Ready Ready

-----Drive Info----- ---LoopA----- ---LoopB-----
Drive      NodeWWN    LED Temp(C) ALPA LoopState ALPA LoopState
  0:0 2000001862c4e410 Green    31 0xe1      OK 0xe1      OK
  0:1 2000001862c4e3f6 Green    31 0xe0      OK 0xe0      OK
  0:2 2000001862c4e3bc Green    28 0xdc      OK 0xdc      OK
  ...
  9:1 2000001862b9affb Green    30 0xa6      OK 0xa6      OK
  9:2 2000001862b9b035 Green    27 0xa5      OK 0xa5      OK
  9:3 2000001862b9b7b5 Green    25 0xa3      OK 0xa3      OK
```

When the power supply is switched off, it will be reported as Off through the `showcage -d` command as well as a degraded alert being posted. The following is a sample of the output:

```
cli% showcage -d cage3 | egrep -i 'Power|\<ps'
Power Supply Info  State Fan State      AC Model
                ps0      OK      OK      OK      MAG
                ps1      Off      OK      OK      MAG
                ps2 Failed      OK Failed      MAG
                ps3      OK      OK      OK      MAG

cli% showalert | grep Cage\ 3
Message          : Cage 3, Power Supply 1 Degraded (Power Supply Off)
Message          : Cage 3, Power Supply 2 Failed (Power Supply Failed, Power
Supply AC Failed)
cli%
```

## NOTES

The power supply model, as shown in the Model column of the `showcage -d` output, is set by service personnel using the `setcage` command, as described in [setcage](#) on page 21.14.

---

**COMMAND**

showcim

**DESCRIPTION**

The `showcim` command displays the CIM server current status, either active or inactive. It also displays the current status of the HTTP and HTTPS ports and their port numbers. In addition, it shows the current status of the SLP port, that is either enabled or disabled.

**SYNTAX**

showcim

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLE**

The following example shows the current CIM status:

```
cli% showcim
-Service- -State- --SLP-- SLPPort -HTTP-- HTTPPort -HTTPS- HTTPSPort PGVer CIMVer
Enabled   Active  Enabled   427 Enabled   5988 Enabled   5989 2.5.1 2.3.1
```

Where the columns are defined as:

- Status. Indicates the CIM server status.
  - ◆ --. Cannot determine the CIM server status.
  - ◆ Enabled. The CIM server is running.
  - ◆ Disabled. The CIM server is not running.

- SLP. Indicates the SLP port state.
  - ◆ --. Cannot determine the SLP port state.
  - ◆ Enabled. The SLP port state is enabled.
  - ◆ Disabled. The SLP port state is disabled.
- SLPPort. The SLP port. Default is 427.
- HTTP. Indicates the HTTP port state.
  - ◆ --. Cannot determine the HTTP port state.
  - ◆ Enabled. HTTP port is enabled.
  - ◆ Disabled. HTTP port is disabled.
- HTTPPort. HTTP port (1024 - 65635). The default value is 5988.
- HTTPS. The HTTPS port state.
  - ◆ --. Cannot determine the HTTPS port state.
  - ◆ Enabled. HTTPS port is enabled.
  - ◆ Disabled. HTTPS port is disabled.
- HTTPSPort. The HTTPS port (1024 - 65535). The default value is 5989.
- PGVer. The Pegasus version.
- CIMVer. Indicates the CIM version running.

## NOTES

None.

---

**COMMAND**

`showclienv`

**DESCRIPTION**

The `showclienv` command displays the CLI environment parameters.

**SYNTAX**

`showclienv`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the currently set CLI environment variables for sample system TestSystem:

```
cli% showclienv
Parameter Value
csvtable      0
nohdtot       0
hafter        -1
listdom       0
editor        emacs
```

**NOTES**

See the `setclienv` command for a complete description of the environment parameters.

---

**COMMAND**

showcpg

**DESCRIPTION**

The `showcpg` command displays Common Provisioning Groups (CPGs) in the system.

**SYNTAX**

`showcpg [options <arg>] [<CPG_name>...|<pattern>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

The following options can not be used together with other options except `-hist` and `-domain`:

`-d`

Show the detail information of CPGs.

`-r`

Specifies that raw space used by the CPGs is displayed.

`-alert`

Indicates whether alerts are posted.

`-alerttime`

Show times when alerts were posted (when applicable).

`-sag`

Specifies that the snapshot administration space autogrowth parameters are displayed.

`-sdg`

Specifies that the snapshot data space autogrowth parameters are displayed.

The following options can be used together with other options:

`-hist`

Specifies that current data from the CPG, as well as the CPG's history data is displayed.

`-domain <domain_name_or_pattern,...>`

Shows only CPGs that are in domains with names matching one or more of the `<domain_name_or_pattern>` argument. This option does not allow listing objects within a domain of which the user is not a member. Patterns are glob-style (shell-style) patterns (see Help on `sub,globpat`).

## SPECIFIERS

`[<CPG_name>...|<pattern>...]`

Specifies that CPGs matching either the specified CPG name or those CPGs matching the specified pattern are displayed. This specifier can be repeated to display information for multiple CPGs. If not specified, all CPGs in the system are displayed. See [Glob-Style Pattern](#) on page 2.4 for more information.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays a system's CPGs:

```
cli% showcpg
```

----- (MB) -----												
		-Volumes-		-Usage-		--- Usr ---		--- Snp ---		-- Adm --		
Id	Name	Warn%	VVs	TPVVs	Usr	Snp	Total	Used	Total	Used	Total	Used
0	cpg1	-	2	0	2	1	15360	15360	32768	0	8192	0
1	cpg2	-	2	2	0	2	0	0	32768	1024	8192	256
2	cpg3	-	1	1	0	1	0	0	32768	512	8192	128
-----												
3	total				2	4	15360	15360	98304	1536	24576	384

The columns for the previous example output are identified as follows:

- Id. The CPG ID.
- Name. The CPG name.
- Warn%. The CPG's allocation warning threshold.
- Volumes. Virtual Volumes (VVs) and Thinly Provisioned Virtual Volumes (TPVVs).
  - ◆ VVs. The Number of VVs (including TPVVs) using the CPG.
  - ◆ TPVVs. The number of TPVVs using the CPG.

- Usage. Virtual Volumes using the CPG.
  - ◆ Usr. The number of VVs whose User Space is using the CPG.
  - ◆ Snp. The number of VVs whose Snapshot Space is using the CPG.
- Usr. User space.
  - ◆ Total. The total logical disk space in MBs in the Usr space.
  - ◆ Used. The total logical disk space in MBs used in Usr space.
- Snp. Snap space.
  - ◆ Total. The total logical disk space in MBs in the Snp space.
  - ◆ Used. The total logical disk space in MBs used in Snp space.
- Adm. Administration space.
  - ◆ Total. The total logical disk space in MBs in the Adm space.
  - ◆ Used. The total logical disk space in MBs used in Adm space.

The following example displays the detailed information of common provisioning groups.

cli% showcpg -d															
----- (MB) -----															
		-Volumes-		-Usage-		--- Usr ---		--- Snp ---		--- Adm ---		--- LD ---			
Id	Name	Warn%	VVs	TPVVs	Ushr	Snp	Total	Used	Total	Used	Total	Used	Ushr	Snp	Adm
0	cpg1	-	2	0	2	1	15360	15360	32768	0	8192	0	4	4	2
1	cpg2	-	2	2	0	2	0	0	32768	1024	8192	256	0	4	3
2	cpg3	-	1	1	0	1	0	0	32768	512	8192	128	0	4	2
-----															
3	total				2	4	15360	15360	98304	1536	24576	384	8	12	7

The columns in the previous example are defined as follows:

- Id. The CPG ID.
- Name. The CPG name.
- Warn%. The CPG's allocation warning threshold.
- Volumes. Virtual Volumes (VVs) and Thinly Provisioned Virtual Volumes (TPVVs).
  - ◆ VVs. The Number of VVs (including TPVVs) using the CPG.
  - ◆ TPVVs. The number of TPVVs using the CPG.



- Usage. Virtual Volumes using the CPG.
  - ◆ **Usr.** The number of VVs whose User Space is using the CPG.
  - ◆ **Snp.** The number of VVs whose Snapshot Space is using the CPG.
- **Usr. User space.**
  - ◆ **Total.** The total logical disk space in MBs in the Usr space.
  - ◆ **Used.** The total logical disk space in MBs used in Usr space.
- **Snp. Snap space.**
  - ◆ **Total.** The total logical disk space in MBs in the Snp space.
  - ◆ **Used.** The total logical disk space in MBs used in Snp space.
- **Adm. Administration space.**
  - ◆ **Total.** The total logical disk space in MBs in the Adm space.
  - ◆ **Used.** The total logical disk space in MBs used in Adm space.
- **LD.** The number of logical disks in Usr, Snp, and Adm space.
  - ◆ **Usr.** The total number of logical disks in the Usr space.
  - ◆ **Snp.** The total number of logical disks in the Snp space.
  - ◆ **Adm.** The total number of logical disks in the Adm space.

The following example displays the common provisioning group's snapshot data space's autogrowth parameters:

```
cli% showcpg -sdg
      -----(MB)-----
Id Name Warn Limit  Grow Args
0  cpg1   70    85 32768 -p -devtype FC
1  cpg2    -    - 32768 -p -devtype FC
2  cpg3    -    - 32768 -p -devtype FC
```

The columns in the previous example are identified as follows:

- **Id.** The CPG ID.
- **Name.** The CPG name.
- **Warn.** The CPG's snapshot data space allocation warning threshold in MBs.

- **Limit.** The CPG's snapshot data allocation limit threshold, or growth limit in MBs.
- **Grow.** The CPG's snapshot data space allocation limit threshold.
- **Args.** The options used in the creation of the CPG.

The following example displays the raw space used by the system's common provisioning groups:

```
cli% showcpg -alert
```

----- Data ----- --- Adm ----												
				- Setting(MB) -		Alerts				Alerts		
Id	Name	Warn%	Total	Warn	Limit	W%	W	L	F	Total	W%	F
0	cpg1	-	32768	70	85	-	-	Y	-	8192	-	-
1	cpg2	-	32768	-	-	-	-	-	-	8192	-	-
2	cpg3	-	32768	-	-	-	-	-	-	8192	-	-
3	cpg4	-	32768	-	-	-	-	-	-	8192	-	-

The columns in the previous example are identified as follows:

- **Id.** The physical disk ID.
- **Name.** The name of the common provisioning group.
- **Warn%.** The allocation warning percentage (see `-aw` option of the `createcpg` or `setcpg` command).
- **Total.** The total logical disk MB space in a CPG.
- **Warn.** The size in MB of the space at which a warning alert is generated.
- **Limit.** The limit size in MB beyond which the space (Data or Adm) will not grow.
- **Alerts.** There are columns for the 4 types of alerts listed below. For the `-alert` option indicates Y if an alert is posted or "-" if no alert has been posted. The `-alerttime` option indicates the time that the alert was posted.
  - ◆ **W%.** An alert corresponding to Warn%.
  - ◆ **W.** An alert corresponding to Warn.
  - ◆ **L.** An alert corresponding to Limit.
  - ◆ **F.** An alert corresponding to a growth failure.

The following example displays the raw space used by common provisioning groups:

cli% showcpg -r

----- (MB) -----																		
		-Volumes-		-Usage-		Usr				Snp				Adm				
Id	Name	Warn%	VVs	TPVVs	Usr	Snp	Total	RTotal	Used	RUsed	Total	RTotal	Used	RUsed	Total	RTotal	Used	RUsed
0	cpg1	-	2	0	2	1	15360	30720	15360	30720	32768	65536	0	8192	24576	0	0	0
1	cpg2	-	2	2	0	2	0	0	0	0	32768	65536	1024	2048	8192	24576	256	768
2	cpg3	-	1	1	0	1	0	0	0	0	32768	65536	512	1024	8192	24576	128	384
-----																		
3	total				2	4	15360	30720	15360	30720	98304	196608	1536	11264	40960	131072	384	1152

The columns in the previous example are identified as follows:

- Id. The CPG ID.
- Name. The CPG name.
- Warn%. The CPG's allocation warning threshold.
- Volumes. Virtual Volumes (VVs) and Thinly Provisioned Virtual Volumes (TPVVs).
  - ◆ VVs. The Number of VVs (including TPVVs) using the CPG.
  - ◆ TPVVs. The number of TPVVs using the CPG.
- Usage. Virtual Volumes using the CPG.
  - ◆ Usr. The number of VVs whose User Space is using the CPG.
  - ◆ Snp. The number of VVs whose Snapshot Space is using the CPG.
- Usr. User space.
  - ◆ Total. The total logical disk space in MBs in the Usr space.
  - ◆ RTotal. The total raw MBs of (raw) space in the Usr space.
  - ◆ Used. The total logical disk space in MBs used in Usr space.
  - ◆ RUsed. The total MBs of (raw) space used in the Usr space.
- Snp. Snp space.
  - ◆ Total. The total logical disk space in MBs in the Snp space.
  - ◆ RTotal. The total raw MBs of (raw) space in the Snp space.
  - ◆ Used. The total logical disk space in MBs used in Snp space.
  - ◆ RUsed. The total MBs of (raw) space used in the Snp space.

- **Adm.** Administration space.
  - ◆ **Total.** The total logical disk space in MBs in the Adm space.
  - ◆ **RTotal.** The total raw MBs of (raw) space in the Adm space.
  - ◆ **Used.** The total logical disk space in MBs used in Adm space.
  - ◆ **RUsed.** The total MBs of (raw) space used in the Adm space.

## NOTES

- When using the `createaldevv` command, the size of the Logical Disk (LD) space created is the first integer multiple of the RAID set size that is large enough to accommodate the requested virtual volume size.

For example, with the default RAID-5 layout with a set size of 768 MB, a requested virtual volume size of 8192 MB causes the creation of LDs with a total size rounded up to an integer multiple of 768 that is 8448 MB. The growth increment of CPGs is similarly rounded up because the growth is done by creating LDs that must be created in units of the LD RAID set size. See the *InForm OS Administrator's Guide* for further details.

- For this command, 1 MB = 1048576 bytes.
- A `Domain` column can be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.

---

**COMMAND**

showdate

**DESCRIPTION**

The `showdate` command displays the date and time for each system node.

**SYNTAX**

showdate

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the date and time for the system node:

```
cli% showdate
Node Date
0    Mon Apr 17 17:07:44 PDT 2007 (US/Pacific)
1    Mon Apr 17 17:07:44 PDT 2007 (US/Pacific)
```

**NOTES**

Set date and time information on nodes by issuing the `setdate` command. The output of the `showdate` command include an abbreviated time zone which may be non-unique. A unique and parenthesized long format time zone name is appended at the end of each row. The long format time zone name is the official time zone string set via `setdate`. See [setdate](#) on page 21.27 for additional information.

---

**COMMAND**

`showdomain`

**DESCRIPTION**

The `showdomain` command displays a list of domains in a system.

**SYNTAX**

`showdomain [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d`

Specifies that detailed information is displayed.

`-sortcol <col> [ , <dir> ][ : <col> [ , <dir> ] ... ]`

Sorts command output based on the column number `<col>`. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays detailed information about a system's domains:

```
cli% showdomain -d
ID Domain      -----CreationTime----- --Comments--- -VVRetentionTimeMax-
1 TestDomain1  2009-08-22 14:23:30 PDT Beef Test                0 Hours
2 TestDomain2  2009-08-22 14:23:30 PDT Chicken Test              0 Hours
3 TestDomain3  2009-08-22 14:23:30 PDT Pork Test                  0 Hours
```

In the example above:

- ID. The domain ID.
- Domain. The domain name.
- CreationTime. The date and time the domain was created.
- Comments. Information about the domain.
- VVRetentionTimeMax. Maximum retention time that may be set for volumes in this domain; if --, the system VVRetentionTimeMax value is used instead.

## NOTES

If the VVRetentionTimeMax is --, then the minimum volume retention time for the system is used instead. If the VVRetentionTimeMax is 0, then the volume retention time in the domain is disabled.

---

**COMMAND**

`showdomainset`

**DESCRIPTION**

The `showdomainset` command lists the host sets defined on the InServ and their members.

**SYNTAX**

`showdomainset [options] [<setname_or_pattern>...]`

`showdomainset -domain [options] [<domainname_or_pattern>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d`

Show a more detailed listing of each set.

`-domain`

Show domain sets that contain the supplied domains or patterns

**SPECIFIERS**

`<setname_or_pattern>...`

An optional list of `<setnames>` or `<patterns>`. If no `<setname>` or `<pattern>` is specified all sets are displayed, otherwise only sets with names matching one or more of the setnames or patterns are displayed. The patterns are glob-style patterns (see help on `sub,globpat`).

`<domainname_or_pattern>...`

Specifies that the domain sets containing domains with the specified names or matching the glob-style patterns should be displayed.

**RESTRICTIONS**

None.



## EXAMPLES

To show all domain sets defined to the system:

```
cli% showdomainset
Id Name      Members
  7 domainset domainset.1
              domainset.2
 23 newset    testdomain
```

Show details of a specific set:

```
cli% showdomainset -d newset
Id Name  Members  Comment
 23 newset testdomain this set has been renamed and now has a comment
```

Show domain sets containing domains matching the pattern domainset.\*:

```
cli% showdomainset -domain domainset.*
Id Name      Members
  7 domainset domainset.1
              domainset.2
```

## NOTES

None.

---

**COMMAND**

showeeprom

**DESCRIPTION**

The `showeeprom` command displays node EEPROM log information.

**SYNTAX**

`showeeprom [options] [<node_ID>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-dead`

Specifies that an EEPROM log for a node that has not started or successfully joined the cluster be displayed. If this option is used, it must be followed by a list of nodes.

**SPECIFIERS**

`<node_ID> . . .`

Specifies the node ID for which EEPROM log information is retrieved. Multiple node IDs are separated with a single space (0 1 2). If no specifiers are used, the EEPROM log for all nodes is displayed.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays the EEPROM log for all nodes:

```
cli% showeeprom
Node: 0
-----
Board revision: 0920-1053-03.01
Assembly: FLH 2007/50 Serial 0039
System serial: 1000183
BIOS version: 2.1.3
OS version: 2.2.4.32
Reset reason: Unknown
Last boot: 2008-02-28 14:58:25 PST
Last cluster join: 2008-02-28 14:58:37 PST
Last panic: Never
Last panic request: Never
Error ignore code: 00
SMI context: 00
Last HBA mode: 2a000000
BIOS state: ff 23 26 27 28 29 2b 80
TPD state: 34 40 ff 2a 2c 2e 30 32
Code 128 (BIOS update) - Subcode 0x2020103 (2020102) Thu Feb 28 14:54:13 2008
Code 25 (PROM Failure) - Subcode 0x1 (0) Fri Feb 15 07:58:48 2008
[Repeat] Fri Feb 15 07:59:20 2008
Code 128 (BIOS update) - Subcode 0x2020102 (2020009) Fri Feb 15 07:58:39 2008
Code 25 (PROM Failure) - Subcode 0x1 (0) Fri Feb 15 07:56:33 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Thu Feb 14 23:30:05 2008
[Repeat] Fri Feb 15 08:36:13 2008
Code 25 (PROM Failure) - Subcode 0x1 (0) Tue Jan 29 13:31:35 2008
[Repeat] Tue Jan 29 13:32:07 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Tue Jan 29 13:30:56 2008
Code 25 (PROM Failure) - Subcode 0x1 (0) Tue Jan 29 13:17:57 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Tue Jan 29 13:17:28 2008
Code 25 (PROM Failure) - Subcode 0x1 (0) Tue Jan 29 13:14:49 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Tue Jan 29 13:04:57 2008
Code 10 (PCI Failure) - Subcode 0xd (30) Mon Jan 14 18:17:12 2008

Node: 1
-----
Board revision: 0920-1053-03.01
Assembly: FLH 2007/50 Serial 0042
System serial: 1000183
BIOS version: 2.1.3
OS version: 2.2.4.32
Reset reason: COLD_POWERON
Last boot: 2008-02-28 14:59:14 PST
Last cluster join: 2008-02-28 14:59:32 PST
Last panic: Never
Last panic request: Never
Error ignore code: 00
SMI context: 00
Last HBA mode: 2a000000
BIOS state: ff 23 26 27 28 29 2b 80
TPD state: 34 40 ff 2a 2c 2e 30 32
Code 128 (BIOS update) - Subcode 0x2020103 (2020102) Thu Feb 28 14:54:18 2008
Code 128 (BIOS update) - Subcode 0x2020102 (2020009) Fri Feb 15 08:10:03 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Thu Feb 14 23:34:54 2008
[Repeat] Fri Feb 15 08:41:02 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Tue Jan 29 13:58:56 2008
Code 15 (PCI_Fibre_Failure) - Subcode 0x0 (0) Tue Jan 29 13:50:54 2008
[Repeat] Tue Jan 29 13:51:02 2008
Code 15 (PCI_Fibre_Failure) - Subcode 0x0 (5) Tue Jan 29 13:50:38 2008
Code 25 (PROM Failure) - Subcode 0x0 (0) Tue Jan 29 13:49:11 2008
Code 20 (AP Init Failure) - Subcode 0x0 (0) Tue Jan 29 13:05:32 2008
Code 25 (PROM Failure) - Subcode 0x0 (0) Wed Jan 23 15:03:20 2008
Code 25 (PROM Failure) - Subcode 0x0 (0) Mon Jan 21 00:14:45 2008
Code 25 (PROM Failure) - Subcode 0x0 (0) Fri Jan 18 11:38:31 2008
Code 15 (PCI_Fibre_Failure) - Subcode 0x0 (5) Tue Jan 15 18:49:06 2008
```

In the example above:

- `Board revision` indicates the 3PAR part number for the node board, including any minor revision codes.
- `Assembly` indicates the code of the assembly, the year or week the node was produced, and the node serial number.
- `System serial` is the cluster serial number.
- `BIOS version` displays the currently installed BIOS version.
- `OS version` displays the currently installed OS version.
- `Reset reason` displays why the board was previously reset. Values can be:
  - ◆ `COLD_POWERON`. The node was powered off and back on.
  - ◆ `EXTERNAL_RESET`. Another node in the cluster forced the reset.
  - ◆ `WATCHDOG_RESET`. A watchdog timer forced the reset.
  - ◆ `PCI_RESET`. The node restarted.
- `Last boot` displays the time the node last started the OS.
- `Last cluster join` displays the time the node last joined the cluster.
- `Last panic` displays the last time another node in the cluster requested this node to take a panic.
- `Error ignore code` is for engineering use only.
- `SMI context` is for engineering use only.
- `Last HBA mode` contains PCI Fibre port settings forcing certain ports to start up in initiator mode and certain ports to start up in target mode.
- `BIOS state` is for engineering use only.
- `TPD state` is for engineering use only.

## NOTES

None.

---

**COMMAND**

showeventlog

**DESCRIPTION**

The `showeventlog` command displays the current system event log.

**SYNTAX**

`showeventlog [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-min <number>`

Specifies that only events occurring after the specified number of minutes are shown. The `<number>` is an integer from 1 through 2147483647.

`-more`

Specifies that you can page through several events at a time.

`-oneline`

Specifies that each event is formatted on one line.

`-d`

Specifies that detailed information is displayed.

`-startt <time>`

Specifies that only events after a specified time are to be shown. The `time` argument can be specified as either `<timespec>`, `<datespec>`, or both.

`<timespec>`

Specified as the hour (`hh`), as interpreted on a 24 hour clock, where minutes (`mm`) and seconds (`ss`) can be optionally specified. Acceptable formats are `hh:mm:ss` or `hhmm`.

`<datespec>`

Specified as the month (`mm` or `month_name`) and day (`dd`), where the year (`yy`) can be optionally specified. Acceptable formats are `mm/dd/yy`, `month_name dd`, `dd month_name yy`, or `yy-mm-dd`. If the syntax `yy-mm-dd` is used, the year must be specified.

`-endt <time>`

Specifies that only events before a specified time are to be shown. The `time` argument can be specified as either `<timespec>`, `<datespec>`, or both.

`<timespec>`

Specified as the hour (`hh`), as interpreted on a 24 hour clock, where minutes (`mm`) and seconds (`ss`) can be optionally specified. Acceptable formats are `hh:mm:ss` or `hhmm`.

`<datespec>`

Specified as the month (`mm` or `month_name`) and day (`dd`), where the year (`yy`) can be optionally specified. Acceptable formats are `mm/dd/yy`, `month_name dd`, `dd month_name yy`, or `yy-mm-dd`. If the syntax `yy-mm-dd` is used, the year must be specified.



**NOTE:** The `pattern` argument in the following options is a regular expression pattern that is used to match against the events each option produces. For each option, the `pattern` argument can be specified multiple times. For example:  
`showeventlog -type Disk.* -type <Tpdctl client>`  
`-sev Major` displays all events of severity `Major` and with a type that matches either the regular expression `Disk.*` or `<Tpdctl client>`.

`-sev <pattern>`

Specifies that only events with severities that match the specified `pattern(s)` are displayed.

`-nsev <pattern>`

Specifies that only events with severities that do not match the specified `pattern(s)` are displayed.

`-class <pattern>`

Specifies that only events with classes that match the specified `pattern(s)` are displayed.

`-nclass <pattern>`

Specifies that only events with classes that do not match the specified `pattern(s)` are displayed.

`-node <pattern>`

Specifies that only events from nodes that match the specified `pattern(s)` are displayed.

`-nnode <pattern>`

Specifies that only events from nodes that do not match the specified pattern(s) are displayed.

`-type <pattern>`

Specifies that only events with types that match the specified pattern(s) are displayed.

`-ntype <pattern>`

Specifies that only events with types that do not match the specified pattern(s) are displayed.

`-msg <pattern>`

Specifies that only events, whose messages match the specified pattern(s), are displayed.

`-nmsg <pattern>`

Specifies that only events, whose messages do not match the specified pattern(s), are displayed.

`-comp <pattern>`

Specifies that only events, whose components match the specified pattern(s), are displayed.

`-ncomp <pattern>`

Specifies that only events, whose components do not match the specified pattern(s), are displayed.

## SPECIFIERS

None.

## RESTRICTIONS

The InForm CLI stores 15 MB of event logs. If the number of logs exceeds the 15 MB limit, old logs are deleted. After a log is deleted, it cannot be recovered.

## EXAMPLES

The following example displays the system events occurred during the last 20 minutes:

```
cli% showeventlog -min 20
Mon Sep 29 09:48:07 2003 PST
Node: 0, Seq: 51, Class: Status change, Severity: Informational, Type: Node booted
Node 0 has booted
Mon Sep 29 09:48:13 2003 PST
Node: 1, Seq: 50, Class: Status change, Severity: Informational, Type: Node booted
Node 1 has booted
Mon Sep 29 09:48:53 2003 PST
Node: 0, Seq: 334, Class: Status change, Severity: Informational, Type: Cage loop
status
Cage 0 is connected to the system on both ports. Autofixing previous alerts.
Mon Sep 29 09:48:53 2003 PST
Node: 0, Seq: 337, Class: Status change, Severity: Informational, Type: Cage loop
status
Cage 1 is connected to the system on both ports. Autofixing previous alerts.
Mon Sep 29 09:48:54 2003 PST
Node: 0, Seq: 419, Class: Status change, Severity: Informational, Type: Notification
VV 0 has started
Mon Sep 29 09:48:54 2003 PST
Node: 0, Seq: 421, Class: Status change, Severity: Informational, Type: Notification
VV 1 has started
```

The following example displays the system eventlog using the `-oneline` option:

```
cli% showeventlog -min 20 -oneline
Time                TZn Node  Seq Class          Severity      Type
Message
Fri Sep 5 19:42:04 2003 PST    0 1154 Status change  Minor         Process has died
Process /opt/tpd/bin/tpdtcl /opt/tpd/bin/tpdtcl.tcl has died on node 0
Fri Sep 5 19:42:04 2003 PST    0 1155 Status change  Informational Change in alert
state Alert 122 changed from state Resolved by System to New
Fri Sep 5 19:42:09 2003 PST    0 1157 Status change  Informational Change in alert
state Alert 122 changed from state New to Resolved by System
Fri Sep 5 19:52:19 2003 PST    0 1192 Status change  Minor         Process has died
Process /opt/tpd/bin/tpdtcl /opt/tpd/bin/tpdtcl.tcl has died on node 0
Fri Sep 5 19:52:19 2003 PST    0 1193 Status change  Informational Change in alert
state Alert 122 changed from state Resolved by System to New
```

## NOTES

The Inform OS stores 15 MB of event logs. If the number of logs exceeds the 15 MB limit, old logs are deleted. After a log is deleted, it cannot be recovered.



---

**COMMAND**

showfirmwaredb

**DESCRIPTION**

The `showfirmwaredb` command displays the current database of firmware levels for possible upgrade. If issued without any options, the firmware for all vendors is displayed.

**SYNTAX**

`showfirmwaredb [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-n <vendor_name>`

Specifies that the firmware vendor from the SCSI database file is displayed.

`-l`

Reloads the SCSI database file into the system.

`-all`

Specifies current and past firmware entries are displayed. If not specified, only current entries are displayed.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting `<dir>` can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns are sorted by values in later columns.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays the current database of firmware levels and prints firmware data:

```
cli% showfirmwaredb
```

Vendor	Prod_rev	Dev_Id	Fw_status	Cage_type	Firmware_File
HITACHI	[C1C1]	DK..DJ-18FC	Current	DC4	/opt/tpd/fw/drive/MDJFC1C1.BIN
HITACHI	[C1C1]	DK..DJ-72FC	Current	DC4	/opt/tpd/fw/drive/MDJFC1C1.BIN
HITACHI	[JLAK]	DK..CJ-18FC	Current	DC4	/opt/tpd/fw/drive/MCJF_LAK.BIN
HITACHI	[JLAK]	DK..CJ-36FC	Current	DC4	/opt/tpd/fw/drive/MCJF_LAK.BIN
HITACHI	[JLAK]	DK..CJ-72FC	Current	DC4	/opt/tpd/fw/drive/MCJF_LAK.BIN
HITACHI	[GLAK]	DK..CJ-18FC	Current	DC4	/opt/tpd/fw/drive/MCJF_LAK.BIN
HITACHI	[GLAK]	DK..CJ-36FC	Current	DC4	/opt/tpd/fw/drive/MCJF_LAK.BIN
HITACHI	[GLAK]	DK..CJ-72FC	Current	DC4	/opt/tpd/fw/drive/MCJF_LAK.BIN
HITACHI	[C0C0]	DK..EJ-36FC	Current	DC4.DC2	/opt/tpd/fw/drive/MEJFC0C0.BIN
HITACHI	[C0C0]	DK..EJ-72FC	Current	DC4.DC2	/opt/tpd/fw/drive/MEJFC0C0.BIN
HITACHI	[C0C0]	DK..EJ-14FC	Current	DC4.DC2	/opt/tpd/fw/drive/MEJFC0C0.BIN
HITACHI	[A6A6]	HUS157336ELF200	Current	ALL	/opt/tpd/fw/drive/MEKFA6A6.BIN
HITACHI	[A6A6]	HUS157373ELF200	Current	ALL	/opt/tpd/fw/drive/MEKFA6A6.BIN
HITACHI	[FA16]	HUS103014FLF210	Current	ALL	/opt/tpd/fw/drive/MFJFFA16.BIN
HITACHI	[FA16]	HUS103030FLF210	Current	ALL	/opt/tpd/fw/drive/MFJFFA16.BIN
HITACHI	[F7A7]	DK..BJ-xxFC	Current	DC4	/opt/tpd/fw/drive/3BJF_7A7.BIN
SEAGATE	[0004]	ST39103FC	Current	DC4	/opt/tpd/fw/drive/0004.lod
SEAGATE	[0004]	ST318203FC	Current	DC4	/opt/tpd/fw/drive/0004.lod
SEAGATE	[0005]	ST318304FC	Current	DC4	/opt/tpd/fw/drive/0005.lod
SEAGATE	[0005]	ST336704FC	Current	DC4	/opt/tpd/fw/drive/0005.lod

```
...
Source file: /var/opt/tpd/scsi_db.cfg
```

## NOTES

The firmware information displayed by the `showfirmwaredb` command is used when issuing the `upgradecage` and `upgradepd` commands. See [upgradecage](#) on page 30.2 and [upgradepd](#) on page 30.4 for more information.

---

**COMMAND**

showhost

**DESCRIPTION**

The `showhost` command displays information about defined hosts and host paths in the system.

**SYNTAX**

`showhost [options <arg>] [<host_name>...|<pattern>...|<host_set>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d`

Shows a detailed listing of the host and path information. This option can only be used with `-agent` and `-domain` options.

`-verbose`

Shows a verbose listing of all host information.

`-chap`

Shows the CHAP authentication information.

`-desc`

Shows the host descriptor information.

`-agent`

Shows information provided by host agent.

`-pathsum`

Shows summary information about hosts and paths. This option cannot be used with the `-d` option.

`-persona`

Shows the host persona settings in effect. This option cannot be used with the `-d` option.

**-listpersona**

Lists the defined host personas. This option cannot be used with the `-d` option. Different host personas support different host operating systems and have different capabilities. Host personas may have the following additional capabilities:

- ◆ **UAREpLun.** Sends a unit attention when the LUN list changes due to adding or removing VLUNs.
- ◆ **RTPG.** Enables the Report Target Port Group (RTPG) command and asymmetric state change unit attention when path counts change due to adding or removing ports in the host's definition.
- ◆ **VolSetAddr.** Enables HPUX Volume Set Addressing (VSA).
- ◆ **SoftInq.** Enables inquiry data formats for hosts such as Egenera and NetApp.
- ◆ **NACA.** Enables Normal Auto Contingent Allegiance (NACA) bit for AIX.
- ◆ **SESLun.** Enables SCSI Enclosure Services (SES) LUN for host agent support.

**-noname**

Shows only host paths (WWNs and iSCSI names) that are not assigned to any host. This option cannot be used with the `-d` option.

**-domain <domainname\_or\_pattern,...>|<domain\_set>**

Shows only hosts that are in domains or domain sets that match one or more of the specifier `<domain_name_or_pattern>` or `set:<domain_set>` arguments. The set name `<domain_set>` must start with `set:`. This option does not allow listing objects within a domain of which the user is not a member.

**SPECIFIERS****<host\_name>...**

Name of the host up to 31 characters in length. This specifier can be repeated to set properties for multiple hosts.

**<pattern>...**

Specifies that information is shown for all hosts matching the specified glob-style pattern. This specifier can be repeated to set properties for multiple hosts using different patterns.

`set:<host_set>...`

Specifies that information about all the hosts that are members of set `<host_set>` should be displayed. The set name `<host_set>` must start with `set:`. Acts as if all members were individually specified as parameters to the comment. May be repeated to specify multiple host sets.

## RESTRICTIONS

None.

## EXAMPLES



**NOTE:** For the following two examples, the Domain column appears only if the `-listdom` global option or `TPDLISTDOM` environment variable was set prior to starting the CLI.

The following example displays detailed host and path information:

```
cli% showhost -d
Id Name          Persona -----WWN/iSCSI_Name----- Port  IP_addr
0 pe750-07-iscsi Generic 210000E08B023F71 1:3:1 0.0.0.0
1 adt           Generic 210000E08B056C21 0:2:1 0.0.0.0
1 adt           Generic 210100E08B256C21 1:2:1 0.0.0.0
-- --           Generic 210000E08B023C71 1:5:1 0.0.0.0
-- --           Generic 210000E08B023F77 0:2:2 0.0.0.0
```

The following example displays host descriptor properties:

```
cli% showhost -desc
----- Host queasy10 -----
Name       : queasy10
Domain     : -
Id         : 0
Location   : Rack 35, Position 8
IP Address : --
OS         : --
Model      : --
Contact    : --
Comment    : --
```

The following example displays host CHAP properties:

```
cli% showhost -chap
Id Name      -Initiator_CHAP_Name- -Target_CHAP_Name-
 0 queasy10 queasy10              s019
 1 foo       --                  --
```

The following example displays all host information:

```
cli% showhost -verbose
Id Name      Persona  -----WWN/iSCSI_Name----- Port  IP_addr
 0 queasy10 Generic  210000E08B027B60              0:0:1 n/a
 1 foo       Generic  myiscsipath.3pardata.com ---   0.0.0.0

Id Name      -Initiator_CHAP_Name- -Target_CHAP_Name-
 0 queasy10 queasy10              s019
 1 foo       --                  --

----- Host queasy10 -----
Name       : queasy10
Domain     : --
Id         : 0
Location   : Rack 35, Position 8
IP Address : --
OS         : --
Model      : --
Contact    : --
Comment    : --

Id Name      Persona_Id Persona_Name Persona_Caps
 0 queasy10      0 Generic      SESLun
 1 foo          0 Generic      SESLun
```

## NOTES

- If host names or patterns are specified, then hosts with names that match any of the patterns are listed. Otherwise all hosts are listed. Patterns are glob-style (shell-style) patterns (see Help on sub, globpat).
- Host descriptor information is available only for hosts that have been assigned a name through the `createhost` command.
- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or

if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables command.

---

**COMMAND**

showhostset

**DESCRIPTION**

The showhostset command lists the host sets defined on the InServ and their members.

**SYNTAX**

showhostset [options] [<setname\_or\_pattern>...]

showhostset -host [options] [<hostname\_or\_pattern>...]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-d

Show a more detailed listing of each set.

-host

Show host sets that contain the supplied hostnames or patterns.

**SPECIFIERS**

<setname\_or\_pattern>...

An optional list of setnames or patterns. If no <setname> or <pattern> is specified all sets are displayed, otherwise only sets with names matching one or more of the setnames or patterns are displayed. The patterns are glob-style patterns (see help on sub,globpat).

<hostname\_or\_pattern>...

Specifies that the sets containing hosts with the specified names or matching the glob-style patterns should be displayed.

**RESTRICTIONS**

None.



## EXAMPLES

Show all host sets defined to the system:

```
cli% showhostset
Id Name      Members
22 myset     -
18 seta      -
  3 sunv40z-09 sunv40z-09-0
                    sunv40z-09-1
                    sunv40z-09-2
                    sunv40z-09-3
```

Show the details of myset only:

```
cli% showhostset -d myset
Id Name  Members Comment
22 myset -          This is an empty set
```

Show the host sets containing host sun40z-09-0:

```
cli% showhostset -host sun40z-09-0
Id Name      Members
  3 sunv40z-09 sunv40z-09-0
                    sunv40z-09-1
                    sunv40z-09-2
                    sunv40z-09-3
```

## NOTES

A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the `TPDLISTDOM` environment variable set. Run `cli -h` and `setclienv -h` for details of the environment variables.

---

**COMMAND**

showinventory

**DESCRIPTION**

Shows information about all the hardware components in the system.

**SYNTAX**

showinventory

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

None.

## NOTES

The inventory is divided into sections and inventory for each of these sections can be shown separately with the following command:

Section	Command
Node	<code>shownode -i</code>
Battery	<code>showbattery -i</code>
Port	<code>showport -i</code>
Cage	<code>showcage -i</code>
Disk	<code>showpd -i</code>

---

**COMMAND**

showiscsisession

**DESCRIPTION**

The showiscsisession command shows the iSCSI sessions.

**SYNTAX**

showiscsisession [option <arg>]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-sortcol <col>[, <dir>][:<col>[, <dir>]...]

Sorts command output based on column number <col>. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting <dir> can be specified as follows:

inc

Sort in increasing order (default).

dec

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:), and rows that have the same earlier columns will be sorted by the values in the later columns.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the iSCSI session information:

```
cli% showiscsisession
N:S:P ---IPAddr---- TPGT TSIH Conns -----iSCSI_Name----- StartTime-----
1:3:1 192.168.2.181 131 13 1 ign.1991-05.com.microsoft:pe750-07 Tue Nov 29 12:43:25 PST 2005
```

In the example above:

- `N:S:P`. The port location in Node:Slot:Port format, where Slot is the PCI slot number in the node and Port is the port number in the PCI card.
- `IPAddr`. The IP address of the Ethernet port.
- `TPGT`. The Target Portal Group Tag.
- `TSIH`. The Target Session Identifying Handle.
- `Conns`. The number of connections for the session.
- `iSCSI_Name`. The iSCSI name of the host.
- `StartTime`. The time the session was started.

## NOTES

None

---

## COMMAND

showld

## DESCRIPTION

The `showld` command displays configuration information about the system's Logical Disks (LDs).

## SYNTAX

```
showld [options <arg>] [<LD_name>|<pattern>...]
```

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

`-cpg <CPG_name>|pattern>...`

Requests that only logical disks in Common Provisioning Groups (CPGs) that match the specified CPG names or patterns be displayed. Multiple CPG names or patterns can be repeated using a comma-separated list (for example `-cpg <CPG_name>,<CPG_name>...`).

`-vv <VV_name>|pattern>...`

Requests that only logical disks mapped to virtual volumes that match any of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list (for example `-vv <VV_name>,<VV_name>...`).

`-domain <domainname>|pattern>...`

Only shows logical disks that are in domains with names that match any of the names or specified patterns. Multiple domain names or patterns can be repeated using a comma-separated list (for example `-vv <domainname_name>,<domainname_name>...`).



**NOTE:** When working with domains, you can use the `-listdom` option to view the Domain column in the command output or for a more global solution, set the `TPDLISTDOM` environment variable.

`-sortcol <col> [ ,<dir>][:<col>[ ,<dir>]...]`

Sorts command output based on the column number `<col>`. Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting `<dir>` as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

`-d`

Requests that more detailed layout information is displayed.

`-p`

Requests that policy information about the logical disk is displayed.

`-state`

Requests that the detailed information is displayed. This is the same as the `-s` option.

`-s`

Requests that the detailed state information is displayed.

## SPECIFIERS

`<LD_name>...`

Requests that information for a specified logical disk is displayed. This specifier can be repeated to display configuration information about multiple logical disks. If not specified, configuration information for all logical disks in the system is displayed.

`<pattern>...`

Specifies that the logical disk matching the specified glob-style pattern is displayed. This specifier can be repeated to display configuration information about multiple logical disks. If not specified, configuration information for all logical disks in the system is displayed. See [Glob-Style Pattern](#) on page 2.4 for more information.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays information for all LDs in a two node system:

```
cli% showld
Id Name          RAID -Detailed_State- Own SizeMB UsedMB Use  Lgct LgId WThru MapV
0 log0.0         1 normal      0/-  20480      0 log   0 ---   Y   N
1 log1.0         1 normal      1/-  20480      0 log   0 ---   Y   N
2 pdsld0.0       1 normal      0/1   8192      0 P,F   0 ---   Y   N
3 admin.usr.0    1 normal      0/1   5120     5120 V    0 ---   N   Y
4 admin.usr.1    1 normal      1/0   5120     5120 V    0 ---   N   Y
5 VV111.usr.0    0 normal      0/1    256      256 V    0 ---   N   Y
6 VV111.usr.1    0 normal      1/0    256      256 V    0 ---   N   Y
```

The following example displays information for all LDs in a four node system:

```
cli% showld
Id Name          RAID State Own      SizeMB UsedMB Use  Lgct LgId WThru MapV
0 log0.0         1 normal 0/-/-/- 20480      0 log   0 ---   Y   N
1 log1.0         1 normal 1/-/-/- 20480      0 log   0 ---   Y   N
2 pdsld0.0       1 normal 0/1/2/3  4096      0 P,F   0 ---   Y   N
3 admin.usr.0    1 normal 0/1/3/2  3584     3584 V    0 ---   N   Y
4 admin.usr.1    1 normal 0/1/2/3  1536     1536 V    0 ---   N   Y
5 admin.usr.2    1 normal 1/0/3/2  5120     5120 V    0 ---   N   Y
6 100_tune.usr.0 1 normal 0/1/2/3   512       512 V    0 ---   N   Y
```

The columns in the previous example are identified as follows:

- Id. The ID of the logical disk.
- Name. The name of the logical disk.
- Domain. The domain to which the logical disks belongs. If the domain does not exist, – is displayed.
- RAID. The RAID type (0 is for RAID-0, 1 is for RAID-1, 5 is for RAID-5, and 6 is for RAID-6).
- Detailed\_State. The current status of the logical disk.
  - ◆ normal. The logical disk has started and is available for use.
  - ◆ orphan. Both the primary owner and backup owner nodes are down, and the logical disk's data is not available.
  - ◆ preserved. Some disks used by the logical disk are missing. Data belonging to the logical disk is saved on the preserved logical disk.
  - ◆ stopped. The logical disk is stopped, and its data is not available.
  - ◆ stopping. The logical disk is being stopped; normally flushes any in-flight data to disk.



- ◆ removing. The logical disk is being deleted.

The following State values indicate that an uncontrolled shutdown has occurred and the logical disk was not properly closed before the shutdown.

- ◆ auto\_check. The logical disk is checked for validity.
  - ◆ checking. The logical disk is checked for validity.
  - ◆ need\_check. The logical disk has been checked, and an inconsistency has been found.
  - ◆ need\_pd. One of the physical disks containing chunklets that make up the logical disk is not yet in the valid state.
- Own. The first number is the logical disk owner node, and the second number is the logical disk backup node.
  - SizeMB. The total size of the logical disk, in MB.
  - UsedMB. The portion of the logical disk that is being used by virtual volumes, in MB.
  - Use. The function of the logical disk.
    - ◆ V. The logical disk is used for a virtual volume.
    - ◆ P , F. The first logical disk is used for preserved data.
    - ◆ P. The logical disk is used for preserved data.
    - ◆ C , SA. The logical disk is used for the snapshot administration space for a common provisioning group (CPG).
    - ◆ C , SD. The logical disk is used for the snapshot data space for a CPG.
    - ◆ log. The logical disk is used as a logging logical disk.
  - Lgct. The number of chunklets that are in logging mode in the logical disk.
  - LgId. The ID of the logging disk that is being used for logging by the logical disk.
  - WThru. Indicates whether the logical disk is in write-through mode.
  - MapV. Indicates if the logical disk is mapped to a virtual volume. Valid values are Y (yes) and N (no).

The following example displays detailed information about the system LDs:

```
cli% showld -d
Id Name          CPG RAID Own  SizeMB RSizeMB RowSz StepKB SetSz Refcnt Avail CAvail -----CreationTime----- -CreationPattern-
0 log0.0         --- 1 0/-  20480 40960   1   256   2     0 mag  mag  2007-11-29 18:44:02 PST -p -devtype FC
1 log1.0         --- 1 1/-  20480 40960   1   256   2     0 mag  mag  2007-11-29 18:44:02 PST -p -devtype FC
2 pdsld0.0       --- 1 0/1   4096   8192   2   256   2     0 mag  mag  2007-11-29 18:44:03 PST -p -devtype FC
3 admin.usr.0    --- 1 0/1   5120  10240   2   256   2     0 mag  mag  2007-11-29 18:44:12 PST -p -devtype FC
4 admin.usr.1    --- 1 1/0   5120  10240   2   256   2     0 mag  mag  2007-11-29 18:44:12 PST -p -devtype FC
5 Collie.usr.0   --- 0 1/0    256    256   1   256   1     0 ch   ch   2008-01-15 16:51:23 PST -p -devtype FC
-----
6                               55552 110848
```

The columns in the previous example are identified as follows:

- **Id.** The ID of the logical disk.
- **Name.** The name of the logical disk.
- **Domain.** The domain to which the logical disks belongs. If the domain does not exist, – is displayed.
- **CPG.** Lists the CPG to which the logical disk belongs. If the logical disk does not belong, --- is displayed.
- **RAID.** The RAID type (0 is for RAID-0, 1 is for RAID-1, 5 is for RAID-5, and 6 is for RAID-6).
- **Own.** The first number is the logical disk owner node, and the second number is the logical disk backup node.
- **SizeMB.** The total size of the logical disk, in MB.
- **RSizeMB.** The logical disk's raw size, in MB.
- **RowSz.** The row size for the logical disk.
- **StepKB.** The step size for the logical disk, in KB (1024 bytes). The number of contiguous bytes that the system accesses before moving to the next chunklet.
- **SetSz.** The set size for the logical disk.
- **Refcnt.** Indicates how many active relocation operations are occurring for the logical disk (reference count).
- **Avail.** Indicates the availability setting specified when the logical disk was created. Availability determines from where spare chunklets can be allocated when one of the logical disk's chunklets fails.
  - ◆ **disk.** A chunklet on the same disk as the failing chunklet can be used as a replacement.

- ◆ **mag.** A chunklet from another disk within the same drive magazine can be used as a replacement.
- ◆ **cage.** A chunklet from another drive cage can be used as a replacement.
- ◆ **port.** A chunklet from another port can be used as a replacement.
- **CAvail.** Indicates the current, or actual, availability: **disk**, **mag**, **port**, or **cage**.
- **Creation Time.** Indicates when the logical disk was created.

The following example displays policy information about all system LDs:

```
cli% showld -p
Id Name          Keep_Preserved Dev_Type K_RPM
0 log0.0          0 FC           10
1 log1.0          0 FC           10
2 pdsld0.0        0 FC           10
3 admin.usr.0     1 FC           10
4 admin.usr.1     1 FC           10
-----
5
```

The columns in the previous example are identified as follows:

- **Id.** The logical disk ID.
- **Name.** The logical disk name.
- **Keep\_Preserved.** Indicates if preserved data is saved for the logical disk. Valid values are 1 (saved) and 0 (not saved).
- **Dev\_Type.** Device type for the physical disk (FC for Fibre Channel or NL for Nearline or SSD for Solid State Device). If no device type is specified, the default is Fibre Channel drives. If a system has only Nearline drives, **-p -devtype NL** must be specified when creating the volumes.
- **K\_RPM.** Shows the device speed used by the majority of the chunklets of the logical disk. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

The following example displays detailed states of each LD:

```
cli% showld -s
Id Name                -State- -Detailed_State-
0 log0.0               normal normal
1 pdsld0.0             normal normal
2 pdsld0.1             normal normal
3 pdsld0.2             normal normal
4 admin.usr.0          normal normal
5 vvfromcpg2.usr.0     normal normal
6 vvfromcpg2.usr.1     normal normal
7 Temple.usr.0         normal normal
...
```

The columns in the previous example are identified as follows:

- Id. The logical disk ID.
- Name. The logical disk name.
- State. The current state of the LD that is one of the following:
  - ◆ normal. The LD is operating normally.
  - ◆ failed. The LD is operating abnormally.
  - ◆ degraded. The LD is in degraded state.
- Detailed\_State. The detailed state of the LD that is one of the following:
  - ◆ nodes\_unavail. The nodes are down.
  - ◆ not\_stated. The LD has stopped and needs to be manually started (see `startld` command).
  - ◆ unavail. The LD state not available.
  - ◆ preserved. One or more sets of the LD are not recoverable because some chunklets are missing. The remaining data from the LD has been preserved.
  - ◆ need\_check. Consistency checking needs to be manually started (see `checkld` command).
  - ◆ check\_failed. The LD check has failed.
  - ◆ need\_pd. The LD needs a PD before it can be started.
  - ◆ check\_ongoing. The LD is currently being checked for consistency.

- ◆ `degraded_avail`. The availability of the LD is degraded.
- ◆ `degraded_perf`. The performance of the LD is degraded.
- ◆ `unknown`. The LD state is unknown.

## NOTES

- For this command, KB=1024 bytes and MB=1048576 bytes.
- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.

---

**COMMAND**

showldch

**DESCRIPTION**

The `showldch` command displays configuration information about the chunklet mapping for one Logical Disk (LD).

**SYNTAX**

`showldch [options <arg>] <LD_name>`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-degraded`

Show only the chunklets in sets that cause the LD availability to be degraded. For example, if the LD normally has cage level availability, but one set has two chunklets in the same cage, then the chunklets in that set are shown. This option cannot be specified with option `-lformat` or `-linfo`.

`-lformat <form>`

Shows the LD's row and set layout on the Physical Disk (PD), where the line format `<form>` is one of:

- ◆ `row` - One line per LD row.
- ◆ `set` - One line per LD set.

`-linfo <info>[,<info>...]`

Specifies the information shown for each LD chunklet, where `<info>` can be one of:

- ◆ `pdpos` - Shows the PD position (default).
- ◆ `pdid` - Shows the PD ID.
- ◆ `pdch` - Shows the PD chunklet.

If multiple `<info>` fields are specified, each corresponding field will be shown separated by a dash (-).

**SPECIFIERS**

<LD\_name>

Specifies the LD name.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays information about chunklets for LD `r1.usr.0`:

```
cli% showldch r1.usr.0
```

Ldch	Row	Set	PdPos	Pdid	Pdch	State	Usage	Media	Sp	From	To
0	0	0	1:3:2	30	0	normal	ld valid	N		---	---
1	0	0	1:3:3	31	3	normal	ld valid	N	10:0	---	---
2	0	1	2:3:1	45	0	normal	ld valid	N		---	---
3	0	1	1:3:0	28	0	normal	ld valid	N		---	---
4	0	2	0:1:2	6	0	normal	ld valid	N		---	---
5	0	2	2:1:3	39	0	normal	ld valid	N		---	---
6	0	3	0:2:0	8	0	normal	ld valid	N		---	---
7	0	3	1:2:2	26	0	normal	ld valid	N		---	---
8	0	4	2:3:3	47	0	normal	ld valid	N		---	---
9	0	4	1:1:2	22	0	normal	ld valid	N		---	---
10	0	5	2:2:3	43	0	normal	ld valid	N		---	---
11	0	5	0:3:0	12	0	normal	ld valid	N		---	---
12	0	6	0:0:0	0	0	normal	ld valid	N		---	---
13	0	6	1:0:2	18	0	normal	ld valid	N		---	---
14	0	7	2:0:1	33	0	normal	ld valid	N		---	---
15	0	7	2:0:2	34	5	normal	ld valid	N	16:1	---	---

The columns in the previous example are identified as follows:

- **Ldch.** The chunklet on the logical disk.
- **Row.** The row that contains the chunklet.
- **Set.** The set that contains the chunklet.
- **PdPos.** The position of the disk in the drive magazine.
- **Pdid.** The ID of the physical disk that contains the chunklet.
- **Pdch.** The physical disk ID for the chunklet.
- **State.** The current status of the chunklet.

- ◆ `normal`. The chunklet is available for use.
- ◆ `normal,smag`. A servicemag operation is being performed on the disks.
- ◆ `stale`. The chunklet is not available for use because of a medium failure or a connectivity failure.
- ◆ `logging`. Writes to the chunklet are being logged to a separate logging logical disk.
- ◆ `playback`. Data is played back from the logging logical disks.
- ◆ `passthru`. The chunklet does not process physical disk errors.
- ◆ `preserved`. Any I/O to the chunklet is written to the preserved logical disks.
- ◆ `preserved playback`. Data is played back from the preserved logical disks.
- **Usage.** The current use of the chunklet.
  - ◆ `available`. The chunklet is not in use.
  - ◆ `ld`. The chunklet is in use by a logical disk.
  - ◆ `relsrsc`. The chunklet is the source of a logical disk relocation operation.
  - ◆ `reltgt`. The chunklet is the target of a logical disk relocation operation.
  - ◆ `synch`. The chunklet is both the source and the target of a logical disk relocation operation (synchronizing the chunklet).
  - ◆ `cmprel`. The system is completing the logical disk relocation operation.
  - ◆ `abtreel`. The system is canceling the logical disk relocation operation.
- **Media.** The current status of the physical disk medium for the chunklet.
  - ◆ `valid`. The chunklet is available for use.
  - ◆ `failed`. The medium has encountered errors and is not available for use.
- **Sp.** The spare status of the chunklet; Y indicates the chunklet is used for spare, N indicates the chunklet is not used as spare.
- **From.** The initial location of the chunklet before relocation.
- **To.** The destination location of the chunklet during relocation



The following example displays information about degraded chunklets for LD `r1.usr.0`.

```
cli% showldch -degraded r1.usr.0
Ldch Row Set PdPos Pdid Pdch State Usage Media Sp From To
  0  0  0 1:3:2 30  0 normal  ld valid N   --- ---
  1  0  0 1:3:3 31  3 normal  ld valid N 10:0 ---
 14  0  7 2:0:1 33  0 normal  ld valid N   --- ---
 15  0  7 2:0:2 34  5 normal  ld valid N 16:1 ---
cli% showldch -degraded r1.usr.1
No degraded sets in LD r1.usr.1
```

The following example displays row and set layout for chunklets for LD `test.usr.0`.

```
cli% showldch -lformat row test.usr.0
-----Set0----- -----Set1-----
row  Ch0  Ch1  Ch2  Ch3  Ch0  Ch1  Ch2  Ch3
  0 2:3:1 1:2:2 0:3:2 1:3:0 0:0:0 2:0:1 1:2:0 0:1:0
  1 2:3:3 1:0:2 0:0:2 2:0:3 0:2:2 1:1:0 2:1:1 1:3:0
  2 0:3:0 2:2:1 1:2:2 2:1:3 0:1:2 1:3:2 2:3:1 0:2:2
  3 0:2:0 1:2:0 2:1:3 1:0:0 2:2:3 0:0:0 1:1:2 0:1:0
  4 2:3:3 1:0:2 0:3:2 2:0:1 0:0:2 1:3:0 2:1:1 1:1:2
```

The following example displays row and set layout including `pdid` and `pdch` for chunklets for LD `test.usr.0`.

```
cli% showldch -lformat row -linfo pdid,pdch test.usr.0
-----Set0----- -----Set1-----
row  Ch0  Ch1  Ch2  Ch3  Ch0  Ch1  Ch2  Ch3
  0 45-7 26-6 14-7 28-6  0-7 33-7 24-6  4-6
  1 47-7 18-7  2-6 35-6 10-7 20-7 37-6 28-8
  2 12-6 41-6 26-8 39-7  6-8 30-8 45-9 10-9
  3  8-7 24-8 39-9 16-8 43-8  0-9 22-7  4-8
  4 47-9 18-9 14-9 33-9  2-8 28-9 37-8 22-9
```

The following example displays row and set layout including `pdid` and `pdpos` for chunklets for LD `test.usr.0`.

```
cli% showldch -lformat set -linfo pdid,pdpos test.usr.0
row set      Ch0      Ch1      Ch2      Ch3
  0   0 45-2:3:1 26-1:2:2 14-0:3:2 28-1:3:0
  0   1  0-0:0:0 33-2:0:1 24-1:2:0  4-0:1:0
  1   0 47-2:3:3 18-1:0:2  2-0:0:2 35-2:0:3
  1   1 10-0:2:2 20-1:1:0 37-2:1:1 28-1:3:0
  2   0 12-0:3:0 41-2:2:1 26-1:2:2 39-2:1:3
  2   1  6-0:1:2 30-1:3:2 45-2:3:1 10-0:2:2
  3   0  8-0:2:0 24-1:2:0 39-2:1:3 16-1:0:0
  3   1 43-2:2:3  0-0:0:0 22-1:1:2  4-0:1:0
  4   0 47-2:3:3 18-1:0:2 14-0:3:2 33-2:0:1
  4   1  2-0:0:2 28-1:3:0 37-2:1:1 22-1:1:2
```

## NOTES

None.

---

**COMMAND**

showldmap

**DESCRIPTION**

The `showldmap` command displays the mapping from a Logical Disk (LD) to Virtual Volumes (VVs).

**SYNTAX**

`showldmap <LD_name>`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

`<LD_name>`

Specifies the LD name.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the region of LD `v0.usr.0` that is used for a VV:

```
cli% showldmap v0.usr.0
Area Start(MB) Length(MB) VVId VVName          VVSp VVOff(MB)
0      0      512      0   v0              usr  0
```

The columns in the previous example are identified as follows:

- `Area`. The ID of the region.
- `Start(MB)`. The offset from the beginning of the logical disk, in MB.
- `Length(MB)`. The length, or size, of the region, in MB.
- `VVId`. The ID of the virtual volume that contains the region.

- `VVName`. The name of the virtual volume that contains the region.
- `VVSp`. The type of the space-user (`usr`), snapshot data (`snp`), or snapshot administration (`adm`)-for which the region is being used.
- `VVOff (MB)`. The offset from the beginning of the space that contains the region.

## NOTES

None.

---

**COMMAND**

`showlicense`

**DESCRIPTION**

The `showlicense` command displays the currently installed license key and current license information.

**SYNTAX**

`showlicense`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-raw`

Specifies that the license key originally entered (the raw license) be displayed. The license key is displayed in a manner that is acceptable input for the `setlicense` command.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays the currently installed license information:

```
cli% showlicense
License key was generated on Mon Sep 14 18:30:24 2009

License features currently enabled:
Dynamic Optimization
InForm Suite
Recovery Manager for Exchange
Recovery Manager for Oracle
System Reporter
Remote Copy
System Tuner
Thin Provisioning (1024G)
VSS Provider for Microsoft Windows

License features enabled on a trial basis:
Virtual Copy                               Expires on September 14, 2012
```

The following example displays the original license key entered on the system:

```
cli% showlicense -raw
60R3-0C1G-60R3-2C1G-60R3-0C9G-70R3-0C1G
60RK-0C0A-FSXZ-8YZ4-Z884-84DW-7CD6-JLKB
7GZA-RRZH-L01W-00AW-FFLR-T848-VWQA-K3F1
GXCJ-G8MG-0XW7-3VGL-EF28-MBDM-03V4-LG2D
YWQ5-KW9G-99GB-3FRL-4FJN-3AH2-T287-SNGD
XN9F-EVGV-25Z9-D2VK-5EXD-KXJ4-JMQE-6JSL
1XA0-9ZTA-H5KR-VG8Y-JBV8-BR72-A1E1-3GS3
ZSG4-32VZ-C2EQ-C69H-B0FR-0XG3-4NAF-G3VA
SV47-NH2R-BL1L-TJKR-31F2-V9QX-0WM0-AWTH
LRB7-8XQD-BLHD-LF68-A4BW-K4QG-ZS7B-Q41X
35Y1-6CZC-KBJ5-6VQE-EC2G-Q6EG-SJV0-KMHC
SAWG-YJ4V-WRE5-3GZA-DGW0-ZYVE-5SVT-8TJX
QR9W-R92D-SSTY-RVWY-1Y2B-YH8D-KRN5-ZZ7Y...
```

## NOTES

Without options, a description of which features are enabled is displayed. When the license was generated and any expiration dates are also shown. If the `-raw` option is specified, the license key prints in a manner that is acceptable input for the `setlicense` command.

---

**COMMAND**

shownet

**DESCRIPTION**

The `shownet` command displays the configuration and status of the administration network interfaces, including the configured gateway and Network Time Protocol (NTP) server.

**SYNTAX**

`shownet [option]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTION**

`-d`

Displays detailed information.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays the status of the system administration network interfaces:

```

root@enodec34:S424# shownet -d
IP Address: 192.168.5.191   Netmask 255.255.252.0
Assigned to nodes: 01
Connected through node 0
Status: Active

Admin interface on node 0
MAC Address:      00:02:AC:43:00:34
RX Packets:      534389   TX Packets:      121669
RX Bytes:        67828134 TX Bytes:        60638375
RX Errors:        0       TX Errors:        0
RX Dropped:      0       TX Dropped:      0
RX FIFO Errors:  0       TX FIFO Errors:  0
RX Frame Errors: 0       TX Collisions:   0
RX Multicast:    0       TX Carrier Errors: 0
RX Compressed:   0       TX Compressed:   0

Remote copy interface in slot 1 on node 0
MAC Address:      00:04:23:C2:1B:72
RX Packets:      0       TX Packets:      0
RX Bytes:        0       TX Bytes:        0
RX Errors:        0       TX Errors:        0
RX Dropped:      0       TX Dropped:      0
RX FIFO Errors:  0       TX FIFO Errors:  0
RX Frame Errors: 0       TX Collisions:   0
RX Multicast:    0       TX Carrier Errors: 0
RX Compressed:   0       TX Compressed:   0

Remote copy interface in slot 1 on node 0
MAC Address:      00:04:23:C2:1B:73
RX Packets:      0       TX Packets:      0
RX Bytes:        0       TX Bytes:        0
RX Errors:        0       TX Errors:        0
RX Dropped:      0       TX Dropped:      0
RX FIFO Errors:  0       TX FIFO Errors:  0
RX Frame Errors: 0       TX Collisions:   0
RX Multicast:    0       TX Carrier Errors: 0
RX Compressed:   0       TX Compressed:   0

Admin interface on node 1
MAC Address:      00:02:AC:43:00:33
RX Packets:      477147   TX Packets:      66399
RX Bytes:        59353196 TX Bytes:        33589695
RX Errors:        0       TX Errors:        0
RX Dropped:      0       TX Dropped:      0
RX FIFO Errors:  0       TX FIFO Errors:  0
RX Frame Errors: 0       TX Collisions:   0
RX Multicast:    0       TX Carrier Errors: 0
RX Compressed:   0       TX Compressed:   0
...

```

## NOTES

Specifying `-d` includes information useful for debugging network issues.



---

**COMMAND**

shownode

**DESCRIPTION**

The `shownode` command displays an overview of the node-specific properties and its component information. Various command options can be used to display the properties of PCI cards, CPUs, physical memory, disk drives, and power supplies.

**SYNTAX**

`shownode [option] [<node_ID>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

The following options are for node summary and inventory information:

`-i`

Shows node inventory information in the table format.

`-d`

Shows node and component information in the table format.

The following options are for node component information. These options cannot be used together with options `-i` and `-d`.

`-verbose`

Specifies detailed information in verbose format. It can be used with any of the following component options:

`-pci`

Displays PCI card information.

`-cpu`

Displays CPU information.

`-mem`

Displays physical memory information.

**-drive**

Displays the disk drive information.

**-ps**

Displays power supply information.

**-s**

Displays the detailed state information for node or power supply (proceed with the **-ps** option)

## SPECIFIERS

**<node\_ID>...**

Displays the node information for the specified node ID(s). This specifier is not required.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays the operating environment status for all nodes in the system:

```
cli% shownode
```

Node	--Name--	-State-	Master	InCluster	---LED---	Control Mem(MB)	Data Mem(MB)	Cache Available(%)
0	enodec34	OK	Yes	Yes	GreenBlk	2048	2048	100
1	enodec33	OK	No	Yes	GreenBlk	2048	2048	100

In the example above:

- Node. The node ID.
- Name. The node name.
- State. The state of the node. Values are as follows:
  - ◆ --. Cannot determine the overall node state.
  - ◆ OK. The node and its components are operating normally.
  - ◆ Degraded. For T-Series and S-Series nodes, the node is degraded due to a missing, failed, or degraded power supply. For E-series nodes, the node is degraded due to a missing or degraded fan.

- ◆ Failed. The node has not initialized, is offline, has mismatching kernel versions has bad drive partitions, is rebooting, or has shutdown.
- Master. Specifies is the node is the master node.
- InCluster. Indicates if the node is in the cluster.
- LED. The node LED information. Values are as follows:
  - ◆ --. The node LED is unknown.
  - ◆ off. The node LED is off.
  - ◆ Green. The kernal is not running.
  - ◆ GreenBlink. The node is in normal state.
  - ◆ Amber. The node is degraded or failed and the kernal is not running.
  - ◆ AmberBlink. The node is degraded or failed.
- Control Mem(MB). The total memory in the node in MB.
- Data Mem(MB). The total data memory in the node in MB.

The following examples display detailed information (`-d` option) for the nodes including their components in a table format. The `shownode -d` command can be used to display the tail information of the nodes including their components in name and value pairs.

```
cli% shownode -d
```

-----Physical Memory-----										
Node	Riser	Slot	SlotID	-Name--	-Usage-	--Type--	--Manufacturer--	-Serial-	-Latency-	Size(MB)
0	n/a	0	J4200	DIMM0	Control	FB-DIMM	Micron Technology	E20BCA01	CL5.0/5.0	2048
0	n/a	1	J4300	DIMM1	Control	FB-DIMM	Micron Technology	E20BCA0C	CL5.0/5.0	2048
0	2-slot	0	J0901	DIMM0.0	Data	DDR_SDRAM	Micron Technology	E010DF6E	CL2.0/2.5	2048
0	4-slot	1	J0900	DIMM1.0	Data	DDR_SDRAM	Micron Technology	E010DF67	CL2.0/2.5	2048
0	4-slot	2	J1300	DIMM2.0	Data	DDR_SDRAM	Micron Technology	E010DF63	CL2.0/2.5	2048
1	n/a	0	J4200	DIMM0	Control	FB-DIMM	Micron Technology	E20BC9AC	CL5.0/5.0	2048
1	n/a	1	J4300	DIMM1	Control	FB-DIMM	Micron Technology	E20BC9AE	CL5.0/5.0	2048
1	2-slot	0	J0901	DIMM0.0	Data	DDR_SDRAM	Micron Technology	E010DF95	CL2.0/2.5	2048
1	4-slot	1	J0900	DIMM1.0	Data	DDR_SDRAM	Micron Technology	E010DF97	CL2.0/2.5	2048
1	4-slot	2	J1300	DIMM2.0	Data	DDR_SDRAM	Micron Technology	E010DF96	CL2.0/2.5	2048

  

-----Internal Drives-----						
Node	Drive	-Manufacturer-	---Model---	-Serial-	-Firmware-	Size(MB) Type
0	0	Seagate	ST9100821AS	5NJ09DF3	3.AAB	95396 SATA
1	0	Seagate	ST9100821AS	5NJ08NA4	3.AAB	95396 SATA

These columns and values are described in detail in the following examples.

```
cli% shownode -mem
```

Node	Riser	Slot	SlotID	-Name--	-Usage-	--Type---	--Manufacturer---	-Serial-	-Latency-	Size(MB)
0	n/a	0	J4200	DIMM0	Control	FB-DIMM	Micron Technology	E20BC997	CL5.0/5.0	2048
0	n/a	1	J4300	DIMM1	Control	FB-DIMM	Micron Technology	E20BC996	CL5.0/5.0	2048
0	2-slot	0	J0901	DIMM0.0	Data	DDR_SDRAM	Micron Technology	E010DF65	CL2.0/2.5	2048
0	4-slot	1	J0900	DIMM1.0	Data	DDR_SDRAM	Micron Technology	E010DF98	CL2.0/2.5	2048
0	4-slot	2	J1300	DIMM2.0	Data	DDR_SDRAM	Micron Technology	E010DF64	CL2.0/2.5	2048
1	n/a	0	J4200	DIMM0	Control	FB-DIMM	Micron Technology	E20BC995	CL5.0/5.0	2048
1	n/a	1	J4300	DIMM1	Control	FB-DIMM	Micron Technology	E20BC992	CL5.0/5.0	2048
1	2-slot	0	J0901	DIMM0.0	Data	DDR_SDRAM	Micron Technology	E010DF6A	CL2.0/2.5	2048
1	4-slot	1	J0900	DIMM1.0	Data	DDR_SDRAM	Micron Technology	E010DF6F	CL2.0/2.5	2048
1	4-slot	2	J1300	DIMM2.0	Data	DDR_SDRAM	Micron Technology	E010DF69	CL2.0/2.5	2048

In the example above:

- Node. The node ID.
- Riser. The location of the DIMM. The DIMM may be located in a 2-slot or 4-slot riser card. If the DIMM is located on the motherboard the value is n/a.
- Slot. The slot number.
- SlotID. The slot ID in JXXX format.
- Name. The location of the DIMM slots printed on the board.
- Usage. The usage of the physical memory, either Control or Data.
- Type. The type of DIMM such as SDRAM, DDR, or DDR2.
- Manufacturer. The manufacturer of the memory.
- Serial. The serial number.
- Latency. The CAS latency.
- Size. The memory size in MB.

```
cli% shownode -drive
...
Node Drive -Manufacturer- ---Model--- -Serial- -Firmware- Size(MB) Type
  0      0 Seagate          ST9100821AS 5NJ09DF3 3.AAB      95396 SATA
  1      0 Seagate          ST9100821AS 5NJ08NA4 3.AAB      95396 SATA
```

In the example above:

- Node. The node ID.
- Drive. The disk drive ID.
- Manufacturer. The manufacturer of the disk drive.
- Model. The model number of the disk drive.
- Serial. The serial number of the disk drive.
- Firmware. The firmware version on the disk drive.
- Size. The drive size in MB.
- Type. The type of disk drive.

```
cli% shownode -pci

-----PCI Cards-----
Node Slot Type -Manufacturer- -Model- -Serial- -Rev- Firmware
  0      0 FC   QLOGIC          2302   D13503   1     3.3.16
  0      1 FC   QLOGIC          2302   D44171   1     3.3.16
  1      0 FC   QLOGIC          2302   C31037   1     3.3.16
```

In the example above:

- Node. The node ID.
- Slot. The slot ID.
- Type. The PCI card type, either Fibre Channel (FC) or Ethernet (Eth).
- Manufacturer. The PCI card manufacturer.
- Model. The model of the PCI card.
- Serial. The serial number of the PCI card.

- Rev. The revision number of the PCI card.
- Firmware. The firmware version on the PCI card.

```
cli% shownode -cpu

-----CPUs-----
Node CPU -Manufacturer- -----Serial----- CPUSpeed(MHz) BusSpeed(MHz)
    0   0 GenuineIntel   000188AAF2AEA667           995         132.76
    0   1 GenuineIntel   0000D8D47BD100A6           995         132.76
    1   0 GenuineIntel   0000F5A65AAD5D6F           995         132.72
    1   1 GenuineIntel   00028FE5F6E4068D           995         132.72
```

In the example above:

- Node. The node ID.
- CPU. The CPU ID.
- Manufacturer. The manufacturer name of the CPU.
- Serial. The serial number of the CPU.
- CPUSpeed. The speed of the CPU in MHz.
- BusSpeed. The bus speed of the CPU in MHz.

```
cli% shownode -ps

-----Power Supplies-----
Node PS -PSState-- FanState ACState DCState -BatState- ChrgLvl(%)
    0   0 NotPresent --      --      --      NotPresent    0
    0   1 OK         OK       OK       OK       OK         100
    1   0 OK         OK       OK       OK       OK         100
    1   1 NotPresent --      --      --      NotPresent    0
```

In the example above:

- Node. The node ID.
- PS. The power supply ID.
- PSState. The power supply state. Values are as follows:
  - ◆ --. Cannot determine the state.
  - ◆ OK. The power supply is operating normally.

- ◆ Failed. The power supply is operating abnormally.
- ◆ NotPresent. The power supply is missing.
- ◆ Degraded. There is a fan failure, or the battery state is not OK.
- FanState. The power supply fan state. Values are as follows:
  - ◆ --. Cannot determine the state.
  - ◆ OK. The fan is operating normally.
  - ◆ Failed. The fan is operating abnormally.
- ACState. The state of the AC power. Values are as follows:
  - ◆ --. Cannot determine the state.
  - ◆ OK. Normal operation.
  - ◆ Failed. Abnormal operation.
- DCState. The state of DC power. Values are as follows:
  - ◆ --. Cannot determine the state.
  - ◆ OK. Normal operation.
  - ◆ Failed. Abnormal operation.
- BatState. The battery state. Values are as follows:
  - ◆ --. Cannot determine the state.
  - ◆ OK. Normal operation.
  - ◆ Failed. Abnormal operation.
  - ◆ NotPresent. The battery is missing.
  - ◆ MaxLifeLow. The maximum battery life is low.
- ChrgLvl. The battery charge level in percentage.

The following example displays the detailed state of the nodes:

```
cli% shownode -s

Node -State- -Detailed_State-
  0 OK      OK
  1 OK      OK
  2 OK      OK
  3 OK      OK
```

In the example above:

- Node. The node ID.
- State. The state of the node. Values are as follows:
  - ◆ OK. The node and its components are operating normally.
  - ◆ Degraded. The node is degraded when the power supply is missing, failed, or degraded (fan failed, battery max life low, failed, expired, or not present).
  - ◆ Failed. The node is either not initialized, offline, kernel revision mismatched, disk partition bad, rebooting, or shut down.
- Detailed\_State. The detailed state of the node. Values are as follows:
  - ◆ tod\_bat\_fail. The time of day battery failed.
  - ◆ invalid\_bat\_config. The battery has an invalid configuration.
  - ◆ link\_error. The node has a link error.
  - ◆ uncorrectable\_mem\_error. The node has an uncorrectable memory error.
  - ◆ multi\_uncorrectable\_mem\_error. The node has multiple uncorrectable memory error.
  - ◆ correctable\_mem\_error. The node has a correctable memory error.
  - ◆ internal\_system\_error. The node has an internal system error.
  - ◆ hardware\_watchdog\_error. The node has a hardware watchdog error.
  - ◆ pci\_error. There is a PCI error in the node.
  - ◆ driver\_software\_error. The node has a driver software error.
  - ◆ cpu\_overheating. The node CPU overheating.



- ◆ `cpu_vrm_overheating`. The node CPU VRM overheating.
- ◆ `control_cache_dimm_overheating`. The node Control Cache DIMM is overheating.
- ◆ `node_offline_due_to_failure`. The node is offline due to failure.
- ◆ `node_shutdown_manually`. The node was shutdown manually.
- ◆ `unknown`. The node state is unknown.

The following example displays the detailed state of the power supplies:

```
cli% shownode -ps -s

Node PS -State- -Detailed_State-
  0  0 OK      OK
  1  1 OK      OK
  2  2 OK      OK
  3  3 OK      OK
```

In the example above:

- Node. The node ID.
- PS. The power supply ID.
- State. The power supply state. Values are as follows:
  - ◆ OK. The power supply is operating normally.
  - ◆ Failed. The power supply is operating abnormally.
  - ◆ NotPresent. The power supply is missing.
  - ◆ Degraded. The fan failure or battery state is not OK.
- Detailed\_State. The power supply detailed state. Values are as follows:
  - ◆ `invalid_bat_count`. The power supply has an invalid battery count.
  - ◆ `dc_failed`. The power supply DC failed.
  - ◆ `ac_failed`. The power supply AC failed.
  - ◆ `fan_failed`. The power supply fan failed.
  - ◆ `charger_overload`. The power supply charger overloaded.
  - ◆ `battery_not_present`. The power supply battery is not present.

- ◆ `invalid_firmware`. The power supply firmware is invalid.
- ◆ `oscillating_presence`. The power supply presence is oscillating.
- ◆ `oscillating_dc`. The power supply DC is oscillating.
- ◆ `oscillating_ac`. The power supply AC is oscillating.
- ◆ `oscillating_fan`. The power supply fan is oscillating.
- ◆ `oscillating_charger`. The power supply charger is oscillating.
- ◆ `oscillating_bat`. The power supply battery is oscillating.
- ◆ `NotPresent`. The power supply is not present.
- ◆ `unknown`. The power supply state is unknown.

## NOTES

None.

---

**COMMAND**

shownodeenv

**DESCRIPTION**

The `shownodeenv` command displays the node operating environment status, including voltages and temperatures.

**SYNTAX**

`shownodeenv [options <arg>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-n <node_ID>...`

Specifies the ID of the node whose environment status is displayed. Multiple node IDs can be specified as a series of integers separated by a space (1 2 3). If no option is used, then the environment status of all nodes is displayed.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays the operating environment status for all nodes in the system:

```
cli% shownodeenv
Node 0
-----
Measurement      Reading  Lo Limit  Hi Limit  Status
CPU0 1.32V:      1.31 V   1.28 V    1.36 V    Within Tolerance
CPU1 1.32V:      1.31 V   1.28 V    1.36 V    Within Tolerance
82563 1.20V:      1.22 V   1.13 V    1.26 V    Within Tolerance
31154 1.30V:      1.29 V   1.22 V    1.37 V    Within Tolerance
82563 1.90V:      1.92 V   1.79 V    2.00 V    Within Tolerance
      3.30V:      3.34 V   3.11 V    3.47 V    Within Tolerance
      PLX 3.30V:   3.32 V   3.12 V    3.47 V    Within Tolerance
      VCC 5.00V:   5.15 V   4.74 V    5.76 V    Within Tolerance
      V_PTT 1.20V: 1.20 V   1.13 V    1.26 V    Within Tolerance
      MCH 1.50V:   1.50 V   1.41 V    1.58 V    Within Tolerance
      ESB 1.50V:   1.51 V   1.41 V    1.58 V    Within Tolerance
      FBD 1.50V:   1.44 V   1.41 V    1.58 V    Within Tolerance
      FBD 1.80V:   1.80 V   1.69 V    1.89 V    Within Tolerance
      VTT_FBD 0.90V: 0.91 V   0.84 V    0.95 V    Within Tolerance
      ESB 1.20V:   1.21 V   1.13 V    1.26 V    Within Tolerance
      PLX 1.00V:   1.00 V   0.94 V    1.05 V    Within Tolerance
Osprey DDR 1.25V: 1.25 V   1.17 V    1.32 V    Within Tolerance
Osprey DDR 2.50V: 2.50 V   2.36 V    2.63 V    Within Tolerance
Osprey Lnk 1.87V: 1.85 V   1.76 V    1.97 V    Within Tolerance
      FPGA 2.50V: 2.50 V   2.36 V    2.63 V    Within Tolerance
      12.00V:      12.06 V   11.37 V    12.62 V    Within Tolerance
      MCH Temp:    38 C     0 C       85 C      Within Tolerance
      Board Temp:  25 C     0 C       70 C      Within Tolerance
      LM94 Temp:   23 C     0 C       65 C      Within Tolerance
      LM87 Temp:   23 C     0 C       65 C      Within Tolerance
      LM87 Ext Temp: 20 C    -10 C     65 C      Within Tolerance
      CPU0 Temp:   24 C     0 C       85 C      Within Tolerance
      CPU1 Temp:   25 C     0 C       85 C      Within Tolerance
      CPU0 VRM Temp: n/a      n/a      n/a      Within Tolerance
      CPU1 VRM Temp: n/a      n/a      n/a      Within Tolerance
      CPU DIMM0 Temp: 37 C     0 C      105 C     Within Tolerance
      CPU DIMM1 Temp: 38 C     0 C      105 C     Within Tolerance
Node 1
-----
Measurement      Reading  Lo Limit  Hi Limit  Status
...
```

## NOTES

None.

---

**COMMAND**

`showpatch`

**DESCRIPTION**

The `showpatch` command displays patches applied to a system.

**SYNTAX**

`showpatch [option <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-hist`

Specifies the history of all patches and updates applied to the system.

`-d <ID>`

Specifies the details on a specified patch ID.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The `showpatch` command is specific to each individual patch and typically displays these fields:

- `Patch ID`. Specifies the patch ID.
- `Release Version`. Specifies TPD or UI release affected by the patch.
- `Synopsis`. Specifies the purpose of patch.
- `Date`. Specifies the build date of patch.
- `Bugs fixed`. Specifies the bugs fixed.
- `Description`. Specifies a detailed description of the problem or fix.
- `Affected Packages`. Specifies the new packages being changed.

- **Obsoletes.** Specifies the patch IDs deleted by this patch.
- **Requires.** Specifies the patch IDs of any other patches required by this patch.
- **Notes.** Specifies any special instructions for the patch.

**NOTES**

This command displays all the patches currently affecting the system if options are not used.

---

**COMMAND**

`showpd`

**DESCRIPTION**

The `showpd` command displays configuration information about a system's Physical Disks (PDs).

**SYNTAX**

`showpd [options] [<PD_ID>...]`

`showpd -listcols`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-listcols`

List the columns available to be shown in the `-showcols` option described below (see `clihelp -col showpd` for help on each column).

The [options] are generally of two kinds: those that select the type of information that is displayed, and those that filter the list of PDs that are displayed.

By default (if none of the information selection options below are specified) the following columns are shown: Id, CagePos, Type, Speed\_KRPM, State, Size\_MB, Free\_MB, Port\_A0, Port\_B0, Port\_A1, Port\_B1.

Options that select the type of information shown include the following:

`-showcols <column>[,<column>...]`

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header. For other options that select the type of information, the column names may not be exactly the same and there may be spanning headers on top of the column names.

Run `showpd -listcol` to list the available columns.

Run `clihelp -col showpd` for a description of each column.

`-i`

Specifies a request for disk (inquiry) data.

The following columns are shown: Id, CagePos, State, Node\_WWN, Manuf, Model, Serial, FW\_Rev.

`-e`

Specifies a request for the disk environment and error information. Note that reading this information places a significant load on each disk.

The following columns are shown: Id, CagePos, Type, State, Rd\_CErr, Rd\_UErr, Wr\_CErr, Wr\_UErr.

`-c`

Specifies a request for chunklet use totals for each PD.

The following columns are shown: Id, CagePos, Type, State, Total\_Chunk, Nrm\_Used\_OK, Nrm\_Used\_Fail, Nrm\_Unused\_Free, Nrm\_Unused\_Uninit, Nrm\_Unused\_Unavail, Nrm\_Unused\_Fail, Spr\_Used\_OK, Spr\_Used\_Fail, Spr\_Unused\_Free, Spr\_Unused\_Uninit, Spr\_Unused\_Fail.

`-state`

Shows detailed information regarding the state of each PD.

If the `-old` option is not specified, the following columns are shown: Id, CagePos, Type, State, Detailed\_State.

`-path`

Shows current and saved path information for disks.

The following columns are shown: ID, CagePos, Type, State, Path\_A0, Path\_A1, Path\_B0, Path\_B1, Order.

Path\_A1 and Path\_B1 are only shown for systems with drive chassis that are connected to four nodes.

`-space`

Shows disk capacity usage information (MB).

The following columns are shown: Id, CagePos, Type, State, Size\_MB, Volume\_MB, Spare\_MB, Free\_MB, Unavail\_MB, Failed\_MB.



```
-sortcol <col> [ ,<dir>][:<col>[ ,<dir>]...]
```

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

```
inc
```

Sort in increasing order (default).

```
dec
```

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

The PD filtering options include:

```
-failed
```

Specifies that only failed physical disks are displayed.

```
-degraded
```

Specifies that only degraded PDs are displayed. If both `-failed` and `-degraded` are specified, the command shows failed disks and degraded disks.

```
-p <pattern>
```

PDs matching the specified pattern are displayed. The following arguments can be specified as patterns for this option:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

```
-nd <item>
```

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0-3). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0-3). The primary path of the disks must be on the specified PCI slot(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0-3). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0-3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-3). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (0, 1, 2). A list can also consist of a single integer (1). If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that only disks on the slots from the slot list are displayed. The slot list is specified as a series of integers separated by commas (0, 1, 2). A list can also consist of a single integer (1).

`-ports <port_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (0, 1, 2). A list can also consist of a single integer (1). If the port list is not specified, all disks on all ports are displayed.

`-w <world-wide_name>`

Specifies the World Wide Name (WWN) of the physical disk. This option and argument must be specified if the `<PD_ID>` specifier is not used and should be the last option in the command line.

## SPECIFIERS

`<PD_ID>`

Specifies a physical disk ID. This specifier must be used if the `-w` option is not specified.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays general information for all system physical disks:

```
cli% showpd
```

				---Size(MB)---		----Ports----	
ID	CagePos	Type	Speed(K)	State	Total	Free	A B
0	0:0:0	NL	7	normal	715008	695552	0:0:1* 1:0:1
1	0:3:0	FC	10	normal	69632	1024	0:0:1 1:0:1*
2	0:4:0	NL	7	normal	715008	695552	0:0:1 1:0:1*
3	0:7:0	FC	10	normal	69632	2560	0:0:1* 1:0:1
-----							
4					1569280	1394688	

The columns in the previous example are identified as follows:

- **ID.** Physical disk ID, as assigned by the system. For unadmitted disk the ID appears as dashes (--).
- **CagePos.** Position of PD, in the format <cage>:<side>.<mag>:<disk> where <cage> is the cage number, <side> is 0 or 1 depending on whether the cage is on the left or right (from the front view) of the disk chassis, <mag> is the magazine number (0 through 4) and <disk> is the disk number (0 through 3) within the magazine. If the position ends with (?) the current position of the drive is not known (for example, the drive may be missing), and the position shown is the last known position.
- **Type.** Device type of the PD.
  - ◆ FC - Fibre Channel
  - ◆ NL - Nearline
  - ◆ SSD - Solid State Drive
- **Speed(K).** The maximum RPM speed of the physical disk.
- **State.** State of the PD can be one of the following:
  - ◆ normal - PD is normal
  - ◆ degraded - The PD is not operating normally. Use `showpd -d` to find out the detail information.

- ◆ new - The PD is new, needs to be admitted before it can be used (see help admitpd)
- ◆ failed - The PD has failed.
- Total. Total size of the disk in MB (1024^2 bytes).
- Free. Free size of the disk in MB (1024^2 bytes). This includes the size that is available to use, uninitialized, and unavailable to use.
- Port\_A. The port location for the A port of the physical disk. If the primary path, the port location is followed by an asterisk (\*). If disabled, the port location is followed by a dash (-).
- Port\_B. The port location for the B port of the physical disk. If the primary path, the port location is followed by an asterisk (\*). If disabled, the port location is followed by a dash (-).

The following example displays only information for Nearline drives:

```
cli% showpd -c -p -devtype NL
```

				----- Normal Chunklets -----						- Spare Chunklets - -					
				- Used - -		----- Unused -----				- Used - -		- Unused - -			
ID	CagePos	Type	State	Total	OK	Fail	Free	Uninit	Unavail	Fail	OK	Fail	Free	Uninit	Fail
0	0:0:0	NL	normal	2793	76	0	2717	0	0	0	0	0	0	0	0
1	0:2:0	NL	normal	344	234	0	4	0	0	0	0	0	34	0	0
2	0:4:0	NL	normal	2793	90	0	2717	0	0	0	0	0	0	0	0
3	0:7:0	NL	normal	272	228	0	10	0	0	0	0	0	34	0	0
-----				-----											
4	total			6202	614		5448						68		

The columns in the previous example are identified as follows:

- ID. Physical disk ID, as assigned by the system. For unadmitted disk the ID appears as dashes (--).
- CagePos. Position of PD, in the format <cage>:<side>.<mag>:<disk> where <cage> is the cage number, <side> is 0 or 1 depending on whether the cage is on the left or right (from the front view) of the disk chassis, <mag> is the magazine number (0 through 4) and <disk> is the disk number (0 through 3) within the magazine. If the position ends with a question mark (?) the current position of the drive is not known (for example, the drive may be missing), and the position shown is the last known position.
- Type. Device type of the PD.
  - ◆ FC - Fibre Channel
  - ◆ NL - Nearline

- ◆ SSD - Solid State Drive
- State. State of the PD can be one of the following:
  - ◆ normal - PD is normal
  - ◆ degraded - The PD is not operating normally. Use `showpd -s` to find out the detail information.
  - ◆ new - The PD is new, needs to be admitted before it can be used (see `help admitpd`)
  - ◆ failed - The PD has failed.
- Total. Total size of the disk in MB ( $1024^2$  bytes).
- Normal Chunklets. This area provides information about chunklets that are not reserved for use as spares.
- Spare Chunklets. This area provides information about chunklets that are reserved for use as spares.
- OK. Number of chunklets with data that is accessible to the system.
- Fail. Number of failed chunklets.
- Free. Number of initialized but currently unused chunklets available for use by logical disks.
- Uninit. Number of chunklets being cleaned.
- Unavail. Number of chunklets that are unavailable for use.

The following example displays inquiry information for all disks:

```
# showpd -i
ID CagePos State ----Node_WWN---- --MFR-- ---Model--- -Serial- -FW_Rev--
0 0:0:0 normal 2210000A330013ED SEAGATE ST3750640NS 3QD0MM1B .AEV,1230
1 0:3:0 normal 2000000C501FD935 SEAGATE ST373207FC 3KT01JF1 XR36
2 0:4:0 normal 2210000A33000CBC SEAGATE ST3750640NS 3QD0T6AN .AEV,1230
3 0:7:0 normal 20000011C60B0AEB SEAGATE ST373207FC 3KT01QZF XR36
4 0:8:0 normal 2210000A33000CCA SEAGATE ST3750640NS 3QD0MTPK .AEV,1230
```

The columns in the previous example are identified as follows:

- ID. Physical disk ID, as assigned by the system. For unadmitted disk the ID appears as dashes (--).

- **CagePos.** Position of PD, in the format < cage> : < side> . < mag> : < disk> where < cage> is the cage number, < side> is 0 or 1 depending on whether the cage is on the left or right (from the front view) of the disk chassis, < mag> is the magazine number (0 through 4) and < disk> is the disk number (0 through 3) within the magazine. If the position ends with a question mark (?) the current position of the drive is not known (for example, the drive may be missing), and the position shown is the last known position.
- **State.** State of the PD can be one of the following:
  - ◆ normal - PD is normal
  - ◆ degraded - The PD is not operating normally. Use `showpd -s` to find out the detail information.
  - ◆ new - The PD is new, needs to be admitted before it can be used (see `help admitpd`)
  - ◆ failed - The PD has failed.
- **Node\_WWN.** Node World Wide Name of the disk.
- **MFR.** Physical disk manufacturer.
- **Model.** Physical disk manufacturer's device ID.
- **Serial.** Physical disk manufacturer's serial number.
- **FW\_Rev.** Physical disk manufacturer's firmware revision tracking string.

The following example displays chunklet use information for all disks:

cli% showpd -c

		----- Normal Chunklets -----								- Spare Chunklets --							
		- Used --				----- Unused -----				- Used - -				- Unused --			
ID	CagePos	Type	State	Total	OK	Fail	Free	Uninit	Unavail	Fail	OK	Fail	Free	Uninit	Fail		
0	0:0:0	NL	normal	2793	76	0	2717	0	0	0	0	0	0	0	0	0	0
1	0:3:0	FC	normal	272	234	0	4	0	0	0	0	0	34	0	0		
2	0:4:0	NL	normal	2793	76	0	2717	0	0	0	0	0	0	0	0	0	0
3	0:7:0	FC	normal	272	228	0	10	0	0	0	0	0	34	0	0		
-----																	
4	total			6130	614		5448						68				

The columns in the previous example are identified as follows:

- **ID.** Physical disk ID, as assigned by the system. For unadmitted disk the ID appears as dashes (--).



- **CagePos.** Position of PD, in the format < cage> : < side> . < mag> : < disk> where < cage> is the cage number, < side> is 0 or 1 depending on whether the cage is on the left or right (from the front view) of the disk chassis, < mag> is the magazine number (0 through 4) and < disk> is the disk number (0 through 3) within the magazine. If the position ends with a question mark (?) the current position of the drive is not known (for example, the drive may be missing), and the position shown is the last known position.
- **Type.** Device type of the PD.
  - ◆ FC - Fibre Channel
  - ◆ NL - Nearline
  - ◆ SSD - Solid State Drive
- **State.** State of the PD can be one of the following:
  - ◆ normal - PD is normal
  - ◆ degraded - The PD is not operating normally. Use `showpd -s` to find out the detail information.
  - ◆ new - The PD is new, needs to be admitted before it can be used (see `help admitpd`)
  - ◆ failed - The PD has failed.
- **Total.** Total size of the disk in MB (1024^2 bytes).
- **Normal Chunklets.** This area provides information about chunklets that are not reserved for use as spares.
- **Spare Chunklets.** This area provides information about chunklets that are reserved for use as spares.
- **OK.** Number of chunklets with data that is accessible to the system.
- **Fail.** Number of failed chunklets.
- **Free.** Number of initialized but currently unused chunklets available for use by logical disks.
- **Uninit.** Number of chunklets being cleaned.
- **Unavail.** Number of chunklets that are unavailable for use.

The following example displays environment and error information for all disks:

```
cli% showpd -e
```

ID	CagePos	Type	State	-Read Error-		-Write Error-		T(C)
				Corr	UnCorr	Corr	UnCorr	
0	0:0:0	FC	normal	0	0	0	0	33
1	0:0:1	FC	normal	0	0	0	0	34
2	0:0:2	FC	normal	0	0	0	0	30
3	0:0:3	FC	normal	0	0	0	0	25
4	0:1:0	FC	normal	0	0	0	0	29
5	0:1:1	FC	normal	0	0	0	0	30
6	0:1:2	FC	normal	0	0	0	0	28
7	0:1:3	FC	normal	0	0	0	0	24
8	1:2:0	FC	normal	0	0	0	0	32
9	1:2:1	FC	normal	0	0	0	0	32
10	1:2:2	FC	normal	0	0	0	0	30
11	1:2:3	FC	normal	0	0	0	0	25
12	1:3:0	FC	normal	0	0	0	0	29
13	1:3:1	FC	normal	0	0	0	0	29
14	1:3:2	FC	normal	0	0	0	0	28
15	1:3:3	FC	normal	0	0	0	0	24

The columns in the previous example are identified as follows:

- **Id.** Physical disk ID, as assigned by the system. For unadmitted disk the ID appears as dashes (--).
- **CagePos.** Position of PD, in the format < cage> : < side> . < mag> : < disk> where < cage> is the cage number, < side> is 0 or 1 depending on whether the cage is on the left or right (from the front view) of the disk chassis, < mag> is the magazine number (0 through 4) and < disk> is the disk number (0 through 3) within the magazine. If the position ends with (?) the current position of the drive is not known (for example, the drive may be missing), and the position shown is the last known position.
- **Type.** Device type of the PD.
  - ◆ FC - Fibre Channel
  - ◆ NL - Nearline
  - ◆ SSD - Solid State Drive
- **Read\_Error\_Corr.** Number of correctable read errors.
- **Read\_Error\_UnCorr.** Number of uncorrectable read errors.
- **Write\_Error\_Corr.** Number of correctable write errors.

- Write\_Error\_UnCorr. Number of uncorrectable write errors.
- T(C). Temperature in degrees celsius.

The following example displays the states of each physical disk:

```
cli% showpd -s
ID CagePos  Type -State- -Detailed_State-
 0 0:0:0    FC   normal  normal
 1 0:0:1    FC   normal  normal
 2 0:0:2    FC   normal  normal
 3 0:0:3    FC   normal  normal
 4 0:1:0    FC   normal  normal
 5 0:1:1    FC   normal  normal
 6 0:1:2    FC   normal  normal
 7 0:1:3    FC   normal  normal
 8 1:2:0    FC   normal  normal
 9 1:2:1    FC   normal  normal
10 1:2:2    FC   normal  normal
11 1:2:3    FC   normal  normal
12 1:3:0    FC   normal  normal
13 1:3:1    FC   normal  normal
14 1:3:2    FC   normal  normal
15 1:3:3    FC   degraded not_available_for_allocations
```

The following example displays the capacity usage totals for all disks:

```
cli% showpd -space
```

----- (MB) -----									
ID	CagePos	Type	-State-	Size	Volume	Spare	Free	Unavail	Failed
0	0:0:0	FC	normal	17152	10496	1280	5376	0	0
1	0:0:1	FC	normal	17152	6400	1280	9472	0	0
2	0:0:2	FC	normal	17152	10496	1280	5376	0	0
3	0:0:3	FC	normal	17152	6400	1024	9728	0	0
4	0:0:0	FC	normal	17152	10752	1024	5376	0	0
5	0:0:1	FC	normal	17152	6400	1024	9728	0	0
6	0:0:2	FC	normal	17152	10496	1024	5632	0	0
7	0:0:3	FC	normal	17152	6400	1024	9728	0	0
8	1:0:0	FC	normal	17152	10496	1024	5632	0	0
9	1:1:1	FC	normal	17152	6400	1024	9728	0	0
10	1:1:2	FC	normal	17152	10496	1024	5632	0	0
11	1:1:3	FC	normal	17152	6400	1024	9728	0	0
12	1:1:0	FC	normal	17152	10496	1024	5632	0	0
13	1:1:1	FC	normal	17152	6400	1024	9728	0	0
14	1:1:2	FC	normal	17152	10496	1024	5632	0	0
15	1:1:3	FC	normal	17152	6400	1024	9728	0	0
-----									
16	total			274432	135424	17152	121856	0	0

The following example displays disk information for the nearline disks on cage 0 and 3 only:

```
cli% showpd -p -cg 0,1 -devtype NL
```

---Size(MB)--- ----Ports----									
ID	CagePos	Type	Speed(K)	State	Total	Free	A	B	
0	0:0:0	NL	7	normal	715008	684288	0:0:1*	1:0:1	
2	0:4:0	NL	7	normal	715008	684288	0:0:1	1:0:1*	
4	0:8:0	NL	7	normal	715008	684288	0:0:1*	1:0:1	
6	0:12:0	NL	7	normal	715008	684288	0:0:1	1:0:1*	
-----									
4	total				2860032	2737152			

The columns in the previous example are identified as follows:

- ID. Physical disk ID, as assigned by the system. For unadmitted disk the ID appears as dashes (--).
- CagePos. Position of PD, in the format <cage>:<side>.<mag>:<disk> where <cage> is the cage number, <side> is 0 or 1 depending on whether the cage is on the left or right (from the front view) of the disk chassis, <mag> is the magazine number (0 through 4) and <disk> is the disk number (0 through 3) within the magazine. If the position ends with (?)

the current position of the drive is not known (for example, the drive may be missing), and the position shown is the last known position.

- **Type.** Device type of the PD.
  - ◆ FC - Fibre Channel
  - ◆ NL - Nearline
  - ◆ SSD - Solid State Drive
- **Speed(K).** The maximum RPM speed of the physical disk. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.
- **State.** State of the PD can be one of the following:
  - ◆ normal - PD is normal
  - ◆ degraded - The PD is not operating normally. Use `showpd -s` to find out the detail information.
  - ◆ new - The PD is new, needs to be admitted before it can be used (see `help admitpd`)
  - ◆ failed - The PD has failed.
- **Total.** Total size of the disk in MB (1024^2 bytes).
- **Free.** Free size of the disk in MB (1024^2 bytes). This includes the size that is available to use, uninitialized, and unavailable to use.
- **Port\_A.** The port location for the A port of the physical disk. If the primary path, the port location is followed by an asterisk (\*). If disabled, the port location is followed by a dash (-).
- **Port\_B.** The port location for the B port of the physical disk. If the primary path, the port location is followed by an asterisk (\*). If disabled, the port location is followed by a dash (-).

The following example displays disk information for the nearline disks on cage 0 and sort the output by the cage position:

```
cli% showpd -sortcol 1 -p -cg 0 -devtype NL
                                ---Size(MB)---  ---Ports---
ID CagePos Type Speed(K) State      Total    Free A      B
  0 0:0:0  NL          7 degraded  715008   684288 0:0:1* 1:0:1
  2 0:4:0  NL          7 degraded  715008   684288 0:0:1  1:0:1*
  4 0:8:0  NL          7 degraded  715008   684288 0:0:1* 1:0:1
  6 0:12:0 NL          7 degraded  715008   684288 0:0:1  1:0:1*
-----
  4 total                                2860032 2737152
```

The following example displays all disk in magazine 0 of cage 0 and 1:

```
cli% showpd -p -mg 0 -cg 0,1
                                --Size(MB)--  ---Ports---
ID CagePos Type Speed(K) State      Total    Free A      B
  0 0:0:0  NL          7 degraded  715008   684288 0:0:1* 1:0:1
  8 1:0:0  FC         10 normal    69632    2560 0:0:2* 1:0:2
-----
  2 total                                784640 686848
```

## NOTES

- In the `showpd` output, when the position of the disk is not valid (for example, in the case of a missing disk), the most recent position might be displayed, followed by a question mark (?).
- To see the device type for a PD (Fibre Channel, Nearline, or Solid State Drive), use the `showpd` command.

- Without the `-i`, `-c`, `-e`, `-s`, `-path` or `-space` options, basic information about the PD is printed. Note that the primary path to the disk is shown by an asterisk(\*) in either the `APort` or `BPort` column, and a minus (-) sign following a path indicates the port is disabled.
- If the `<PD_ID>` specifier and the `-w` option is not specified, all disks are displayed.

---

**COMMAND**

showpdata

**DESCRIPTION**

The showpdata command displays information about the preserved data in the system.

**SYNTAX**

showpdata [ <LD\_name> ]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

[ <LD\_name> ]

Requests that preserved data for a specific Logical Disk (LD) is displayed. This specifier is not required on the command line. If not specified, then the amount of preserved data by LD is displayed.

**RESTRICTIONS**

Access to all domains is required to run this command if no LD name is specified. If an LD name is specified, then access to the domain of the LD is required.

**EXAMPLES**

The following example displays preserved data information for LD admin.usr.3:

```
cli% showpdata admin.usr.3
showpdata v0.usr.0
Preserved LD raid sets
      Ldname      Set      PreservedCnt
      admin.usr.3    0          1
No preserved chunklets
```

**NOTES**

None.



---

**COMMAND**

showpdch

**DESCRIPTION**

The `showpdch` command displays the status of selected Physical Disk (PD) chunklets.

**SYNTAX**

`showpdch [options]`

`showpdch [options] <pdid>|<pattern>... (deprecated usage)`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

The following filtering options specify conditions that a chunklet should meet in order for the chunklet to be displayed. Multiple conditions can be specified and a chunklet will be displayed if any of the specified conditions are met. By default (if no filtering options are specified) only those chunklets that are mapped to an LD are displayed.

`-a`

Specifies that information about all chunklets is displayed.

`-fail`

Specifies that information about failed chunklets (media-failed chunklets, disk-failed chunklets, or chunklets marked as failed by the operating system) is displayed.

`-mov`

Specifies that information about chunklets that have moved, are scheduled to move, or are moving, is displayed.

`-from <pdid,...>`

Specifies chunklets that have moved or are to be moved from the matching PD IDs.

`-cln`

Specifies that information for clean chunklets is displayed.

`-cng`

Specifies that information for chunklets that are being cleaned by the system, is displayed.

`-tgt`

Specifies that information for chunklets marked as targets of relocation is displayed.

`-src`

Specifies that information about chunklets that are marked as sources of relocation, is displayed.

`-spr`

Specifies that information about chunklets that are marked as spares, is displayed.

`-log`

Specifies that information about chunklets that are logging is displayed.

`-sync`

Specifies that information about chunklets that are synchronizing with their RAID sets is displayed.

`-sortcol <col> [ ,<dir>][:<col>[ ,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

`-p <pattern>`

Specifies a pattern for disks. Patterns are used to select disks that are used for creating LDs. If no pattern is specified, the option defaults to all Fibre Channel (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks matching the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL) and FC drives.



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0–7). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0–7). The primary path of the disks must be on the specified PCI slot(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–4). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the `CagePos` column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0–7). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0–3). The specified drive position(s) must contain disks.

`-dk <item>`

Specifies one or more PDs. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1,2,3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that PDs with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that PDs with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that PDs with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that PDs with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that PDs identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that PDs must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

`<PD_ID> | <pattern>...` (deprecated usage)

Specifies the PD IDs or the PD ID pattern for disks whose chunklets are to be displayed.

## RESTRICTIONS

If no option is specified, the `showpdch` command defaults to displaying information about chunklets used for logical disks (`-ld`).

## EXAMPLES

The following example displays information about chunklets on PD 1:

```
cli% showpdch 1
Pdid Chnk          LdName LdCh  State Usage Media Sp Cl From To
  1    0    vvfromcpg2.usr.1    0 normal   ld valid  N  N  --- ---
  1    1    vvfromcpg2.usr.1    4 normal   ld valid  N  N  --- ---
  1    2    vvfromcpg2.usr.1    8 normal   ld valid  N  N  --- ---
-----
Total chunklets: 3
```

The columns in the previous example are identified as follows:

- **Pdid.** The physical disk on which the chunklets reside.
- **Chnk.** The chunklet number.
- **LdName.** The name of the logical disk that is using the spare chunklet.

- **LdCh.** The position of the chunklet on the logical disk.
- **State.** The state of the chunklet as identified by the kernel.
  - ◆ **logging.** I/O to the chunklet is written to the logging logical disk.
  - ◆ **playback.** Data is played back from the logging logical disks.
  - ◆ **passthru.** Chunklets do not process physical disk errors.
  - ◆ **preserved.** Any I/O to the chunklet is written to the preserved logical disks.
  - ◆ **preserved playback.** Data is played back from the preserved logical disks.
  - ◆ **stale.** The chunklet is not available for use because of a medium failure or a connectivity failure.
  - ◆ **normal.** The chunklet is available for use.
  - ◆ **normal ,smag.** A servicemag operation is performed on the disks.
  - ◆ **none.** Chunklets were not used by any logical disk.
- **Usage.** Shows whether the spare chunklet is in use by a logical disk.
  - ◆ **available.** The chunklet is available for use as a spare or as a logical disk.
  - ◆ **ld.** The chunklet is in use by a logical disk.
  - ◆ **synch.** The chunklet is both the source and the target of a logical disk relocation operation (synchronizing the chunklet).
  - ◆ **cmprcl.** The system is completing the logical disk relocation operation.
  - ◆ **reldsrc.** Relocation source. The data has been moved to another chunklet.
  - ◆ **reltgt.** Relocation target. The data in the chunklet has been moved from another spare chunklet.
  - ◆ **abtrcl.** Abort relocation. The system is canceling the logical disk relocation operation.
- **Media.** The current status of the physical disk medium for the chunklet.
  - ◆ **valid.** The chunklet is available for use.
  - ◆ **failed.** The medium has encountered errors and is not available for use.
- **Sp.** The spare status of the chunklet; **Y** indicates the chunklet is reserved for spare, **N** indicates a previously free chunklet selected by the system as a spare.

- C1. The clean status of the chunklet. N indicates that the chunklet is in-use. Y indicates that the chunklet is clean. Cg indicates that the chunklet is being cleaned.
- From. The initial location of the chunklet before relocation in the syntax <PD\_ID> : <chunklet\_number>.
- To. The destination location of the chunklet during relocation in the syntax <PD\_ID> : <chunklet\_number>.

## NOTES

None.

---

## COMMAND

showpdvv

## DESCRIPTION

The `showpdvv` command displays the virtual volumes that are mapped to a particular physical disk.

## SYNTAX

```
showpdvv [options] [<PD_ID> [:<chunklet>]...]
```

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

`-sum`

Shows number of chunklets used by virtual volumes for different space types for each physical disk.

`-p <pattern>`

PDs matching the specified pattern are displayed. The following arguments can be specified as patterns for this option:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (`item`). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0-3). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (`item`). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0-3). The primary path of the disks must be on the specified PCI slot(s).



`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0-3). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0-3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the *CagePos* column of *showpd* output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-3). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

`-sortcol <col>[, <dir>][:<col>[, <dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

inc

Sort in increasing order (default).

dec

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

## SPECIFIERS

<PD\_ID>

Specifies the physical disk ID using an integer. This specifier is not required if `-p` option is used, otherwise it must be used at least once on the command line.

[ :<chunklet> ]

Specifies the chunklet number in a physical disk to which virtual volumes are mapped. This specifier is not required. This specifier is not required. This specifier cannot be used along with `-p` option.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays that the user space virtual volumes `v0` and `v1`, respectively, are mapped to PD 44:

```
cli% showpdvv 44
PDId CagePos Type Speed(K) VVId VVName VVSp
  44 2:3:2   FC         10  109 v0      usr
  44 2:3:2   FC         10  110 v1      usr
```

The following example displays the summary output for volumes mapped to physical disk 55:

```
cli% showpdvv -sum 55
                                     ----Chunklets----
PDId CagePos Type Speed(K) VVId VVName Adm Snp Usr Total
  55 4:9:0   FC         15    2 v0      0  0  2    2
  55 4:9:0   FC         15    3 v1      0  0 20   20
```

## NOTES

None.

---

**COMMAND**

showport

**DESCRIPTION**

The `showport` command displays information about ports in the system.

**SYNTAX**

The syntax of the `showport` command can be one of the following:

■ `showport [options] [-failed] [-sortcol <col>[,<dir>]]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-i

Show hardware and inventory formation.

-c

Displays all devices connected to each loop (by position on the loop). Position 0 is the cage connected directly to the node on this loop.

Position n is the (n + 1) cage on the loop. Thus a "pos" value of 1 means this cage is the second cage away from the node on the loop. Example:

Node --> 0th cage --> 1st cage --> 2nd cage --> 3rd cage

-par

Displays a parameter listing such as the configured data rate of a port and the maximum data rate that the card supports. Also shown is the type of attachment (Direct Connect or Fabric Attached) and whether the unique\_nwwn and VCN capabilities are enabled.

-rc

Displays information that is specific to the Remote Copy ports.

-rcfc

Displays information that is specific to the Fibre Channel Remote Copy ports.

`-rcip`

Displays information specific to the Ethernet Remote Copy ports.

`-iscsi`

Displays information about iSCSI ports.

`-iscsiname`

Displays iSCSI names associated with iSCSI ports.

`-sfp`

Displays information about hot pluggable SFPs attached to ports.

`-ddm`

Displays information about the SFPs DDM. This option must be used with the `-sfp` option.

`-d`

Displays detailed information about the SFPs attached to ports. This option is used with the `-sfp` option.

`-failed`

Displays only failed ports.

`-state`

Displays the detailed state information.

`-s`

Displays the detailed state information. This option has been deprecated and will be removed in a future release.

`-sortcol <col>[, <dir>][:<col>[, <dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting `<dir>` can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by the values in the later columns.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays information about all ports in the system:

```
cli% showport
N:S:P      Mode      State  ----Node_WWN---- -Port_WWN/HW_Addr-  Type
0:0:1 initiator   ready  2FF70002AC00000A   20010002AC00000A   disk
0:0:2 initiator   ready  2FF70002AC00000A   20020002AC00000A   disk
0:0:3 initiator   ready  2FF70002AC00000A   20030002AC00000A   disk
0:0:4 initiator   loss_sync  2FF70002AC00000A   20040002AC00000A   free
0:2:1 initiator   ready  2FF70002AC00000A   20210002AC00000A   host
0:2:2 target      ready  2FF70002AC00000A   20220002AC00000A   host
0:4:1 peer        offline -          0002B3A99F7E      rcip
0:5:1 initiator   loss_sync  2FF70002AC00000A   20510002AC00000A   free
0:5:2 initiator   loss_sync  2FF70002AC00000A   20520002AC00000A   free
1:0:1 initiator   ready  2FF70002AC00000A   21010002AC00000A   disk
1:0:2 initiator   ready  2FF70002AC00000A   21020002AC00000A   disk
1:0:3 initiator   ready  2FF70002AC00000A   21030002AC00000A   free
1:0:4 initiator   loss_sync  2FF70002AC00000A   21040002AC00000A   free
1:2:1 initiator   ready  2FF70002AC00000A   21210002AC00000A   host
1:2:2 initiator   ready  2FF70002AC00000A   21220002AC00000A   disk
1:3:1 target      ready  -              00C0DD077977      iscsi
1:3:2 target      loss_sync -              00C0DD077979      iscsi
1:4:1 peer        offline -          0002B3BCE77D      rcip
1:5:1 target      ready  2FF70002AC00000A   21510002AC00000A   host
1:5:2 target      loss_sync  2FF70002AC00000A   21520002AC00000A   free
```

The columns in the previous example are identified as follows:

- **N:S:P.** The physical position of the port, in the syntax node:slot:port.
- **Mode.** Indicates whether the port is an initiator or a target. Initiators connect to disks (default) and target ports connect to hosts or fabrics. A target port that has yet to be initialized by the system appears as suspended. Ethernet ports use peer mode.
- **State.** State of the port. Possible values vary according to whether the port is an Ethernet or Fibre Channel port.
  - ◆ **ready.** The port is online and ready for use.

- ◆ `loss_sync`. The port is not physically connected to anything.
- ◆ `config_wait`. Firmware has yet to be initialized.
- ◆ `login_wait`. Fibre Channel adapter is attempting port and process logins with all loop ports.
- ◆ `error`. Fibre Channel adapter has experienced an unrecoverable error.
- ◆ `non_participate`. Port is logically isolated from the Fibre Channel loop.
- ◆ `offline`. The port is offline.
- ◆ `pending_reset`. A reset of the port has been scheduled, but not yet completed.
- `Node_WWN`. The WWN that belongs to the controller node identified in the N:S:P column.
- `Port_WWN/HW_Addr`. The WWN that belongs to the controller node port identified in the N:S:P column.
- `Type`. Indicates the port connection type.
  - ◆ `host`. Port is connected to hosts. This port can also be used for Remote Copy over IP (RCIP).
  - ◆ `disk`. Port is connected to disks.
  - ◆ `free`. Port is not connected.

The following example displays hardware and connection information about the ports' Fibre Channel settings:

```
cli% showport -i
N:S:P Brand Model Rev Firmware Serial
1:2:1 3PAR FC044X 08 1.02.N.5 00981760001d6187
1:2:2 3PAR FC044X 08 1.02.N.5 00981760001d6187
1:2:3 3PAR FC044X 08 1.02.N.5 00981760001d6187
1:2:4 3PAR FC044X 08 1.02.N.5 00981760001d6187
```



The following example displays all devices connected to each port and information about the ports' Fibre Channel settings:

```
cli% showport -c
N:S:P      Mode      Device Pos Config      Topology Rate Cls Mode_change
1:2:1 initiator cage4  0  valid private_loop 2Gbps  3    allowed
1:2:2 initiator cage5  0  valid private_loop 2Gbps  3    allowed
1:2:3 initiator cage7  0  valid private_loop 2Gbps  3    allowed
1:2:4 initiator cage6  0  valid private_loop 2Gbps  3    allowed
```

The columns in the previous example are identified as follows:

- **N:S:P.** The physical position of the port, in the syntax node:slot:port.
- **Mode.** Indicates whether the port is an initiator or a target. Initiators connect to disks (default) and target ports connect to hosts or fabrics. A target port that has yet to be initialized by the system appears as suspended.
- **Device.** The device connected to the port, for example cage0.
- **Pos.** The position of the device, 0 or 1.
- **Config.** The validity of the port configuration. Values are valid or invalid.
- **Topology.** The link topology. Valid values are private\_loop, public\_loop, point\_to\_point, and fabric.
- **Rate.** Data transfer rate between the port and a host or a disk. Values can be 1Gbps, 2Gbps, 4Gbps, or n/a.
- **Cls.** The Fibre Channel class of the port., either 2 or 3.
- **Mode\_change.** Indicates whether a port mode change is allowed or prohibited.

The following example displays each system ports' configuration:

```
cli% showport -par
N:S:P Connmode ConnType CfgRate MaxRate Class2  UniqNodeWwn VCN      IntCoal
0:0:1 disk      loop    auto    2Gbps   disabled n/a        enabled enabled
0:0:2 disk      loop    auto    2Gbps   disabled n/a        enabled enabled
0:3:1 host      loop    auto    2Gbps   disabled n/a        enabled enabled
0:3:2 host      loop    auto    2Gbps   disabled n/a        enabled enabled
1:0:1 disk      loop    auto    2Gbps   disabled n/a        enabled enabled
1:0:2 disk      loop    auto    2Gbps   disabled n/a        enabled enabled
1:3:1 host      loop    auto    2Gbps   disabled n/a        enabled enabled
1:3:2 host      loop    auto    2Gbps   disabled n/a        enabled enabled
-----
8
```

The columns in the previous example are identified as follows:

- N:S:P. The physical position of the port, in the syntax node:slot:port.
- ConnType. Indicates the type of connection configured for the port; loop for loop, point for point-to-point, or lp for loop or point-to-point.
- CfgRate. The configurable bit rate of data transfer between the port and a host or a disk; either auto, 1Gbps, 2Gbps, or 4Gbps.
- MaxRate. Indicates the maximum bit rate of transfer between the HBA and the host or disk.
- Class2. Identifies whether class 2 is enabled (ack0 or ack1) or disabled.
- VCN. VLUN State Change Notification support setting. Valid values are n/a, enabled or disabled. When enabled and in public loop or fabric topology, a Registered State Change Notification (RSCN) message is issued to the fabric controller whenever a VLUN is created or removed.
- Persona. Indicates the port persona setting. An asterisk(\*) to the left indicates the persona cannot be trusted because one or more underlying attributes have been changed.
- IntCoal. Indicates whether the interrupt coalesce setting is Enabled or Disabled.

The following example displays information about all Remote Copy ports:

```
cli% showport -rc
N:S:P State      HwAddr      Rate Type
5:3:2 ready 25320002AC000006 2Gbps rcip
4:1:1 ready 24110002AC000006 2Gbps rcip
0:5:1 ready 0002B39B2013 100Mbps rcip
```

The columns in the previous example are identified as follows:

- **N:S:P.** The physical position of the port, in the syntax `node:slot:port`.
- **State.** State of the port.
  - ◆ **ready.** The port is online and ready for use.
  - ◆ **loss\_sync.** The port is not physically connected to anything.
  - ◆ **config\_wait.** Firmware has yet to be initialized.
  - ◆ **login\_wait.** Fibre Channel adapter is attempting port and process logins with all loop ports.
  - ◆ **error.** Fibre Channel adapter has experienced an unrecoverable error.
  - ◆ **non\_participate.** Port is logically isolated from the Fibre Channel loop.
  - ◆ **offline.** The port is offline.
- **HwAddr.** A unique identifier of the port hardware used for Remote Copy connection. For an RCIP port, it is the MAC address of the port.
- **Rate.** Data transfer rate (bitrate) for the Remote Copy interface.
- **Type.** Indicates the port connection type.
  - ◆ **rcip.** Port is used for Remote Copy over IP (RCIP).

The following example displays the detailed state of the port:

```
cli% showport -s
N:S:P --State-- -Detailed_State-
0:0:1 ready      loss_sync
0:0:2 ready      loss_sync
0:3:1 ready      offline
0:3:1 ready      offline
0:5:1 ready      loss_sync
-----
10
```

The columns in the previous example are identified as follows:

- N:S:P. The physical position of the port, in the syntax node:slot:port.
- State. State of the port.
  - ◆ OK. Port SFP is operating normally.
  - ◆ Degraded. Port SFP is not operating normally.
- Detailed State.
  - ◆ config\_wait. The link configuration is waiting.
  - ◆ alpa\_wait. The link ALPA is waiting.
  - ◆ login\_wait. The link login is waiting.
  - ◆ ready. The link is online and ready for use.
  - ◆ loss\_sync. The link is not physically connected to anything.
  - ◆ error. The link has an error.
  - ◆ non\_participate. The link is not participating.
  - ◆ taking\_coredump. The link is a taking a coredump.
  - ◆ offline. The link is offline.
  - ◆ fw\_dead. The links firmware is dead.
  - ◆ link\_idle\_for\_reset. The link is idle and ready for reset.
  - ◆ dhcp\_in\_progress. The link DHCP is in progress.
  - ◆ pending\_reset. The link reset is pending.

- ◆ unknown. The link state is unknown.

The following example displays information about RCIP ports:

```
cli% showport -rcip
N:S:P   State ---HwAddr--- IPAddr Netmask Gateway MTU Rate Duplex AutoNeg
0:1:1   offline 000423C21B72    -      -      -      - n/a   n/a   n/a
0:1:1   offline 000423C21B73    -      -      -      - n/a   n/a   n/a
1:2:1   offline 000423ADE95E    -      -      -      - n/a   n/a   n/a
1:2:1   offline 000423ADE95F    -      -      -      - n/a   n/a   n/a
```

The columns in the previous example are identified as follows:

- N:S:P. The physical position of the port, in the syntax node:slot:port.
- State. State of the port.
  - ◆ ready. The port is online and ready for use.
  - ◆ loss\_sync. The port is not physically connected to anything.
  - ◆ config\_wait. Firmware has yet to be initialized.
  - ◆ login\_wait. Fibre Channel adapter is attempting port and process logins with all loop ports.
  - ◆ error. Fibre Channel adapter has experienced an unrecoverable error.
  - ◆ non\_participate. Port is logically isolated from the Fibre Channel loop.
  - ◆ offline. The port is offline.
- HwAddr. A unique identifier of the port hardware used for Remote Copy connection. For an RCIP port, it is the MAC address of the port.
- IPAddr. The IP address of the Remote Copy interface.
- Netmask. Netmask for the Ethernet port.
- Gateway. Gateway address for the Remote Copy interface.
- MTU. Maximum Transfer Unit (MTU) size for the specified Remote Copy interface (default is 1500). The largest supported value is 9000 and the smallest is 100.
- Rate. Data transfer rate for the Remote Copy interface.
- Duplex. Values can be either Full or Half.

- AutoNeg. Values can either be Yes or No.

The following example displays information about iSCSI ports:

```
cli% showport -iscsi
N:S:P State      IPAddr      Netmask      Gateway TPGT  MTU  Rate  DHCP  iSNS_Prim iSNS_Sec iSNS_Port
1:3:1 ready      192.168.9.163 255.255.255.0 192.168.9.1 131 1500 1Gbps  0    0.0.0.0 0.0.0.0 3205
1:3:2 loss_sync  0.0.0.0      0.0.0.0      0.0.0.0 132 1500  n/a   0    0.0.0.0 0.0.0.0 3205
```

The following example displays information about iSCSI names associated with iSCSI ports:

```
cli% showport -iscsiname
N:S:P IPAddr -----iSCSI_Name-----
1:3:1 0.0.0.0 iqn.2000-05.com.3pardata:21310002ac00000a
1:3:2 0.0.0.0 iqn.2000-05.com.3pardata:21320002ac00000a
```

The following examples display both standard and detailed information about SFPs attached to ports:

```
cli% showport -sfp
N:S:P -State- -Manufacturer- MaxSpeed(Gbps) TXDisable TXFault RXLoss DDM
0:0:1 OK      FINISAR_CORP.      2.10 No      No      No      Yes
0:0:2 OK      FINISAR_CORP.      2.10 No      No      Yes     Yes
0:3:2 OK      SIGMA-LINKS        2.10 No      No      Yes     Yes
1:0:1 OK      FINISAR_CORP.      2.10 No      No      No      Yes
1:0:2 OK      FINISAR_CORP.      2.10 No      No      Yes     Yes
cli%

root@enodea5:~# showport -sfp -d
-----Port 0:0:2-----
N:S:P      : 0:0:2
State      : Degraded
Manufacturer : PICOLIGHT
Part Number : PL-XPL-VE-S24-31
Serial Number : 425EF1E6
Revision   : N/A
MaxSpeed(Gbps): 2.10
Qualified  : No
TX Disable : --
TX Fault   : --
RX Loss    : --
RX Power Low : No
DDM Support : No

-----Port 1:3:1-----
N:S:P      : 1:3:1
State      : OK
Manufacturer : FINISAR_CORP.
Part Number : FTLF8519P2BNL
Serial Number : U76031S
Revision   : A
MaxSpeed(Gbps): 2.10
Qualified  : Yes
TX Disable : No
TX Fault   : No
RX Loss    : No
RX Power Low : No
DDM Support : Yes
```

## NOTES

See [Restrictions](#) on page 10.16 for important information regarding port modes and port pair protection.

**COMMAND**

showportarp

**DESCRIPTION**

The `showportarp` command shows the ARP table for iSCSI ports in the system.

**SYNTAX**

`showportarp [ <N>:<S>:<P> ]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

`[ <N>:<S>:<P> ]`

Specifies the port for which information about devices on that port are displayed.

node

Specifies the node.

slot

Specifies the PCI bus slot in the specified node.

port

Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot.

If `<N>:<S>:<P>` is not specified, the ARP table for all iSCSI ports is displayed.

**EXAMPLES**

The following example displays the ARP table for the iSCSI ports in the system.

```
cli% showportarp
N:S:P      HwAddr      IPAddr  Time(mins)
1:3:1  00E07BF8BF87    192.168.9.1      0
1:3:1  001143CD039A    192.168.8.151    0
```

In the example output above:



- **HwAddr.** The MAC address of a remote host discovered through the MAC address resolution process.
- **IPAddr.** The IP address of the remote host to which the port is attempting to connect.
- **Time (mins).** The amount of time (in minutes) that the entry has been in the table. When the entry has been in the table for 20 minutes, it is removed.

## NOTES

None.

---

**COMMAND**

showportdev

**DESCRIPTION**

The `showportdev` command displays detailed information about either all devices or only arbitrated loop devices on a Fibre Channel port.

**SYNTAX**

showportdev loop|all|ns <N:S:P>

**AUTHORITY**

Super, Service, Edit, Browse

**SUBCOMMANDS**

loop|all|ns

loop

Specifies that information is returned for arbitrated loop devices that are attached to the specified port. If this subcommand is not specified, then the `all` subcommand must be specified on the command line.

all

Specifies that information for all devices attached to the specified port is returned. If this subcommand is not specified, then the `loop` subcommand must be specified on the command line.

ns

Specifies that information for the switch name server database is returned. This subcommand is only for use with fabric-attached topologies.

**OPTIONS**

None.

**SPECIFIERS**

<N:S:P>

Specifies the port for which information about devices on that port are displayed.

node

Specifies the node.

slot

Specifies the PCI bus slot in the specified node.

port

Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot.

## RESTRICTIONS

The `loop` subcommand is functional only in a private loop topology.

## EXAMPLES

The following example displays information about all devices attached to Fibre Channel port 1:0:2:

```
cli% showportdev all 1:0:2
```

PtId	LpID	Hadr	Node_WWN	Port_WWN	ftrs	svpm	bbct	flen
0xd3	0x0c	0x00	2FF70002AC000013	21020002AC000013	0x8800	0x0022	n/a	0x0800
0xef	0x00	0xef	2000000087002078	2200000087002078	0x8800	0x0012	n/a	0x0800
0xe8	0x01	0xe8	2000000087002515	2200000087002515	0x8800	0x0012	n/a	0x0800
0xe4	0x02	0xe4	20000000870024CB	22000000870024CB	0x8800	0x0012	n/a	0x0800
0xe2	0x03	0xe2	20000000870028AE	22000000870028AE	0x8800	0x0012	n/a	0x0800
0xe1	0x04	0xe1	2000000087002224	2200000087002224	0x8800	0x0012	n/a	0x0800
0xe0	0x05	0xe0	2000000087003019	2200000087003019	0x8800	0x0012	n/a	0x0800
0xdc	0x06	0xdc	5000087000190E9F	5000087000190EA1	0x8800	0x0012	n/a	0x0800
0xda	0x07	0xda	2000000087002397	2200000087002397	0x8800	0x0012	n/a	0x0800

The columns in the previous example are identified as follows:

- PtId. The ID of the port.
- LpID. The ID of the loop.
- Hadr. The hard address.
- Node\_WWN. The WWN of the node.
- Port\_WWN. The WWN of the port.
- ftrs. Common features that are located in PLOGI\_ACC common word 1, bits 31-16.
- svpm. The service parameters that are located in PRLI word 3, bits 15-0.
- bbct. The buffer to buffer credit that is located in PLOGI common word 0, bits 15-0.
- flen. The maximum received frame length that is located in PLOGI\_ACC class-3 word 1, bits 15-0.

## NOTES

None.

**COMMAND**

showportisns

**DESCRIPTION**

The `showportisns` command shows iSNS host information for iSCSI ports in the system.

**SYNTAX**

`showportisns [<N>:<S>:<P>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

`[<N>:<S>:<P>]`

Specifies the port for which information about devices on that port are displayed.

`node`

Specifies the node.

`slot`

Specifies the PCI bus slot in the specified node.

`port`

Specifies the iSCSI port number of the PCI card in the specified PCI bus slot.

If not specified, iSNS host information for all iSCSI ports is displayed.

**EXAMPLES**

The following example displays hosts discovered by the iSCSI port from the iSNS server in the system.

```
cli% showportisns
N:S:P  Host_IPAddr  -----Host_iSCSI_Name-----      Host_alias
1:3:1  192.168.2.181          ign.1991-05.com.microsoft:pe750-07 <MS SW iSCSI Initiator>
1:3:1  192.168.17.33  ign.1991-05.com.microsoft:dt-ashok-xp.hq.3pardata.com <MS SW iSCSI Initiator>
```

Where

- `N:S:P`. Represents the node:slot:port of the iSCSI port.
- `Host_IPAddr`. The IP address of a remote host.
- `Host_iSCSI_Name`. Represents the iSCSI name of the host.
- `Host_alias`. Represents the iSCSI alias of the host.

**NOTES**

None.

---

**COMMAND**

`showportlesb`

**DESCRIPTION**

The `showportlesb` command displays Fibre Channel Link Error Status Block (LESB) counters (the number of errors accumulated for Fibre Channel devices). The LESB is composed of six counters that can measure Fibre Channel signal integrity or status.

**SYNTAX**

The syntax for the `showportlesb` command can include one of the following arguments:

- `showportlesb reset`
- `showportlesb compare [all|<N:S:P>]`
- `showportlesb single|both <N:S:P>`
- `showportlesb hist [options <arg>] <N:S:P>`
- `showportlesb diffhist [options <arg>] <N:S:P>`

**AUTHORITY**

Super, Service, Edit, Browse

**SUBCOMMANDS**

`reset`

Specifies that internal 3PAR counters are checked against current LESB counters and event alerts are raised as necessary. All ports of the internal 3PAR counters are reset.

`compare`

Specifies that internal 3PAR counters are checked against current LESB counters.

`single|both`

The `single` subcommand specifies that counters for the indicated port are displayed. The `both` subcommand specifies that counters for both ports, if the device is dual ported, are displayed.

`hist`

Displays the history of LESB raw counters on the specified port of loop devices.

`diffhist`

Displays differences between historical samples of LESB counters on the specified port of loop devices.

## OPTIONS

`-startt <time>`

Specifies that samples taken of LESB counters should commence after the indicated time (<time>). Time can be specified in hours or as a specific date. When specifying the time in hours, the following formats can be used:

- ◆ `hh[:mm[:ss]]`, where `hh` is the hour, `mm` is the minute (optional), and `ss` is the second (optional).
- ◆ `hhmm`, where `hh` is interpreted as a 24 hour clock.

When specifying the time as a date, the following formats can be used:

- ◆ `mm/dd[/yy]`, where `mm` is the month, `dd` is the day, and `yy` is the year (optional).
- ◆ `monthname dd[,yy]`, where `dd` is the day and `yy` is the year (optional).
- ◆ `dd monthname [yy]`, where `dd` is the day and `yy` is the year (optional).
- ◆ `yy-mm-dd`, where `yy` is the year, `mm` is the month, and `dd` is the day.

`-endt <time>`

Specifies that samples taken of LESB counters cease after the indicated time (<time>). Time can be specified in hours or as a specific date. When specifying in hours, the following formats can be used:

- ◆ `hh[:mm[:ss]]`, where `hh` is the hour, `mm` is the minute (optional), and `ss` is the second (optional).
- ◆ `hhmm`, where `hh` is interpreted as a 24 hour clock.

When specifying as a date, the following formats can be used:

- ◆ `mm/dd[/yy]`, where `mm` is the month, `dd` is the day, and `yy` is the year (optional).
- ◆ `monthname dd[,yy]`, where `dd` is the day and `yy` is the year (optional).
- ◆ `dd monthname [yy]`, where `dd` is the day and `yy` is the year (optional).
- ◆ `yy-mm-dd`, where `yy` is the year, `mm` is the month, and `dd` is the day.



## SPECIFIERS

<N:S:P>

Specifies the port for which information about devices on that port are displayed. This specifier is required for the `single`, `both`, `hist`, and `diffhist` subcommands, and optional for the `compare` subcommand. If this specifier is not used with the `compare` subcommand, then all ports are compared. The port is specified as follows:

`node`

Specifies the node.

`slot`

Specifies the PCI bus slot in the specified node.

`port`

Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays the reset of internal counters and ports:

```
cli% showportlesb reset
LESB reset completed
```

The following example displays the comparison of all ports:

```
cli% showportlesb compare
Port <1:0:1>

Loop <1:0:1>      Time since last save: 0:00:28
  ID ALPA LinkFail LossSync LossSig PrimSeq InvWord InvCRC
<1:0:1> 0xef      1      19      19      0      0      0
pd7    0x6d      1       5       0      0     270     0
pd6    0x72      1       4       0      0     524     0
pd5    0x73      1       4       0      0     335     0
pd4    0x76      1       4       0      0     334     0
pd3    0x79      1       4       0      0     401     0
pd2    0x80      1       4       0      0     344     0
pd1    0x81      1       4       0      0     270     0
pd0    0x88      1       4       0      0     401     0
Port <0:0:1>

Loop <0:0:1>      Time since last save: 0:00:28
  ID ALPA LinkFail LossSync LossSig PrimSeq InvWord InvCRC
<0:0:1> 0xef      1      19      19      0      0      0
pd7    0x6d      1       5       0      0     465     0
pd6    0x72      1       5       0      0     890     0
pd5    0x73      1       4       0      0     969     0
pd4    0x76      1       5       0      0     761     0
pd3    0x79      1       4       0      0     815     0
pd2    0x80      1       4       0      0     925     0
pd1    0x81      1       7       0      0    3283     0
pd0    0x88      1     258       0      0     269     0

LESB compare completed
```

The following example displays the counters for port 1:0:2:

```
cli% showportlesb single 1:0:2
```

	ID	ALPA	LinkFail	LossSync	LossSig	PrimSeq	InvWord	InvCRC
cage1		0x1	3	4	0	0	755	0
pd12		0xef	1245	39201	0	0	156804	0
pd13		0xe8	1	1608	0	0	6432	0
pd14		0xe4	1	1586	0	0	6344	0
pd15		0xe2	1	1588	0	0	6352	0
pd16		0xe1	2	5088	0	0	20352	0
pd17		0xe0	1	1596	0	0	6384	0
pd18		0xdc	1	1595	0	0	6380	0
pd19		0xda	1	1596	0	0	6384	0
pd20		0xd9	2	5047	0	0	20188	0
pd21		0xd6	1	1604	0	0	6416	0
pd22		0xd5	1	1609	0	0	6436	0
pd23		0xd4	1	1616	0	0	6464	0
<1:0:2>		0xd3	0	1	1	0	0	0
cage0		0x18	1	165	0	0	26	0
pd8		0x67	2	865794	0	0	3463176	0
pd9		0x66	1	63932	0	0	255728	0
pd10		0x65	1	61572	0	0	246288	0
pd11		0x63	1	61525	0	0	246100	0
pd4		0x6c	2	67006	0	0	268024	0
pd5		0x6b	1	63474	0	0	253896	0
pd6		0x6a	1	63471	0	0	253884	0
pd7		0x69	1	63598	0	0	254392	0
pd0		0x72	2	65863	0	0	263452	0
pd1		0x71	1	64024	0	0	256096	0
pd2		0x6e	1	63942	0	0	255768	0
pd3		0x6d	1	63897	0	0	255588	0

The columns in the previous example are identified as follows:

- ID. The device ID.
- ALPA. Arbitrated Loop Physical Address.
- LinkFail. The Fibre Channel loop either has a loss of signal (electrical or optical) or a loss of synchronization that is greater than the timeout period.
- LossSync. Fibre Channel data is not valid though there is a signal (electrical or optical).
- LossSig. Loss of signal (electrical or optical) to the receiver port of a Fibre Channel node.

- **PrimSeq.** Primitive Sequence Protocol Error. There were errors during the transmission of a Fibre Channel primitive sequence. This might indicate an error during the loop recovery or initialization.
- **InvWord.** Invalid Transmission Word. Illegal Fibre Channel transmission word received.
- **InvCRC.** Invalid Cyclical Redundancy Check. Data corruption in the Fibre Channel frame.

## NOTES

- If the `both` subcommand is specified, the resulting output text is greater than 80 columns wide.
- Internal 3PAR counters are checked every 10 minutes.
- Each LESB counter is a 32-bit, unsigned integer.
- LESB counters on Fibre Channel devices cannot be reset.

---

## COMMAND

showrcopy

## DESCRIPTION

The `showrcopy` command displays details of the Remote Copy configuration.

## SYNTAX

The syntax for the `showrcopy` command can be one of the following:

- `showrcopy [options <arg>] [links]`
- `showrcopy [options <arg>] [groups [<name_or_pattern>]]`
- `showrcopy [options <arg>] [targets [<name_or_pattern>]]`

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

`-d`

Displays more detailed configuration information.

`-domain <domainname_or_pattern>[,<domainname_or_pattern>...]`

Shows only Remote Copy groups whose virtual volumes are in domains with names matching one or more of the specified domain name or pattern.

## SPECIFIERS

`links`

Specifies all Remote Copy links.

`groups [<name_or_pattern>]`

Specifies either all Remote Copy volume groups or a specific Remote Copy volume group(s) by name or glob-style patterns.

`targets [<name_or_pattern>]`

Specifies either all target definitions or a specific target(s) definition by name or glob-style.

RESTRICTIONS

This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

EXAMPLES

The following example displays output from the `showrcopy` command:

```
cli% showrcopy
Remote Copy System Information
Status: Started, Normal

Target Information

Name      ID Type Status Options Policy
InServ2 9  IP  ready          mirror_config

Link Information

Target  Node Address      Status Options
InServ2 0    10.100.33.11 Up
InServ2 1    10.101.33.11 Up
receive 0    10.100.33.11 Up
receive 1    10.101.33.11 Up

Group Information

Name          Target      Status      Role      Mode      Options
sync_group_1 InServ2      Started     Primary   Sync
  LocalVV      ID RemoteVV    ID SyncStatus LastSyncTime
  localvv.0    391 remotevv.0  351 Syncing (25%) Thu Dec 14 17:37:40 PST 2006
  localvv.1    392 remotevv.1  352 Syncing (40%) Thu Dec 14 17:37:40 PST 2006

Name          Target      Status      Role      Mode      Options
sync_group_2.r11 InServ2      Started     Secondary Sync
  LocalVV      ID RemoteVV    ID SyncStatus LastSyncTime
  remotevv.0    401 localvv.0    361 Syncing      NA
  remotevv.1    402 localvv.1    362 Syncing      NA
```

The following values can appear in the `Link Information Status` field:

- Not Started. Link is not started or is being started, such as when its node is down or Remote Copy is stopped (through the `stoprcopy` command).
- Down. Link is down and will attempt to restart.
- Up. Link is up and running.

The following values can appear in the `Group Information Syncstatus` field:

- New. Volume is configured, but has not yet been started.
- Syncing. Volume is currently synchronizing.
- NotSynced. The volume is not synchronized, likely the result of an initial sync failure or some other failure.
- Stopped. The volume was synced the last time the group was started, but the group is currently stopped. There might be writes that have not been sent to the secondary site.
- Stale. Volume was previously synchronized, but a previous synchronization attempt failed. Thus, the secondary has a valid copy, just not a valid copy from the last synchronization attempt.

The following example displays output from the `showrcopy groups <pat>` command, where `<pat>` is specified as `b*` and `l*`:

```
cli% showrcopy groups b* l*

Remote Copy System Information
Status: Started, Normal

Group Information

Name          Target      Status    Role      Mode      Options
bart          bf_mirror  Started   Primary   Periodic   Last-Sync Tue Jul 18 14:12:59
PDT 2006 , Period 30m
  LocalVV      ID      RemoteVV  ID      SyncStatus  LastSyncTime
  bart-tp-ws   9      bart-ws   3924    Synced      Tue Jul 18 14:13:09 PDT 2006

Name          Target      Status    Role      Mode      Options
lisa          bf_mirror  Started   Primary   Periodic   Last-Sync Tue Jul 18 14:09:44
PDT 2006 , Period 30m, over_per_alert
  LocalVV      ID      RemoteVV  ID      SyncStatus  LastSyncTime
  lisa-ws      2      lisa-ws   3922    Synced      Tue Jul 18 14:10:32 PDT 2006

cli%
```

In the example above:

- Name. The name of the group.
- Target. The target to which the group is mirrored.
- Status. The following values can appear in the group `Status` field:
  - ◆ New. Group that has not yet been started.

- ◆ Starting. Currently attempting to start the group.
- ◆ Started. Group is started (has Remote Copy running).
- ◆ Stopped. Group was stopped.
- Role. The role of the group as either Primary or Secondary.
- Mode. The mode of the group as either Periodic or Synchronous.
- Options. The options set for the group.
- LocalVV and ID. The name and ID of this system.
- RemoteVV and ID. The name and ID of the target system.
- SyncStatus. The following values can appear in the SyncStatus field:
  - ◆ New. Volume is configured as a primary volume, but has not yet been started.
  - ◆ Remote. Volume is configured as a secondary volume, but has not yet been started.
  - ◆ Syncing. Volume is currently synchronizing.
  - ◆ Synced. The primary and secondary volumes are in sync.
  - ◆ NotSynced. The volume is not synchronized, likely the result of an initial sync failure or some other failure.
  - ◆ Stopped. The volumes were previously synchronized, but may be out of sync due to a group being stopped.
  - ◆ Stale. Volume was previously synchronized, but a previous synchronization attempt failed. Thus the secondary has a valid copy, just not a valid copy from the last synchronization attempt.
- LastSyncTime. The time at which the last volume synchronization was completed.

## NOTES

- If the `showrcopy` command is used with no specifiers, all configuration information is displayed.
- The `showrcopy` command also displays the group's domain names if the global `-listdom` option is used or if the `TPDLISTDOM` environment variable is set.
- If `showrcopy links|groups|targets` is used without specifying a name, information for all links, groups, or targets is displayed.



- See Help on `sub,globpat` for information on glob-style patterns.
- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `-listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable set. See Help on `sub,globpat` for information on glob-style patterns. Please run `cli -h` and `setclienv -h` for details of the environment variables.

**COMMAND**

showrcrtransport

**DESCRIPTION**

The `showrcrtransport` command shows status and information about end-to-end transport for Remote Copy in the system.

**SYNTAX**

`showrcrtransport [options]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-rcip`

Show information about Ethernet end-to-end transport.

`-rcfc`

Show information about Fibre Channel end-to-end-transport.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays status and information about the end-to-end transport for all Remote Copy configured ports:

```
cli% showrcrtransport
N:S:P Peer_Address      Address State Type
0:5:1      - 10.100.5.87  new rcip
1:5:1      - 10.101.6.87  new rcip
```

The columns in the previous example are identified as follows:

- `N:S:P`. The physical position of the port, in the syntax `node:slot:port`.

- **Peer\_Address.** IP address of the Remote Copy storage server port connected to this Remote Copy port.
- **Address.** IP address of the Remote Copy storage server port.
- **State.** Remote copy end-to-end transport state. Can include one of the following results:
  - ◆ **new.** Configuration is not completed on this port.
  - ◆ **incomplete.** Configuration is not yet completed on the peer port.
  - ◆ **ready.** Configuration is completed on this port and the peer port; transport is ready for use.
  - ◆ **missing.** A configured transport was disconnected.
- **Type.** Indicates the port connection type.
  - ◆ **rcip.** Port is used for Remote Copy over IP (RCIP).
  - ◆ **rcfc.** Port is used for Remote Copy over FC (RCFC).

The following example displays status and information about the end-to-end transport for RCIP ports:

```
cli% showretransport -rcip
```

N:S:P	State	HwAddr	IPAddress	PeerIPAddress	Netmask	Gateway	MTU	Rate	Duplex
0:1:1	new	000423C21B72	192.168.25.226	-	255.255.255.0	-	1500	1Gbps	Full
0:1:1	new	000423C21B73	192.168.25.226	-	255.255.255.0	-	1500	n/a	Half
1:2:1	new	000423ADE95E	192.168.25.227	-	255.255.255.0	-	1500	1Gbps	Full
1:2:1	new	000423ADE95F	192.168.25.227	-	255.255.255.0	-	1500	n/a	Half

The columns in the previous example are identified as follows:

- **N:S:P.** The physical position of the port, in the syntax `node:slot:port`.
- **State.** Remote copy end-to-end transport state. Can include one of the following results:
  - ◆ **new.** Configuration is not completed on this port.
  - ◆ **incomplete.** Configuration is not yet completed on the peer port.
  - ◆ **ready.** Configuration is completed on this port and the peer port; transport is ready for use.
  - ◆ **missing.** A configured transport was disconnected.
- **HwAddr.** Hardware address of the Ethernet port indicated in the **N:S:P** column.

- **IPAddress.** IP address of the Ethernet port indicated in the N:S:P column.
- **PeerIPAddress.** IP address of the peer Ethernet port to which the port indicated in the N:S:P column is connected.
- **Netmask.** Netmask for the IP address.
- **Gateway.** Address of the gateway.
- **MTU.** Maximum Transfer Unit (MTU) size for the specified Remote Copy interface (default is 1500). The largest supported value is 9000 and the smallest is 100.
- **Rate.** Actual bit rate of the port indicated in the N:S:P column.
- **Duplex.** Values can be either Full or Half.

## NOTES

None.

---

**COMMAND**

showrsv

**DESCRIPTION**

The `showrsv` command displays SCSI reservation and registration information for Virtual Volume Logical Unit Numbers (VLUNs) bound for a specified port.

**SYNTAX**

```
showrsv [options <arg>] [<VV_name>]
```

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-l <scsi3|scsi2>`

Specifies that either SCSI-3 persistent reservation or SCSI-2 reservation information is displayed. If this option is not specified, information about both SCSI-2 and SCSI-3 reservation will be shown.

`-host <hostname>`

Displays reservation and registration information only for Virtual Volumes (VVs) that are visible to a particular host.

**SPECIFIERS**

`[<VV_name>]`

Specifies the VV name, using up to 31 characters.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays information about reservation and registration information for VLUNs bound with host name w2k\_emx1\_cisco.

```
cli% showrsv -host w2k_emx1_cisco
```

	VVname	Host	Owner	Port	ReservationType
w2k_clusterd.10	w2k_emx1_cisco	10000000C92B9909	1:3:1	SCSI-3(6)	
w2k_clusterd.11	w2k_emx1_cisco	10000000C92B9909	1:3:1	SCSI-3(6)	
w2k_clusterd.12	w2k_emx1_cisco	10000000C92B9909	1:3:1	SCSI-3(6)	
w2k_clusterd.13	w2k_emx1_cisco	10000000C92B9909	1:3:1	SCSI-3(6)	

The columns in the previous example are identified as follows:

- VVname. Name of volume exported on this port to the host in Host column.
- Host. Host name of the host connected to this port.
- Owner/Registrant. WWN of the host that has the reservation or registration.
- Port. Port name in N:S:P format.
- ReservationType. The type of reservation being established either SCSI-2 or SCSI-3. SCSI-3 reservations can be as follows:
  - ◆ 1. (Write exclusive). All read requests allowed. Write requests allowed for the reservation holder only.
  - ◆ 3. (Exclusive access). Read and write requests allowed for the reservation holder only.
  - ◆ 5. (Write exclusive-registrants only). All read requests allowed. Write requests allowed for registered initiators only.
  - ◆ 6. (Exclusive access-registrants only). Read and write requests allowed for registered initiators only.
  - ◆ 7. (Write exclusive-all registrants). All read requests allowed. Write requests allowed for registered initiators only.
  - ◆ 8. (Exclusive access-all registrants). Read and write requests allowed for registered initiators only.

## NOTES

None.

---

**COMMAND**

showsched

**DESCRIPTION**

The `showsched` command shows the state of tasks currently scheduled on the system.

**SYNTAX**

`showsched [options][<schedname>|<pattern>]`

**AUTHORITY**

Super, Service

**OPTIONS**

`-all`

Specifies that all scheduled tasks will be displayed.

`-sortcol <col> [,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). rows with the same information in them as earlier columns will be sorted by values in later columns.

**SPECIFIERS**

`<schedname>|<pattern>...`

Specifies the schedules with the specified name (31 character maximum), matching the glob-style pattern. This specifier can be repeated to display configuration information about multiple schedules. This specifier is not required. If not specified, configuration information for all non-system scheduled tasks in the system is displayed.

RESTRICTIONS

None.

EXAMPLES

The following example displays information of all scheduled tasks.

cli% showsched -all

----- Schedule -----										
SchedName	File/Command	Min	Hour	DOM	Month	DOW	CreatedBy	Status	Alert	NextRunTime
move_back_chunklet1	moverelocpd -f 0,7,14	17	2	*	*	0	3parsvc	active	Y	2009-03-15 03:17:00 PDT
move_back_chunklet2	moverelocpd -f 1,8,15	17	2	*	*	1	3parsvc	active	Y	2009-03-16 02:17:00 PDT
move_back_chunklet3	moverelocpd -f 2,9	17	2	*	*	2	3parsvc	active	Y	2009-03-17 02:17:00 PDT
move_back_chunklet4	moverelocpd -f 3,10	17	2	*	*	3	3parsvc	active	Y	2009-03-18 02:17:00 PDT
move_back_chunklet5	moverelocpd -f 4,11	17	2	*	*	4	3parsvc	active	Y	2009-03-12 02:17:00 PDT
move_back_chunklet6	moverelocpd -f 5,12	17	2	*	*	5	3parsvc	active	Y	2009-03-13 02:17:00 PDT
move_back_chunklet7	moverelocpd -f 6,13	17	2	*	*	6	3parsvc	active	Y	2009-03-14 02:17:00 PDT
remove_expired_vvs	removevv -f	27	*	*	*	*	3parsvc	active	Y	2009-03-11 17:27:00 PDT
-----										
8	total									

The columns in the previous example are identified as follows:

- SchedName. Name of the scheduled task.
- File/Command. The file or command that has been scheduled.
- Schedule. WWN of the host that has the reservation or registration.
  - ◆ Min. The minute of the hour that a scheduled task will execute.
  - ◆ Hour. The hour of the day that a scheduled task will execute.
  - ◆ DOM. The day of the month that a scheduled task will execute.
  - ◆ Month. The month that a scheduled task will execute.
  - ◆ DOW. The day of the week that a scheduled task will execute.
- CreatedBy. The name of the user that created the task.
- Status. Indicates whether the task is active or has been suspended.
- Alert. Indicates whether the task will generate an alert.
- NextRunTime. The time when the next scheduled task will run.

NOTES

If the -all option is not specified then scheduled system tasks created by the 3parsvc user will not be displayed.



---

**COMMAND**

showsnmpmgr

**DESCRIPTION**

The `showsnmpmgr` command displays a list of registered SNMP managers for receiving traps.

**SYNTAX**

showsnmpmgr

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays a list of registered managers:

cli% showsnmpmgr		
Manager	IP	Port
82.185.98.76		162
82.162.9.7		8004
5def:2008:abcd::161a		9162

In the example above:

Host IP. The trap manager's IP address.

Port. The trap manager's port number.

## EXIT STATUS

The following codes are returned indicating success or failure:

- 0 indicates that the command was successful.
- 1 indicates that the command failed.

## NOTES

Managers are registered using the `addsnmpmgr` command. See [addsnmpmgr](#) on page 4.2 for additional information.

---

**COMMAND**

showsnmppw

**DESCRIPTION**

The `showsnmppw` command displays the SNMP community string passwords.

**SYNTAX**

`showsnmppw [options]`

**AUTHORITY**

Super, Edit, Browse

**OPTIONS**

`-rw` | `-r` | `-w`

Specifies that the read-write (`-rw`), read-only (`-r`), or write-only (`-w`) password is displayed.

If not specified, the read/write community string password is displayed.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays all SNMP passwords:

```
cli% showsnmppw  
[password]
```

**EXIT STATUS**

The following codes are returned indicating success or failure:

- 0 indicates that the command was successful.
- 1 indicates that the command failed.

**NOTES**

SNMP passwords are registered using the `setsnmppw` command. See [setsnmppw](#) on page 21.68 for more information.

---

## COMMAND

showspace

## DESCRIPTION

The `showspace` command displays estimated free space for logical disk creation.

## SYNTAX

`showspace [options <arg>]`

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

The following options are used to select the logical disk creation parameters used for space calculation:

`-cpg <CPG_name>|<pattern>`

Specifies that logical disk creation parameters are taken from CPGs that match the specified CPG name or pattern, indicated with a glob-style pattern (see [Glob-Style Pattern](#) on page 2.4 for more information). Multiple CPG names or patterns can be specified using a comma separated list, for example `cpg1,cpg2,cpg3`. Only the `-hist` option can be specified with the `-cpg` option.

`-hist`

Specifies that free space history over time for CPGs specified with the `-cpg` option is displayed. This option can only be used if the `-cpg` option is specified.

`-t <RAID_type>`

Specifies the RAID type of the logical disk: `r0` for RAID-0, `r1` for RAID-1, `r5` for RAID-5, or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklets>`

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row using an integer from 1 through 2147483647. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size in kilobytes using 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port|cage|mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the chunklet characteristics, either *first* (attempt to use the lowest numbered available chunklets) or *last* (attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is *first*.

`-p <pattern>`

Specifies a pattern for candidate disks. Patterns are used to disks that are used for creating LDs. If no pattern is specified, the option defaults to all Fibre Channel (FC) disks.

If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL) and FC, and Solid State Drive (SSD) drives:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (*item*). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0-7). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (*item*). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0-7). The primary path of the disks must be on the specified PCI slot(s).

`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (*item*). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–4). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (*item*). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the *CagePos* column of *showpd* output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (*item*). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0–3). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (*item*). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0–3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more disks. Disks are identified by one or more integers (*item*). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0–3). Disks must match the specified ID(s).

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel or NL, Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.

`-rpm <number>`

Specifies that disks must be of the indicated speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

None.

## RESTRICTIONS

The `-cpg` and `-hist` options cannot be used with any other option.



## EXAMPLES

The following example displays the estimated free space for a RAID-1 LD:

```
cli% showspace -t r1
--Estimated(MB)---
RawFree UsableFree
  13824      6912
```

## NOTES

- The `showspace` command includes chunklets in the process of being initialized as shown by the Normal, Unused, Uninit column of `showpd -c`. The space represented by these chunklets will not actually be available until the initialization process is complete.
- The space calculated is an estimate and not an exact figure.
- For this command MB = 1048576 bytes.

**COMMAND**

showspare [-used]

**DESCRIPTION**

The `showspare` command displays information about chunklets in the system that are reserved for spares and previously free chunklets selected for spares by the system.

**SYNTAX**

showspare

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-used

Show only used spare chunklets. By default all spare chunklets are shown.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays information about spare chunklets on chunklets 53 through 56 on physical disk 4:

```
cli% showspare
PdId Chnk LdName LdCh State Usage Media Sp Cl From To
  4   53 ronnie   0 normal   ld valid  N  N 2:37 ---
  4   54 james   28 normal   ld valid  N  N 0:29 ---
  4   55 dio     28 normal   ld valid  N  N 0:32 ---
  4   56 rocks    0 normal   ld valid  N  N 0:38 ---
```

The columns in the previous example are identified as follows:

- PdId. The physical disk on which the chunklets reside.
- Chnk. The chunklet number.

- **LdName.** The name of the logical disk that is using the spare chunklet.
- **LdCh.** The position of the chunklet on the logical disk.
- **State.** The state of the chunklet as identified by the kernel.
  - ◆ **logging.** I/O to the chunklet is written to the logging logical disk.
  - ◆ **playback.** Data is played back from the logging logical disks.
  - ◆ **passthru.** Chunklets do not process physical disk errors.
  - ◆ **preserved.** Any I/O to the chunklet is written to the preserved logical disks.
  - ◆ **preserved playback.** Data is played back from the preserved logical disks.
  - ◆ **stale.** The chunklet is not available for use because of a medium failure or a connectivity failure.
  - ◆ **normal.** The chunklet is available for use.
  - ◆ **normal , smag.** A servicemag operation is performed on the disks.
  - ◆ **none.** Chunklets were not used by any logical disk.
- **Usage.** Shows whether the spare chunklet is in use by a logical disk.
  - ◆ **available.** The chunklet is available for use as a spare or as a logical disk.
  - ◆ **ld.** The chunklet is in use by a logical disk.
  - ◆ **synch.** The chunklet is both the source and the target of a logical disk relocation operation (synchronizing the chunklet).
  - ◆ **cmprcl.** The system is completing the logical disk relocation operation.
  - ◆ **reldsrc.** Relocation source. The data has been moved to another chunklet.
  - ◆ **reltgt.** Relocation target. The data in the chunklet has been moved from another spare chunklet.
  - ◆ **abtrecl.** Abort relocation. The system is canceling the logical disk relocation operation.
- **Media.** The current status of the physical disk medium for the chunklet.
  - ◆ **valid.** The chunklet is available for use.
  - ◆ **failed.** The medium has encountered errors and is not available for use.

- **Sp.** The spare status of the chunklet; **Y** indicates the chunklet is reserved for spare, **N** indicates a previously free chunklet selected by the system as a spare.
- **Cl.** The clean status of the chunklet. **N** indicates that the chunklet is in-use. **Y** indicates that the chunklet is clean. **Cg** indicates that the chunklet is being cleaned.
- **From.** The initial location of the chunklet before relocation.
- **To.** The destination location of the chunklet during relocation.

## NOTES

The `showpdch` command is a more general and versatile command that can be used instead of `showspare`.

---

**COMMAND**

showsshkey

**DESCRIPTION**

The showsshkey command will displays all SSH public keys that have been set with setsshkey.

**SYNTAX**

showsshkey

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example shows the SSH public keys that have been set by the current user:

```
cli% showsshkey
ID Key
0 ssh-dss AAAAB3NzaC1kc3MAAACBALDVpIVrNEGJn0tFDCSEltLXdGVqBkb8Dm/vziYKXShc6
EIEu3kdD7pK9yhTEKx8q9urwgtKwyYH1KMWYp042kuwIM12UzIvkV6DqbA8Z4aOH81WnBm7pEscD
zHdq6IZJ47MsZLj1xXO2RYyWTFwsbpKYLXoUrrj8vDqCtcVFGqvAAAAFQDtTx3ZLckWAI9Pp/ jgt
JRI+z63wAAAIAtFIyscMhuyma3yMvne8MyUyPlmJUzpKLHKwCNfvwL49bfXSoczitmY+ENHTgtWd
uRaX0uIMmeDeeO5qUUKDfxp5t3KURh147QcU4lNmLCWvR1S5ep1ff0+D7PVjJtRw1ZxLtfQMyGWz
YCMzSRnk4lpD6R392A8XLSzvqrLXat3wAAAIeA1+crC528VaZSiSbeW/FybPx3EhvvV/Zyv6FTFI
OE5+bsWbeQZP33M3yuwTvVuLvphSzpxTKQy1TOAK/Q6XmsvE9ubUdki5X1rIuFvOzU88KJ0X1c+
XPw5+NaI8VXat74YR7dSBE5sbC3EKhz142fd7IH+nHfpHiHQIQSiIIFyw= user@example.com
```

**NOTES**

The public SSH keys that are displayed by this command are used to determine that the private keys that are paired with them are valid, and are not useful in determining the contents of that private key. As such, the contents of the public keys are not sensitive, and having access to them will not aid someone attempting to access an account to which they should not have access.

---

**COMMAND**

showsys

**DESCRIPTION**

The `showsys` command displays the InServ system properties includes system name, model, serial number, and system capacity information, etc.

**SYNTAX**

`showsys [options]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d`

Specifies that more detailed information about the system is displayed.

`-param`

Specifies that the system parameters are displayed.

`-fan`

Displays the system fan information.

`-space`

Displays the system capacity information in MB (1048576 bytes).

`-domainspace`

Displays the system capacity information broken down by domain in MB (1048576 bytes).

`-desc`

Displays the system descriptor properties.

`-mgmtoldports`

Displays the value of the `MgmtOldPorts` parameters (see help for `setsys` command for more information of this parameter). This option is deprecated and will be removed when the `MgmtOldPorts` parameter is removed.

`-devtype FC|NL|SSD`

Displays the system capacity information where the disks must have a device type string matching the specified device type; either Fibre Channel (FC), Nearline (NL), Solid State Drive (SSD). This option can only be issued with the `-space` option.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays information about an InServ Storage Server:

```
cli% showsys
                                     -----(MB)-----
ID -Name- ---Model--- -Serial-  Nodes Master TotalCap AllocCap FreeCap FailedCap
15 s015   InServ T400 1000015      2      0 1063424   208384   855040         0
```

In the previous example:

- ID. The system ID.
- Name. The system name.
- Model. The model type of the InServ server.
- Serial. The system serial number.
- Nodes. The number of nodes in the system.
- Master. The master node ID.
- TotalCap. The total system capacity in MB.
- AllocCap. The allocated system capacity in MB.
- FreeCap. The free system capacity in MB.
- Failed Cap. The failed system capacity in MB.



The following example displays the system descriptor properties of an InServ Storage Server:

```
cli% showsys -desc
-----System s36-----
System Name : Your Name
Location    : Your Facility Address
Owner       : Your Company Name
Contact     : Joe Admin
Comment     : Your Notes
```

The following example displays the system capacity in MB for an InServ Storage Server:

```
cli% showsys -space
-----System Capacity (MB)-----
Total Capacity      : 1063424
  Allocated          : 208384
    Volumes          : 103937
      Base Volumes   : 4266
        User         : 4266
          Copy        : 0
            Admin     : 0
      CPGs (TPVVs & CPVVs) : 98304
        Copy         : 65536
          Used        : 0
            Unused    : 65536
            Admin     : 32768
              Used    : 0
                Unused : 32768
          Unmapped    : 1367
    System            : 104448
      Internal        : 104448
        Spare         : 0
          Used        : 0
            Unused    : 0
    Free              : 855040
      Initialized     : 855040
        Uninitialized : 0
          Failed       : 0
```

The following example displays more detailed (-d option) information about the same storage server:

```
cli% showsys -d
-----General-----
System Name      :      S424
System Model     :      InServ E200
Serial Number    :      1100424
System ID        :      424
Number of Nodes  :      2
Master Node      :      0
Nodes Online     :      0,1
Nodes in Cluster :      0,1

-----System Capacity (MB)-----
Total Capacity   :      6277120
Allocated Capacity :      687872
Free Capacity    :      5589248
Failed Capacity  :      0

-----System Fan-----
Primary Node ID   :      0
Secondary Node ID :      1
State             :      OK
LED               :      Green
Speed             :      Normal

Primary Node ID   :      1
Secondary Node ID :      0
State             :      OK
LED               :      Green
Speed             :      Normal

-----System Descriptors-----
Location         :
Owner            :
Contact          :
Comment          :
```

The following example shows system parameters (`-param` option) for an InServ Storage Server:

```
cli% showsys -param
System parameters from configured settings

---Parameter---      ---Value---
RawSpaceAlertFC       :                0
RawSpaceAlertNL       :                0
RemoteSyslog          :                1
RemoteSyslogHost      :    192.168.6.15
SparingAlgorithm      :           Minimal
CopySpaceReclaim      :                0
EventLogSize          :                3M
VVRetentionTimeMax    :    336 Hours
```

## NOTES

- See [setsys](#) on page 21.75 for information on setting the threshold parameters indicated by the Value column in the output for `showsys -param`.
- In the output for `showsys -param`, `(from configured settings)` indicates that the system parameters displayed have been successfully read from the Persistent Repository (PR). If the PR is not available (most likely because of problems with the admin volume), the output reads `(from default settings)` and the values displayed would indicate the system defaults. When `(from default settings)` is displayed, system parameters can not be updated.
- The system capacity information may have some overlaps among Volumes, System, Failed Capacities.
- If the `VVRetentionTimeMax` is 0, then the volume retention time in the system is disabled.

---

**COMMAND**

showsysmgr

**DESCRIPTION**

The `showsysmgr` displays startup state information about the system manager.

**SYNTAX**

showsysmgr

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the startup state information about the system manager. In this example, the system manager reports that it is up and running.

```
cli% showsysmgr
System is up and running from Thu May 24 15:39:22 PDT 2007
```

If the system has experienced a power failure, issuing the `showsysmgr` command displays the following message:

```
cli% showsysmgr
System is recovering from a previous powerfailure. Please use the CLI
commands for 'showvv', 'showld', 'showpd' to see any unstarted vvs, lds,
pds.
```

If the system has attempted powerfail recovery three times, and during the recovery attempts encountered kernal panics because of hardware or software errors, the following message is displayed. See [Notes](#) in the following section for additional information.

```
cli% showsysmgr
System is recovering from a previous power failure.
Please use the 'showvv', 'showld', 'showpd' CLI commands to
check for not started vvs, lds, pds.
Use force_idewipe to wipe pfail partititon and restart system with all
lds/vvs being checked. This can cause some data to be lost.
```

If the system has attempted powerfail recovery, and during the recovery encountered kernal panics because of hardware or software errors, the error could be associated with a specific volume. The following message is displayed. See [Notes](#) in the following section for additional information.

```
cli% showsysmgr
System is recovering from a previous power failure.
Please use the 'showvv', 'showld', 'showpd' CLI commands to
check for not started vvs, lds, pds.
Use force_iderecovery to start pfail recovery.
VVs with the id(s) 1 will lose data.
```

## NOTES

- If the system has experienced a power failure, issue the `showvv`, `showld`, or `showpd` commands to determine if any physical disks, logical disks, or virtual volumes are unstarted. If the system stays in the powerfail state for longer than 15 minutes, the `setsysmgr` command can be issued with caution. See the [setsysmgr](#) command for more information.



**CAUTION:** Issuing the `setsysmgr force_idewipe` command can result in data loss.

- If the system has experienced a power failure, any cached data is saved to the IDE partition of each node. Upon restoration of power, the saved cached data is reapplied (powerfail recovery). During powerfail recovery, if kernal panics are encountered because of hardware or software errors, the powerfail recovery process fails. In such cases, the system attempts recovery three times. After three attempts, the system waits for user intervention. Issue the `setsysmgr force_idewipe` command to force the system to restart and check all logical

disks and virtual volumes upon restarting. See [setsysmgr](#) on page 21.79 for more information.

- During powerfail recovery, a kernel panic because of hardware or software errors might indicate a damaged volume. In such cases, after the system restarts, it waits for user intervention. Issue the `setsysmgr force_iderecovery` command to force the system to start another powerfail recovery. Any saved data for virtual volumes with IDs displayed in the `showsysmgr` command output can be lost. Those virtual volumes are checked when the system restarts. See [setsysmgr](#) on page 21.79 for more information.

---

**COMMAND**

`showtarget`

**DESCRIPTION**

The `showtarget` command displays information about unrecognized targets.

**SYNTAX**

`showtarget`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays information about unrecognized targets:

```
cli% showtarget
no unknown targets listed
```

In the previous example, there are no unknown targets. If unknown targets are found, information (port, node WWN, port WWN, and state) about each target found is displayed.

**NOTES**

Some hosts might appear as unknown targets.

---

**COMMAND**

`showtask`

**DESCRIPTION**

The `showtask` command displays information about tasks. By default, this command displays all non-system tasks on the system within the last 24 hours.

**SYNTAX**

`showtask [options <arg>] [<task_name> | <pattern>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-all`

Specifies that all scheduled tasks will be displayed. Unless the `-all` option is specified, system tasks are not displayed.

`-done`

Display includes only tasks that are successfully completed.

`-failed`

Display includes only tasks that are unsuccessfully completed.

`-active`

Display includes only tasks that are currently in progress.

`-t <hours>`

Show only tasks that started within the past `<hours>`, where `<hours>` is an integer from 1 through 9999999.

`-type <task_type | pattern>`

Specifies that specified patterns are treated as glob-style patterns and that all tasks whose types match the specified pattern are displayed. To see the different task types use the `showtask` column help.



`-d <task_ID>...`

Show detailed task status for specified tasks. Tasks must be explicitly specified using their task IDs `<task_ID>`. Multiple task IDs can be specified. This option cannot be used in conjunction with other options.

`-sortcol <col>[, <dir>][:<col>[, <dir>]...]`

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays all tasks on a system. In this case, the last task is still in progress.

```
cli% showtask
Id   Type  Name   Status Phase Step  -----StartTime-----  -----FinishTime-----
1  tune_vv testr1  Done   0/0   0/0 Wed Oct 06 18:44:05 EDT 2004 Wed Oct 06 18:44:57 EDT 2004
2  tune_vv testr1  Done   0/0   0/0 Wed Oct 06 19:44:34 EDT 2004 Wed Oct 06 19:45:10 EDT 2004
3  tune_vv testr1  Active 2/3   5/8 Wed Oct 06 19:49:31 EDT 2004
```

The columns in the previous example are identified as follows:

- Id. The task ID for the displayed task.
- Type. The task type. Task types are as follows:
  - ◆ vv\_copy. Track physical copy operations (`createvvcopy` command).

- ◆ `move_regions`. Track region move operations (`movereg` command).
  - ◆ `promote_sv`. Track virtual copy promotes (`promotesv` command). Requires a 3PAR Virtual Copy license.
  - ◆ `remote_copy_sync`. Track Remote Copy volume group synchronizations (`syncrcopy` and `startrcopygroup` commands). Requires a 3PAR Remote Copy license.
  - ◆ `tune_vv`. Track 3PAR System Tuner volume tuning operations (`tunealddv` command). Requires a 3PAR Dynamic Optimization license.
  - ◆ `tune_vv_restart`. Track restarted 3PAR System Tuner volume tuning operations (`tunealddv -restart` command). Requires a 3PAR Dynamic Optimization license.
  - ◆ `tune_vv_rollback`. Track rolling back of 3PAR System Tuner volume tuning operation that was previously interrupted (`tunealddv -rollback` command). Requires a 3PAR Dynamic Optimization license.
- **Name.** The name of the system object that is being operated on. When there are multiple objects, those objects are not identified individually but as a group (for example, multiple LDs or multiple CPGs).
  - **Status.** The task state. Task states are as follows:
    - ◆ **Done.** The task has completed successfully.
    - ◆ **Active.** The task is still in process.
    - ◆ **Cancelled.** The task was canceled by you.
    - ◆ **Failed.** The task failed to complete because of a reason other than user cancelation.
  - **Phase.** For Active tasks, indicates the number of completed phases and the total number of phases in the current step, using the syntax `<#completed_phase>/<#total_phases>`. Note that each step is composed of an arbitrary number of phases.
  - **Step.** For Active tasks, indicates the number of completed steps and the total number of steps in the current task, using the syntax `<#completed_step>/<#total_steps>`. Note that each task is composed of an arbitrary number of steps.
  - **Start Time.** Indicates the time that the task was started.
  - **Finish Time.** For Done, Cancelled, and Failed tasks, indicates the time that the task stopped because of completion, cancelation, or failure.

The following example shows details about a specific task using the task ID (task ID is 3).

```
cli% showtask -d 3
Id Type          Name                Status Phase Step -----StartTime----- -----FinishTime-----
3 snapspace_accounting ss_accounting done    ---  --- 2009-02-27 11:37:46 PST 2009-02-27 11:37:48 PST

Detailed status:
2009-02-27 11:37:46 PST Created      task.
2009-02-27 11:37:46 PST Started      snapshot usage data collection process for VVs
2009-02-27 11:37:46 PST Updated      snapshot usage data for VV North
2009-02-27 11:37:47 PST Updated      snapshot usage data for VV VV_Tech_1.0020.0000.ro
2009-02-27 11:37:48 PST Updated      snapshot usage data for VV VV_Tech_1
2009-02-27 11:37:48 PST Finished     snapshot usage data collection process.
```

## NOTES

- See the *InForm OS CLI Administrator's Manual* for additional information and examples regarding task management and task management commands, as well as information about Dynamic Optimization.
- By default, this command shows all tasks that started within the last 24 hours.
- The system stores information for the most recent 1000 tasks. Task ID numbers roll at 9999.

---

**COMMAND**

showtemplate

**DESCRIPTION**

The showtemplate command displays existing templates that can be used for Virtual Volume (VV), Logical Disk (LD), or Common Provisioning Group (CPG) creation.

**SYNTAX**

```
showtemplate [options] [<template_name_or_pattern>...]
```

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-t VV|LD|CPG

Specifies that the template type displayed is a VV, LD, or CPG template.

-fit

Specifies that the properties of the template is displayed to fit within 80 character lines.

**SPECIFIERS**

<template\_name\_or\_pattern>

Specifies the name of a template, using up to 31 characters or glob-style pattern for matching multiple template names. If not specified, all templates are displayed.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the properties of all templates in CLI line format:

```
cli% showtemplate
Name Type Other Options
Test1 VV -nro -desc "1st Shot"
CPGTemplate CPG -nro -ro -t r1 -ha cage -ssz 2 -ss 256 -sdgs 0
```

**NOTES**

If no options are specified, all existing templates are displayed.

---

**COMMAND**

showtoc

**DESCRIPTION**

The `showtoc` command displays the system table of contents summary that provides a summary of the system's resources.

**SYNTAX**

showtoc

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the system table of contents.:

```
cli% showtoc
gltab toc generation: 313
Toc header magic num: 42444854
Toc in use len: 70944
Toc not in use len: 0
Toc on disk len: 38912
Toc version: 69
Toc generation: 313
Toc pd_entries: 80
Toc ld_entries: 5
Toc vv_entries: 1
Toc cage_entries: 4
```

## NOTES

None.

---

**COMMAND**

showtocgen

**DESCRIPTION**

The `showtocgen` command displays the table of contents generation number.

**SYNTAX**

showtocgen

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the table of contents generation number:

```
cli% showtocgen
Table of Contents generation number: 4292
```

**NOTES**

The table of contents generation number increases each time there is a change in the system configuration.

---

## COMMAND

`showuser`

## DESCRIPTION

The `showuser` command displays information about one or all users, including the username, authority level, and system resources to which a given user has access.

The `showuser` command shows account information for local users (those created with the `createuser` command) or shows the SSH key information for local and LDAP users (enabled by the `setauthparam` command).

## SYNTAX

The syntax for the `showuser` command can be the following:

- `showuser [-oneline] [<user_name>]`
- `showuser -k`

## AUTHORITY

Super, Edit, Browse, Service

## OPTIONS

`-k`

Shows when users have set a public SSH key.

`-oneline`

Shows all information about a user on one line.

## SPECIFIERS

`<user_name>`

Specifies your login name, using up to 31 characters. This specifier is not required on the command line. If no specifier is used, information about all users is displayed.

## RESTRICTIONS

Because the output is only for local users and is further limited when the user is at the Browse or Service authority level, an LDAP user cannot see any other user account information. The `showuserconn` command can be used to see the privilege level of the current user and whether they are local or LDAP.



## EXAMPLES

```
cli% showuser root
Username Domain Privilege Default
root      all      super      N
cli%
```

The previous example indicates user `root` being a member of all domains. Valid output includes `all`, or `<domain_name>`. The columns are identified as follows:

- Username. Your user name.
- Domain. The domain to which the user belongs.
- Privilege. The privilege level required to run the command.
- Default. Indicates whether the domain is the default domain of the user.

## NOTES

- Users are created using the `createuser` command. See [createuser](#) on page 11.77 for more information.
- For the first usage (without the `-k` option), the `showuser` command displays account information for all local users when the `<username>` is not provided or the account information for just the local user with the name `<username>` when it is provided. The output for users with Browse or Service level privileges is limited to the user's own accounts and if the `<username>` is provided, it must be the user's name. The output shows the users' username, domain, privilege, and whether the domain is the default for the user.
- For the second usage (`-k`), users who have set a public SSH key with the `setsshkey` command are displayed, one per line. For users with Browse or Service privilege levels, the output is limited to the current user.
- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.

**COMMAND**

showuseracl

**DESCRIPTION**

The showuseracl command displays a user’s access control list (ACL).

**SYNTAX**

showuseracl

**AUTHORITY**

Super, Edit, Browse, Service

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays ACL information about all system users:

```
cli% showuseracl
User          Operation Object_Names_or_Patterns
suser         updatevv avvro*
buser         updatevv avvr*,vv0,cpvv0,tpvv0 -f
ruser         updatevv vv0 cpvv0 tpvv0 avv*
```

The columns in the previous example are identified as follows:

- User. Your user name.
- Operation. The command each user is allowed to execute.
- Object\_Names\_or\_Patterns. The object on which the operations are performed.

**NOTES**

None.

---

**COMMAND**

`showuserconn`

**DESCRIPTION**

The `showuserconn` command displays information about users who are currently connected (logged in) to the InServ Storage Server.

**SYNTAX**

`showuserconn [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-current`

Shows all information about the current connection only.

`-sortcol <col> [,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting `<dir>` can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays information about user connections:

```
cli% showuserconn
  Id Name  --IP_Addr--- Domain Level ----Connected_since---- Current
23180 root  192.168.17.7 all      super 2007/05/31 15:02:53 PDT current
-----
      1 total
```

The columns in the previous example are identified as follows:

- **Id.** Your InServ Storage Server identification number.
- **Name.** Indicates the user name under which you logged in.
- **IP\_Addr.** The IP address where the user connection exists.



**NOTE:** The `Domain` column is displayed only *if the CLI was started with the `-listdom global` option or with the `LISTDOM` environment variable*. Refer to the *InForm OS CLI Administrator's Manual* for instructions on using global options and setting environment variables.

- **Domain.** Indicates the domain in which the user has privileges.
- **Level.** Indicates the user's privilege level.

## NOTES

A `Domain` column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.

---

**COMMAND**

`showversion`

**DESCRIPTION**

The `showversion` command displays information about the storage server software. Use this command to determine if your system is using the latest software version.

**SYNTAX**

`showversion [options]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-a`

Show all component versions.

`-b`

Show build levels.

`-s`

Show release version number only (useful for scripting).

If no options are specified, the overall version of the software is displayed.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays comprehensive information about the system software:

```
cli% showversion
Release version 2.3.1.60 (DEVEL)
Patches:  None

Component Name          Version
CLI Server              2.3.1
CLI Client              2.3.1
System Manager          2.3.1
Kernel                  2.3.1
TPD Kernel Code         2.3.1
```

## NOTES

- When displaying all versions, for certain components multiple versions might be installed. In such cases, multiple lines are displayed.
- If no options are specified, the overall version of the software is displayed.

---

## COMMAND

showvlun

## DESCRIPTION

The `showvlun` command displays information about Virtual Volume Logical Unit Numbers (VLUNs) in the system, such as all active and template VLUNs. The display is divided into two sections: the upper provides information about active VLUNs and the lower provides information about VLUN templates.

## SYNTAX

```
showvlun [options <arg>]  
showvlun -listcols
```

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

`-listcols`

List the columns available to be shown in the `-showcols` option described below (see `clihelp -col showvlun` for help on each column).

The [options] are generally of two kinds: those that select the type of information that is displayed, and those that filter the list of VLUNs that are displayed.

By default (if none of the information selection options below are specified) the following columns are shown: LUN VVName HostName Host\_WWN Port Type

Options that select the type of information shown include the following:

`-showcols <column>[,<column>...]`

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run `'showvlun -listcol'` to list the available columns.

Run `'clihelp -col showvlun'` for a description of each column.

`-lvw`

Show the World Wide Name (WWN) of the Virtual Volume (VV) associated with the VLUN.

`-a`

Shows only active VLUNs.

`-t`

Shows only VLUN templates.

`-pathsum`

Shows path summary information for active VLUNs.

`-host {<hostname|pattern>}...`

Specifies that only VLUNs exported to hosts that match the hostname or glob-style patterns (see Help on sub,globpat).

`-v {<VV_name>|<pattern>}...`

Requests that only Logical Disks (LD) mapped to VVs that match any of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list (for example `-vv <VV_name>,<VV_name>...`).

`-l {<LUN|pattern>}...`

Specifies that only exports to the specified LUN are displayed. This specifier can be repeated to display information for multiple LUNs.

`-nodes <nodelist>`

Requests that only VLUNs for specific nodes are displayed. The node list is specified as a series of integers separated by commas (for example `0,1,2`). The list can also consist of a single integer (for example `1`).

`-slots <slotlist>`

Requests that only VLUNs for specific slots are displayed. The slot list is specified as a series of integers separated by commas (for example `0,1,2`). The list can also consist of a single integer (for example `1`).

`-ports <portlist>`

Requests that only VLUNs for specific ports are displayed. The port list is specified as a series of integers separated by commas (for example `0,1,2`). The list can also consist of a single integer (for example `1`).



`-domain {<domain_name|pattern>}...`

Shows only the VLUNs whose VVs are in domains with names that match one or more of the `<domainname_or_pattern>` options. This option does now allow listing objects within a domain of which the user is not a member. Multiple domain names or patterns can be repeated using a comma-separated list.

`-sortcol <col> [,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number `<col>`. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting `<dir>` can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

## **SPECIFIERS**

None.

## **RESTRICTIONS**

None.

## EXAMPLES

The following example displays all active and template VLUNs:

```
cli% showvlun
Active VLUNs
Lun  VVname      Host  ----Host_WWN----  Port  Type
  0  tpv.0  pe750-07  210000E08B056C21  0:2:1  host
  1  tpv.1  pe750-07  210000E08B056C21  0:2:1  host
  2  tpv.2  pe750-07  210000E08B056C21  0:2:1  host
  3  tpv.3  pe750-07  210000E08B056C21  0:2:1  host
  0  tpv.0  pe750-07  210100E08B256C21  1:2:1  host
  1  tpv.1  pe750-07  210100E08B256C21  1:2:1  host
  2  tpv.2  pe750-07  210100E08B256C21  1:2:1  host
  3  tpv.3  pe750-07  210100E08B256C21  1:2:1  host
  0  test.0      sun  210000E08B023F71  0:2:2  host
  1  test.1      sun  210000E08B023F71  0:2:2  host
  2  test.2      sun  210000E08B023F71  0:2:2  host
  3  test.3      sun  210000E08B023F71  0:2:2  host
  0  test.0      sun  210000E08B023C71  1:5:1  host
  1  test.1      sun  210000E08B023C71  1:5:1  host
  2  test.2      sun  210000E08B023C71  1:5:1  host
  3  test.3      sun  210000E08B023C71  1:5:1  host
-----
16

VLUN Templates
Lun  VVname      Host  ----Host_WWN----  Port  Type
  0  tpv.0  pe750-07  -----  ---  host
  1  tpv.1  pe750-07  -----  ---  host
  2  tpv.2  pe750-07  -----  ---  host
  3  tpv.3  pe750-07  -----  ---  host
  0  test.0      sun  -----  ---  host
  1  test.1      sun  -----  ---  host
  2  test.2      sun  -----  ---  host
  3  test.3      sun  -----  ---  host
-----
8
```

The following example displays all active and template VLUNs using glob-style patterns:

```
cli% showvln -v *.2,*.3
Active VLUNs
Lun  VVname      Host  ----Host_WWN----  Port Type
  2  tpv.2  pe750-07  210000E08B056C21  0:2:1 host
  3  tpv.3  pe750-07  210000E08B056C21  0:2:1 host
  2  tpv.2  pe750-07  210100E08B256C21  1:2:1 host
  3  tpv.3  pe750-07  210100E08B256C21  1:2:1 host
  2  test.2      sun  210000E08B023F71  0:2:2 host
  3  test.3      sun  210000E08B023F71  0:2:2 host
  2  test.2      sun  210000E08B023C71  1:5:1 host
  3  test.3      sun  210000E08B023C71  1:5:1 host
-----
8

VLUN Templates
Lun  VVname      Host  ----Host_WWN----  Port Type
  2  tpv.2  pe750-07  -----          --- host
  3  tpv.3  pe750-07  -----          --- host
  2  test.2      sun  -----          --- host
  3  test.3      sun  -----          --- host
-----
4
```

The columns in the previous examples are identified as follows:

- **Lun.** The LUN to which the virtual volume is exported.
- **VVname.** The name of the exported virtual volume.
- **Host.** The name of the host to which the virtual volume is exported. In the *Templates* section, the port-presents rule contains no value.
- **Host\_WWN.** The WWN of the host to which the virtual volume is exported. In the *Templates* section, dashes mean this column does not apply.
- **Port.** The port (node:slot:port) on which the host to which the virtual volume is exported appears. In the *Templates* section, the host-sees rule contains dashes.
- **Type.** The type of export rule that is applied when the virtual volume is exported. Possible entries are *port* for port-presents, *host* for host-sees, *hostset* for host-set VLUNs, *matched-set* for matched-set.

## NOTES

- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.
- See [Glob-Style Pattern](#) on page 2.4 for more information.

---

## COMMAND

`showvv`

## DESCRIPTION

The `showvv` command displays information about all Virtual Volumes (VVs) or a specific VV in a system.

## SYNTAX

- `showvv [options <arg>] [<VV_name|pattern|VV_set>...]`
- `showvv -listcols`

## AUTHORITY

Super, Service, Edit, Browse

## OPTIONS

`-listcols`

List the columns available to be shown in the `-showcols` option described below (see `clihelp -col showvv` for help on each column).

The [options] are generally of two kinds: those that select the type of information that is displayed, and those that filter the list of VVs that are displayed.

By default (if none of the information selection options below are specified) the following columns are shown:

Id, Name, Prov, Type, CopyOf, BsId, Rd, State, Adm\_Rsvd\_MB, Snp\_Rsvd\_MB,  
Usr\_Rsvd\_MB, VSize\_MB.

Options that select the type of information shown include the following:

`-showcols <column>[,<column>...]`

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run `showvv -listcol` to list the available columns.

Run `clihelp -col showvv` for a description of each column.

-d

Displays detailed information about the VVs. The following columns are shown:

ID, Name, Rd, Mstr, Prnt, Roch, Rwch, PPrnt, PBlkRemain, VV\_WWN, CreationTime.

-pol

Displays policy information about the VV. The following columns are shown:

ID, Name, Policies.

-s

Displays Logical Disk (LD) space use by the VVs. The following columns are shown:

ID, Name, Prov, Type, Adm\_Rsvd\_MB, Adm\_Used\_MB, Snp\_Rsvd\_MB, Snp\_Used\_MB, Snp\_Used\_Perc, Warn\_Snp\_Perc, Limit\_Snp\_Perc, Usr\_Rsvd\_MB, Usr\_Used\_MB, Usr\_Used\_Perc, Warn\_Usr\_Perc, Limit\_Usr\_Perc, Tot\_Rsvd\_MB, VSize\_MB.



**NOTE:** For snapshot (vcopy) VVs, the Adm\_Used\_MB, Snp\_Used\_MB, Usr\_Used\_MB and the corresponding \_Perc columns have an (\*) before the number for two reasons: to indicate that the number is an estimate that must be updated using the updatesnapspace command, and to indicate that the number is not included in the total for the column since the corresponding number for the snapshot's base VV already includes that number.

-r

Displays raw space used by the VVs. The following columns are shown:

ID, Name, Prov, Type, Adm\_RawRsvd\_MB, Adm\_Rsvd\_MB, Snp\_RawRsvd\_MB, Snp\_Rsvd\_MB, Usr\_RawRsvd\_MB, Usr\_Rsvd\_MB, Tot\_RawRsvd\_MB, Tot\_Rsvd\_MB, VSize\_MB.

-zone

Displays mapping zone information for VVs. The following columns are shown:

ID, Name, Prov, Type, VSize\_MB, Adm\_Zn, Adm\_Free\_Zn, Snp\_Zn, Snp\_Free\_Zn, Usr\_Zn, Usr\_Free\_Zn.

-g

Displays the volume's SCSI geometry settings for VVs. The following columns are shown:

ID, Name, SPT, HPC, SctSz.

**-alert**

Indicates whether alerts are posted on behalf of the VVs. The following columns are shown:

ID, Name, Prov, Type, VSize\_MB, Snp\_Used\_Perc, Warn\_Snp\_Perc,  
Limit\_Snp\_Perc, Usr\_Used\_Perc, Warn\_Usr\_Perc, Limit\_Usr\_Perc,  
Alert\_Adm\_Fail\_Y, Alert\_Snp\_Fail\_Y, Alert\_Snp\_Wrn\_Y, Alert\_Snp\_Lim\_Y,  
Alert\_Usr\_Fail\_Y, Alert\_Usr\_Wrn\_Y, Alert\_Usr\_Lim\_Y.

**-alerttime**

Shows times when alerts were posted (when applicable). The following columns are shown:

ID, Name, Alert\_Adm\_Fail, Alert\_Snp\_Fail, Alert\_Snp\_Wrn, Alert\_Snp\_Lim,  
Alert\_Usr\_Fail, Alert\_Usr\_Wrn, Alert\_Usr\_Lim.

**-cpprog**

Shows the physical copy and promote progress. The following columns are shown:

ID, Name, Prov, Type, CopyOf, VSize\_MB, Copied\_MB, Copied\_Perc.

**-cpgalloc**

Shows CPGs from which a VV is allocated, if any. The following columns are shown:

ID, Name, Prov, Type, UsrCPG, SnpCPG.

**-state**

Shows the detailed state information for the VVs. The following columns are shown:

ID, Name, Prov, Type, State, Detailed\_State.

**-hist**

Shows the history information of the VVs. The following columns are shown:

ID, Name, Prov, Type, CreationTime, ExpirationEndTime, ExpirationTime,  
SpaceCalcTime, Comment.

**-rcopy**

This option appends two columns, RcopyStatus and RcopyGroup, to any of the display options above.

**-notree**

Do not display VV names in tree format. Unless either the **-notree** or the **-sortcol** option described below are specified, the VVs are ordered and the names are indented in tree format to indicate the virtual copy snapshot hierarchy.

```
-sortcol <col> [ ,<dir>][:<col>[ ,<dir>]...]
```

Sorts command output based on the column number <col>. Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting <dir> can be specified as follows:

```
inc
```

Sort in increasing order (default).

```
dec
```

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by the values in later columns.

A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the `TPDLISTDOM` environment variable set.

The VV filtering option include:

```
-p <pattern>
```

Pattern for matching virtual volumes to show (see below for description of <pattern>). If the `-p` option is specified multiple times, each instance of <pattern> adds additional candidate disks that match that pattern.

A <pattern> is one or more of the following:

```
-cpg {<CPG_name|pattern>}...
```

Displays only VVs with `UsrCPG` or `Snpcpg` names that matches one or more of the <CPG\_name> or <pattern> specifier. Multiple <CPG\_name> or <pattern> specifiers can be repeated using a comma-separated list (for example `-cpg <CPG_name> , <CPG_name> ...`).

```
-prov {<prov|pattern>}...
```

Displays only VVs with `prov` (provisioning) values that match the <prov> or <pattern> specifier. Multiple <prov> or <patterns> specifiers can be repeated using a comma-separated list (for example `-prov <prov> , <prov> ...`).

```
-type <type|pattern>...
```

Displays only VVs of types that match the <type> or <pattern> specifier. Multiple <type> or <patterns> specifiers can be repeated using a comma-separated list (for example `-type <type> , <type> ...`).



`-host {<host_name|pattern>}...`

Displays only VVs that are exported as VLUNs to hosts with names that match one or more of the specified `<host_name>` or `<patterns>`. See [Specifiers](#) for additional information on patterns as used with this command. Multiple `<host_name>` or `<pattern>` specifiers can be repeated using a comma-separated list (for example `-host <host_name>,<host_name>...`).

`-baseid {<baseid|pattern>}...`

Displays only VVs with `bsid` column that matches one or more of the `<baseid>` or `<pattern>` specifiers. Multiple `<baseid>` or `<pattern>` specifiers can be repeated using a comma-separated list (for example `-baseid <baseid>,<baseid>...`).

`-copyof {<vv_name|pattern>}...`

Displays only VVs with `CopyOf` columns that match one or more of the `<copyof>` or `<pattern>` specifiers. Multiple groups or patterns can be repeated using a comma-separated list (for example `-rcopygroup <group_name>,<group_name>...`).

`-rcopygroup {<group_name|pattern>}...`

Displays only VVs that are in Remote Copy groups that match one or more of the specified groups or patterns. Multiple groups or patterns can be repeated using a comma-separated list (for example `-rcopygroup <group_name>,<group_name>...`).

`-policy <policy|pattern>...`

Show only VVs whose policy matches the one or more of the policies or patterns.

`-domain {<domain_name|pattern>}...`

Shows only VVs that are in domains with names matching one or more of the specified domains or patterns. This option does now allow listing objects within a domain of which the user is not a member. Multiple domain names or patterns can be repeated using a comma-separated list (for example `-domain <domain_name>,<domain_name>...`).

`-expired`

Show only VVs that have expired.

`-exp <time>{d|D|h|H}`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either `d` or `D` for day and `h` or `H` for hours following the entered time value.

`-retained`

Shows only VVs that have a retention time.

`-failed`

Shows only failed VVs.

The following VV filtering options are deprecated and will be removed in a future release. Instead, use the appropriate `-p <pattern>` option described above.

`-cpg <CPG_name|pattern>...`

Show only VVs whose `UsrCPG` or `Snpcpg` matches the one or more of the `<CPG_name>` or `<pattern>` specifier. This option is deprecated without the preceding `-p` option. Use the `-p -cpg` option described above instead.

`-tpvv`

Show only Thin Provisioned VVs (TPVVs). This option is deprecated. Use the `-p -prov tp*` option described above instead.

`-host <hostname|pattern>...`

Show only VVs that are exported as VLUNs to hosts with names that match one or more of the `<hostname>` or `<pattern>` specifier. This option is deprecated without the preceding `-p` option. Use `-p -host` option described above instead.

`-rcopygroup <groupname|pattern>...`

Show only VVs that are in remote copy groups that match one or more of the `<groupname>` or `<pattern>` specifier. This option is deprecated without the preceding `-p` option. Use `-p -rcopygroup` option described above instead.

## SPECIFIERS



**NOTE:** If one or more <VV\_name|pattern> are specified, then the virtual volumes with names that match any of the patterns are listed. Otherwise, all virtual volumes are listed. The virtual volumes shown include snapshots. Patterns are glob-style (shell-style) patterns (see Help for `sub, globpat`).

<VV\_name|pattern|VV\_set>...

Specifies the VV name (up to 31 characters in length) matching the glob-style pattern or that are members of the supplied VV set will have their information displayed. The VV\_set name must start with `set :`. This specifier can be repeated to display configuration information about multiple VVs. This specifier is not required on the command line. If not specified, configuration information for all VVs in the system is displayed.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays information about all virtual volumes:

```
cli% showvv
```

Id	Name	Prov	Type	CopyOf	BsId	Rd	-Detailed_State-	---Rsvd(MB)---	Adm	Snps	Usrs	-(MB)-	VSize
0	admin	full	base	---	0	RW	normal		0	0	10240	10240	
1	par	full	base	Scott	1	RW	normal	128	512	16384		256	
2	copy	cpvv	vcopy	Jeff	2	RW	normal	0	0	16384		5120	
3	vvcp.1.2	tpvv	base	par	1	RO	normal	--	--	16384		10240	
4	tpvv	full	base	Eric	4	RW	normal	128	512	65536		10240	
-----													
5	total							256	1024	43008		36096	

The columns in the previous example are identified as follows:

- ID. The ID of the virtual volume.
- Name. The name of the virtual volume.
- Prov. The provisioning for the VV. Can be one of the following:
  - ◆ full. Fully provisioned VV, either with no Snp (snapshot) space or with statically allocated Snp space.

- ◆ **tpvv.** Thin Provisioned VV, with space for the base volume allocated from the `Usr` space that is associated with the `UsrCPG`. Snapshots allocate space from the `Snp` space associated with the `SnpCPG` (if any).
  - ◆ **cpvv.** Commonly Provisioned VV. The `Usr` space for this VV is fully provisioned and the `Snp` space is associated with a `SnpCPG`.
  - ◆ **tpsd.** An old-style Thin Provisioned VV (created on a 2.2.4 release or earlier) where both the base VV and snapshot data are allocated from the `Snp` space associated with the `UsrCPG`.
  - ◆ **snp.** Since the VV is Type `vcopy` (snapshot) it is the provisioning that is associated with the base VV.
- **Domain.** The domain to which the virtual volume belongs. Valid values are `-` or `<domain_name>`. If the domain does not exist, `-` is displayed. The Domain column can be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable.
  - **Type.** Indicates the copy type of virtual volume and can be one of the following:
    - ◆ **base.** Base volume (not a copy).
    - ◆ **pcopy.** Physical copy (full copy).
    - ◆ **vcopy.** Snapshot copy (virtual copy).
  - **CopyOf.** Displays the virtual volume from which the copy was made.
    - ◆ **---**. The virtual volume is a base volume.
    - ◆ **name.** The name of the virtual volume from which a copy was made.
  - **BsID.** If the virtual volume is a base volume or a physical copy, the base volume ID number is the same as the ID number in the `Id` column. If the virtual volume is a virtual copy, the column shows the base volume from which the copy was made.
  - **Rd.** Indicates whether the virtual volume is read/write (RW) or read-only (RO).
  - **State.** Indicates the current status of the virtual volume.
    - ◆ **preserved.** A logical disk used by the virtual volume has unavailable RAID sets. The data belonging to the virtual volume is saved on the preserved logical disk. Host access to the virtual volume is unavailable until RAID sets are made available.

- ◆ `started`. The virtual volume has been started and is available for use.
- ◆ `started, stl`. The virtual volume is stale.
- ◆ `started, cpf`. The virtual volume copy or promote has failed.
- ◆ `not_started`. The virtual volume has not yet been started (perhaps because one of its underlying logical disks has not yet been started).
- ◆ `pmt`. A promote is currently occurring on the base volume.

The following `State` values indicate that an uncontrolled shutdown has occurred and the volume was not properly closed before the shutdown. Volumes are checked for internal validity before being started.

- ◆ `auto_check`. The virtual volume is checked for validity; however, it is waiting in a queue for computing resources to become available on the primary owner node.
  - ◆ `checking`. The virtual volume is being checked for validity.
  - ◆ `need_check`. The virtual volume has been checked, and an inconsistency has been found.
  - ◆ `need_ld`. One of the underlying logical disks is not in the normal state. After all logical disks are in the normal state, the status changes to `auto_check` and then to `checking`.
- `AdmMB`. The snapshot administrator space in megabytes.
  - `SnapMB`. The snapshot space in megabytes.
  - `userMB`. Your user space in megabytes.
  - `VSize`. Virtual size of volume in MB (1024^2 bytes).

The following example displays detailed information about the system's virtual volumes:

```
cli% showvv -d
Id  Name Rd  Mstr Prnt  Roch Rwch PPrnt PBlkRemain -----VV_WWN----- -----CreationTime-----
0  admin RW 1/-/-  ---  ---  ---  ---  - 50002AC0000001A8 Thu May 24 20:22:07 PDT 2007
7   vv1 RW 1/0/-  ---  ---  ---  ---  - 50002AC0000701A8 Tue Aug 28 16:41:01 PDT 2007
8   vv2 RW 0/1/-  ---  ---  ---  ---  - 50002AC0000801A8 Tue Aug 28 16:41:26 PDT 2007
9   vv3 RW 1/0/-  ---  ---  ---  ---  - 50002AC0000901A8 Tue Aug 28 16:41:37 PDT 2007
-----
4
```

The columns in the output above are identified as follows:

- `Id`. The virtual volume's ID.

- **Name.** The virtual volume's name.
- **Rd.** The read/write permissions for the volume. Values can be RW (read and write) or RO (read only).
- **Mstr.** The master node for the volume responsible for assigning mapping zones.
- **Prnt.** The ID of the volume's parent in the snapshot tree.
- **Roch.** The ID of the read only child volume in the snapshot tree.
- **Rwch.** The ID of the read write child volume in the snapshot tree.
- **PPrnt.** The ID of the volume's physical parent volume.
- **PBlkRemain.** The number of remaining blocks to be copied from the parent volume (for physical copy destination).
- **VV\_WWN.** The volume's World Wide Name.
- **CreationTime.** The creation date and time of the volume.

The following example displays the distribution of space for each virtual volume:

```
cli% showvv -s
```

		---Adm---		-----Snp-----				-----Usr-----							
		--(MB)--		--(MB)--		-(% VSize)--		--(MB)--		-(% VSize)--		-----(MB)-----			
Id	Name	Prov	Type	Rsvd	Used	Rsvd	Used	Wrn	Lim	Rsvd	Used	Wrn	Lim	Tot_Rsvd	VSize
0	admin	full	base	0	0	0	0	0.0	--	10240	10240	100.0	--	10240	10240
13	BasketDome	cpvv	base	0	0	0	0	0.0	2 5	256	256	100.0	--	256	256
2 total				0	0	0	0	0.0	2 5	10496	10496	100.0		10496	10496

The columns in the previous example are identified as follows:

- **ID.** The ID of the virtual volume.
- **Name.** The name of the virtual volume.
- **Prov.** The provisioning for the VV. Can be one of the following:
  - ◆ **full.** Fully provisioned VV, either with no Snp (snapshot) space or with statically allocated Snp space.
  - ◆ **tpvv.** Thin Provisioned VV, with space for the base volume allocated from the Usr space that is associated with the UsrCPG. Snapshots allocate space from the Snp space associated with the SnpCPG (if any).
  - ◆ **cpvv.** Commonly Provisioned VV. The Usr space for this VV is fully provisioned and the Snp space is associated with a SnpCPG.

- ◆ **tpsd.** An old-style Thin Provisioned VV (created on a 2.2.4 release or earlier) where both the base VV and snapshot data are allocated from the **Snp** space associated with the **UsrCPG**.
- ◆ **snp.** Since the VV is Type **vcopy** (snapshot) it is the provisioning that is associated with the base VV.
- **Adm\_Rsvd\_MB.** Adm (snapshot admin or SA) reserved space in megabytes.
- **Adm\_Used\_MB.** Adm used space in megabytes.
- **Snp\_Rsvd\_MB.** Snp (snapshot data or SD) reserved space in megabytes.
- **Snp\_Used\_MB.** Snp used space in megabytes.
- **Snp\_Used\_Perc.** Snp used space as a percentage of the virtual size (**VSize\_MB**).
- **Warn\_Snp\_Perc.** Percentage of VV virtual size (**VSize\_MB**) at which the **Snp\_Used\_MB** will generate a warning alert.
- **Limit\_Snp\_Perc.** Percentage of VV virtual size (**VSize\_MB**) at which the **Snp\_Used\_MB** will be prevented from growing.
- **Usr\_Rsvd\_MB.** Usr (User) reserved space in megabytes.
- **Usr\_Used\_MB.** Usr used space in megabytes.
- **Usr\_Used\_Perc.** Usr used space as a percentage of the virtual size (**VSize\_MB**).
- **Warn\_Usr\_Perc.** Percentage of VV virtual size (**VSize\_MB**) at which the **Usr\_Used\_MB** will generate a warning alert.
- **Limit\_Usr\_Perc.** Percentage of VV virtual size (**VSize\_MB**) at which the **Usr\_Used\_MB** will be prevented from growing.
- **Tot\_Rsvd\_MB.** Total reserved space (**Adm\_Rsvd\_MB** + **Snp\_Rsvd\_MB** + **Usr\_Rsvd\_MB**)
- **Tot\_RawRsvd\_MB.** Total raw reserved space (**Adm\_RawRsvd\_MB** + **Snp\_RawRsvd\_MB** + **Usr\_RawRsvd\_MB**).

The following example displays policy information for virtual volume 0:

```
cli% showvv -p 0
Id  Name Domain                      Policies
  0  admin    -                        stale_ss,system
-----
  1
```

The columns in the previous example are identified as follows:

- **Id.** The virtual volume ID.
- **Name.** The virtual volume name.
- **Domain.** The domain to which the virtual volume belongs. Valid values are – or <domain\_name>. If the domain does not exist, – is displayed. The Domain column can be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable.
- **Policies.** The current policy assigned to the displayed virtual volume.

The following example displays the snapshot space information:

```
cli% showvv -s -p -prov tp* -host hname
```

The columns in the previous example are identified as follows:

- **ID.** The ID of the virtual volume.

The following example displays the detailed state information:

```
cli% showvv -state
Id Name  Prov Type State -Detailed_State-
  0 admin full base normal normal
-----
  1 total
```

The columns in the previous example are identified as follows:

- **ID.** The ID of the virtual volume.
- **Name.** The virtual volume name.
- **Prov.** The provisioning for the VV. Can be one of the following:
  - ◆ **full.** Fully provisioned VV, either with no Snp (snapshot) space or with statically allocated Snp space.
  - ◆ **tpvv.** Thin Provisioned VV, with space for the base volume allocated from the `Usr` space that is associated with the `UsrCPG`. Snapshots allocate space from the `Snp` space associated with the `SnpCPG` (if any).



- ◆ **cpvv.** Commonly Provisioned VV. The `Usr` space for this VV is fully provisioned and the `Snp` space is associated with a `SnpCPG`.
- ◆ **tpsd.** An old-style Thin Provisioned VV (created on a 2.2.4 release or earlier) where both the base VV and snapshot data are allocated from the `Snp` space associated with the `UsrCPG`.
- ◆ **–.** Since the VV is Type `vcopy` (snapshot) it is the provisioning that is associated with the base VV.
- **Type.** Indicates the copy type of virtual volume and can be one of the following:
  - ◆ **base.** Base volume (not a copy).
  - ◆ **pcopy.** Physical copy (full copy).
  - ◆ **vcopy.** Snapshot copy (virtual copy).
- **State.** The state of the virtual volume.
  - ◆ **normal.** The virtual volume is operating normally.
  - ◆ **failed.** The virtual volume is operating abnormally.
  - ◆ **degraded.** The virtual volume is degraded.
- **Detailed\_State.** The detailed state of the logical disk.
  - ◆ **lds\_not\_started.** The logical disks are not started.
  - ◆ **not\_started.** The virtual volumes are not started.
  - ◆ **needs\_check.** The logical disks needs to be checked for consistency.
  - ◆ **snapdata\_invalid.** The Snapdata is invalid.
  - ◆ **preserved.** One or more sets of a logical disk is unavailable due to missing chunklets and the remaining virtual volume data is preserved in a preserved data logical disk.
  - ◆ **stale.** Parts of the virtual volume contain old data because a copy-on-write operation failed.
  - ◆ **copy\_failed.** A promote or copy to this volume failed.
  - ◆ **degraded\_avail.** The availability of the virtual volume is degraded.
  - ◆ **degraded\_perf.** The performance of the virtual volume is degraded.
  - ◆ **promoting.** The volume is currently the target of a promote.

- ◆ `copy_target`. The volume is currently the target of a copy.
- ◆ `tuning`. The volume is currently tuning.
- ◆ `closing`. The volume is currently closing.
- ◆ `removing`. The volume is currently removing.
- ◆ `creating`. The volume is currently creating.
- ◆ `copy_source`. The volume is a copy source.
- ◆ `unknown`. The volume state is unknown.

The following example displays raw space usage information for thin provisioned VVs exported to host `hname` and all VVs in the tree with the base VV of ID 50.

cli% showvv -r -p -prov tp\* -host halfdome -p -baseid 11

----- (MB) -----												
		----Adm----		----Snp-----		----Usr-----		----Tot-----				
Id	Name	Prov	Type	RawRsvd	Rsvd	RawRsvd	Rsvd	RawRsvd	Rsvd	RawRsvd	Rsvd	VSize
11	HalfDome	full	base	0	0	0	0	512	256	512	256	256
-----												
1	total			0	0	0	0	512	256	512	256	256

The columns in the previous example are identified as follows:

- **ID**. The ID of the virtual volume.
- **Name**. The virtual volume name.
- **Prov**. The provisioning for the VV. Can be one of the following:
  - ◆ **full**. Fully provisioned VV, either with no Snp (snapshot) space or with statically allocated Snp space.
  - ◆ **tpvv**. Thin Provisioned VV, with space for the base volume allocated from the `Usr` space that is associated with the `UsrCPG`. Snapshots allocate space from the `Snp` space associated with the `SnpCPG` (if any).
  - ◆ **cpvv**. Commonly Provisioned VV. The `Usr` space for this VV is fully provisioned and the `Snp` space is associated with a `SnpCPG`.
  - ◆ **tpsd**. An old-style Thin Provisioned VV (created on a 2.2.4 release or earlier) where both the base VV and snapshot data are allocated from the `Snp` space associated with the `UsrCPG`.

- ◆ --. Since the VV is Type vcopy (snapshot) it is the provisioning that is associated with the base VV.
- Type. Indicates the copy type of virtual volume and can be one of the following:
  - ◆ base. Base volume (not a copy).
  - ◆ pcopy. Physical copy (full copy).
  - ◆ vcopy. Snapshot copy (virtual copy).
- Adm\_RawRsvd\_MB. Raw (including RAID overhead) Adm reserved space in MB (1024^2 bytes).
- Adm\_Rsvd\_MB. Adm (snapshot admin or SA) reserved space in MB Adm
- Snp\_RawRsvd\_MB. Raw (including RAID overhead) Snp reserved space in MB (1024^2 bytes).
- Snp\_Rsvd\_MB. Snp (snapshot data or SD) reserved space in MB (1024^2 bytes).
- User\_RawRsvd\_MB. Raw (including RAID overhead) User reserved space in MB (1024^2 bytes).
- User\_Rsvd\_MB. Usr (User) reserved space in MB (1024^2 bytes).
- Tot\_RawRsvd\_MB. Total raw reserved space (Adm\_RawRsvd\_MB + Snp\_RawRsvd\_MB + Usr\_RawRsvd\_MB).
- Tot\_Rsvd\_MB. Total reserved space (adm\_Rsvd\_MB + Snp\_Rsvd\_MB + Usr\_Rsvd\_MB).

The following example displays information on volume grown and reclaimed space using the `-showcols` option:

```
cli% showvv -showcols
Grown_Adm_MB,Grown_Snp_MB,Grown_Snp_SD_MB,Reclaimed_Adm_MB,Reclaimed_Snp_MB,Reclaimed_Snp_SD_MB,Name
tpvvgrp.?
Grown_Adm_MB Grown_Snp_MB Grown_Snp_SD_MB Reclaimed_Adm_MB Reclaimed_Snp_MB Reclaimed_Snp_SD_MB Name
0 259072 0 0 0 15616 tppvgrp.0
0 480256 60160 0 0 216064 tppvgrp.1
0 451584 66048 0 0 230400 tppvgrp.2
0 381952 18176 0 0 134272 tppvgrp.3
0 440320 57984 0 0 201984 tppvgrp.4
0 342016 6784 0 0 89600 tppvgrp.5
0 428032 5875712 0 0 18196608 tppvgrp.6
```

The columns in the previous example are identified as follows:

- Grown\_Adm\_MB. The amount of administrator space that the volume has grown.

- `Grown_Snp_MB`. The amount of snapshot space that the volume has grown.
- `Grown_Snp_SD_MB`. The amount of snapshot SD space that the volume has grown.
- `Reclaimed_Adm_MB`. The amount of administrator space that the volume has had reclaimed.
- `Reclaimed_Snp_MB`. The amount of snapshot space that the volume has had reclaimed.
- `Reclaimed_Snp_SD_MB`. The amount of snapshot SD space that the volume has had reclaimed.
- `Name`. The virtual volume name.

## NOTES

- The per-snapshot space is approximate and must be calculated using the `updatesnapspace` command.
- If the `-showcols` option is used, the full column names are shown in the header otherwise the header contains abbreviated column names.
- For `<cpgname_or_pattern>`, `<hostname_or_pattern>`, `<groupname_or_pattern>`, and `<domainname_or_pattern>`, the patterns are glob-style (shell-style) patterns (see `Help on sub,globpat`).
- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.
- The `showvv` command has new properties that provide statistics on the growth and reclamation of `usr`, `snp`, and `adm` space within a volume. These properties can be displayed using the `-showcols` option. The new properties, as shown by the `-listcols` option are `Grown_Adm_MB`, `Grown_Snp_MB`, `Grown_Snp_SD_MB`, `Reclaimed_Adm_MB`, `Reclaimed_Snp_MB`, or `Reclaimed_Snp_SD_MB`. These statistics are only available via the CLI and only via the `-showcols` option.
- For this command: 1MB = 1048576 bytes.

---

**COMMAND**

showvmap

**DESCRIPTION**

The `showvmap` command displays information about how Virtual Volume (VV) regions are mapped to logical disks.

**SYNTAX**

`showvmap <VV_name>`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

None.

**SPECIFIERS**

`<VV_name>`

Specifies the VV name, using up to 31 characters in length.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays information about how VV `VV1` is mapped:

```
cli% showvmap VV1
Space      Start(MB)  Length(MB)  LdId    LdName          LdOff(MB)
adm 0         0         256         25      VV1.adm.0        0
    1         256         256         24      VV1.adm.1        0
snp 0         0         224         23      VV1.snp.0        0
    1         224         192         22      VV1.snp.1        0
usr 0         0         256         27      VV1.usr.0        0
    1         256         256         26      VV1.usr.1        0
    2         512         256         27      VV1.usr.0        256
    3         768         256         26      VV1.usr.1        256
    4        1024         256         27      VV1.usr.0        512
    5        1280         256         26      VV1.usr.1        512
    6        1536         256         27      VV1.usr.0        768
    7        1792         256         26      VV1.usr.1        768
```

The columns in the previous example are identified as follows:

- `Space`. The type of the space (`usr`), snapshot data (`snp`), or snapshot administration (`adm`) for which the region is being used followed by the ID of the region. The number to the right of the space type are the chunklets for each space.
- `Start(MB)`. The offset from the beginning of the region, in MB.
- `Length(MB)`. The length, or size, of the region in MB.
- `LdId`. The ID of the logical disk that contains the region.
- `LdName`. The name of the logical disk that contains the region.
- `LdOff(MB)`. The offset from the beginning of the logical disk that contains the region, in MB.

## NOTES

None.

---

**COMMAND**

showvvpd

**DESCRIPTION**

The `showvvpd` command displays Virtual Volume (VV) distribution across Physical Disks (PDs).

**SYNTAX**

`showvvpd [option <arg>] {<VV_name|pattern>}...`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on the column number (<col>). Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting <dir> as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by the values in later columns.

**SPECIFIERS**

`<VV_name> | <pattern>`

Specifies the VV with the specified name (31 character maximum) or matches the glob-style pattern for which information is displayed. This specifier can be repeated to display configuration information about multiple VVs. This specifier is not required. If not specified, configuration information for all VVs in the system is displayed.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays information about multiple VVs. In this case multi.0, multi.1, multi.2 and multi.3 were created using the `-cnt 4` option. Using `multi.*`, the aggregate chunklets for all the physical disks is displayed:

```
cli% showvvpd multi.*
Id Cage_Pos SA SD usr total
0 0:0:0 0 0 0 0
1 0:0:1 0 0 0 0
2 0:0:2 0 0 0 0
3 0:0:3 0 0 0 0
4 0:1:0 0 0 0 0
5 0:1:1 0 0 0 0
6 0:1:2 0 0 0 0
7 0:1:3 0 0 0 0
8 0:2:0 0 0 0 0
9 0:2:1 0 0 0 0
10 0:2:2 0 0 0 0
11 0:2:3 0 0 0 0
12 0:3:0 0 0 0 0
13 0:3:1 0 0 0 0
14 0:3:2 0 0 0 0
15 0:3:3 0 0 0 0
16 1:0:0 0 0 4 4
17 1:0:1 0 0 4 4
18 1:0:2 0 0 4 4
19 1:0:3 0 0 4 4
20 1:1:0 0 0 4 4
21 1:1:1 0 0 4 4
22 1:1:2 0 0 4 4
23 1:1:3 0 0 4 4
24 1:2:0 0 0 4 4
25 1:2:1 0 0 4 4
26 1:2:2 0 0 4 4
27 1:2:3 0 0 4 4
28 1:3:0 0 0 4 4
29 1:3:1 0 0 4 4
30 1:3:2 0 0 4 4
31 1:3:3 0 0 4 4
32 2:0:0 0 0 4 4
33 2:0:1 0 0 4 4
34 2:0:2 0 0 4 4
35 2:0:3 0 0 4 4
36 2:1:0 0 0 4 4
37 2:1:1 0 0 4 4
38 2:1:2 0 0 4 4
39 2:1:3 0 0 4 4
40 2:2:0 0 0 4 4
41 2:2:1 0 0 4 4
42 2:2:2 0 0 4 4
43 2:2:3 0 0 4 4
44 2:3:0 0 0 4 4
45 2:3:1 0 0 4 4
46 2:3:2 0 0 4 4
47 2:3:3 0 0 4 4
-----
48 total 0 0 128 128
```



For the example above, if you only specified a single volume you would get the same result. That is because the four VVs are interleaved across the same logical disks and share the same chunklets.

The following example displays the distribution of space for a single VV (`multi.0`):

```
cli% showvvpd multi.0
Id Cage_Pos SA SD usr total
0 0:0:0 0 0 0 0
1 0:0:1 0 0 0 0
2 0:0:2 0 0 0 0
3 0:0:3 0 0 0 0
4 0:1:0 0 0 0 0
5 0:1:1 0 0 0 0
6 0:1:2 0 0 0 0
7 0:1:3 0 0 0 0
8 0:2:0 0 0 0 0
9 0:2:1 0 0 0 0
10 0:2:2 0 0 0 0
11 0:2:3 0 0 0 0
12 0:3:0 0 0 0 0
13 0:3:1 0 0 0 0
14 0:3:2 0 0 0 0
15 0:3:3 0 0 0 0
16 1:0:0 0 0 4 4
17 1:0:1 0 0 4 4
18 1:0:2 0 0 4 4
19 1:0:3 0 0 4 4
20 1:1:0 0 0 4 4
21 1:1:1 0 0 4 4
22 1:1:2 0 0 4 4
23 1:1:3 0 0 4 4
24 1:2:0 0 0 4 4
25 1:2:1 0 0 4 4
26 1:2:2 0 0 4 4
27 1:2:3 0 0 4 4
28 1:3:0 0 0 4 4
29 1:3:1 0 0 4 4
30 1:3:2 0 0 4 4
31 1:3:3 0 0 4 4
32 2:0:0 0 0 4 4
33 2:0:1 0 0 4 4
34 2:0:2 0 0 4 4
35 2:0:3 0 0 4 4
36 2:1:0 0 0 4 4
37 2:1:1 0 0 4 4
38 2:1:2 0 0 4 4
39 2:1:3 0 0 4 4
40 2:2:0 0 0 4 4
41 2:2:1 0 0 4 4
42 2:2:2 0 0 4 4
43 2:2:3 0 0 4 4
44 2:3:0 0 0 4 4
45 2:3:1 0 0 4 4
46 2:3:2 0 0 4 4
47 2:3:3 0 0 4 4
-----
48 total 0 0 128 128
```

The following example displays information for a specific column:

```
cli% showvvpd -sortcol 5,dec junk
Id Cage_Pos SA SD usr total
6 0:1:2 1 1 1 3
19 1:0:3 1 1 1 3
27 1:2:3 1 1 1 3
30 1:3:2 1 1 1 3
43 2:2:3 1 1 1 3
0 0:0:0 1 1 0 2
3 0:0:3 1 1 0 2
7 0:1:3 0 1 1 2
8 0:2:0 1 1 0 2
9 0:2:1 1 1 0 2
10 0:2:2 1 1 0 2
11 0:2:3 1 1 0 2
15 0:3:3 1 1 0 2
17 1:0:1 1 1 0 2
18 1:0:2 1 1 0 2
20 1:1:0 1 1 0 2
21 1:1:1 1 1 0 2
22 1:1:2 1 1 0 2
24 1:2:0 1 1 0 2
26 1:2:2 1 1 0 2
29 1:3:1 1 1 0 2
31 1:3:3 1 1 0 2
33 2:0:1 1 1 0 2
34 2:0:2 1 1 0 2
39 2:1:3 1 1 0 2
40 2:2:0 1 1 0 2
42 2:2:2 1 1 0 2
44 2:3:0 1 1 0 2
45 2:3:1 1 1 0 2
46 2:3:2 0 1 1 2
1 0:0:1 0 1 0 1
2 0:0:2 0 1 0 1
4 0:1:0 0 1 0 1
5 0:1:1 0 1 0 1
12 0:3:0 0 1 0 1
13 0:3:1 1 0 0 1
14 0:3:2 1 0 0 1
16 1:0:0 0 0 1 1
25 1:2:1 0 1 0 1
28 1:3:0 0 1 0 1
32 2:0:0 0 1 0 1
35 2:0:3 0 1 0 1
37 2:1:1 0 1 0 1
38 2:1:2 1 0 0 1
47 2:3:3 1 0 0 1
23 1:1:3 0 0 0 0
36 2:1:0 0 0 0 0
41 2:2:1 0 0 0 0
-----
48 total 32 40 8 80
```

## NOTES

- The command prints the number of chunklets in each physical disk that contain data for the Snapshot Admin (SA), Snapshot Data (SD), User space as well as the total for all spaces

for all the VVs that match the <VV\_name> or <pattern> list. This can be useful in determining how evenly the VV is striped across the disks.

- Not all the data in the chunklets is necessarily allocated exclusively to the selected VV. In some cases, only part of a chunklet could be data that is allocated for the VVs.

---

**COMMAND**

showvvset

**DESCRIPTION**

The `showvvset` command lists the Virtual Volume (VV) sets defined on the InServ and their members.

**SYNTAX**

```
showvvset [options] [<setname_or_pattern>...]  
showvvset -vv [options] [<vvname_or_pattern>...]
```

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-d

Show a more detailed listing of each set.

-vv

Show VV sets that contain the supplied vnames or patterns

**SPECIFIERS**

<setname\_or\_pattern>...

An optional list of setnames or patterns. If no <setname> or <pattern> is specified all sets are displayed, otherwise only sets with names matching one or more of the setnames or patterns are displayed. The patterns are glob-style patterns (see help on sub,globpat).

<vvname\_or\_pattern>...

Specifies that the sets containing virtual volumes with the specified names or matching the glob-style patterns should be displayed.

**RESTRICTIONS**

None.

## EXAMPLES

To show all VV sets defined to the system:

```
cli% showvvset
Id Name  Members
  0 oravv oravv.0
           oravv.1
           oravv.2
           oravv.3
           oravv.4
           oravv.5
           oravv.6
           oravv.7
           oravv.8
           oravv.9
 20 sia-1 test
           ttpvv.rw
           test-sv
           bar
```

Show the sia-1 set only, with detail:

```
cli% showvvset -d sia-1
20 sia-1 test      This set has a comment
           ttpvv.rw
           test-sv
           bar
```

Show VV sets containing VVs matching test\*:

```
cli% showvvset -vv test*
Id Name  Members
 20 sia-1 test
           ttpvv.rw
           test-sv
           bar
```

## NOTES

A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.



# 23

## Shutdown Commands

---

In this chapter

<code>shutdownnode</code>	<b>23.2</b>
<code>shutdownsys</code>	<b>23.4</b>

---

**COMMAND**

`shutdownnode`

**DESCRIPTION**

The `shutdownnode` command shuts down a system node.

**SYNTAX**

`shutdownnode halt|reboot [option] <node_ID>`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-f`

Forces the operation so that the command does not require confirmation before proceeding.

**SPECIFIERS**

`<node_ID>`

Specifies the node, identified by its ID, to be shut down.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example shuts down system node 0 and halts the restarting of the node:

```
cli% shutdownnode halt 0
```

**NOTES**

- When issuing the `shutdownnode` command without the `-f` option, the system manager executes a set of validation checks before proceeding with the shutdown.



- If any of the following conditions exists, the shutdown operation will not proceed:
  - ◆ The system software upgrade is in progress.
  - ◆ The target node is not online.
  - ◆ The system is processing tasks that should not be interrupted, such as `tunevv` and remote copy related tasks.
  - ◆ If any other node is online but not yet integrated into the cluster.
  - ◆ If another shutdown node operation is already in progress.
  - ◆ If the shutdown node operation will result in the system shutdown due to loss of quorum.
  - ◆ One or more orphaned logical disks exist on the system that cannot be preserved.
  - ◆ One or more admin LDs cannot be reset, resulting in the kernel being unable to access meta data from those LDs.
  - ◆ One or more data (user or snap) LDs cannot be reset, causing their associated VLUNs to become inaccessible to host applications.

---

**COMMAND**

`shutdownsys`

**DESCRIPTION**

The `shutdownsys` command shuts down an entire system.

**SYNTAX**

`shutdownsys halt|reboot`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**SUBCOMMANDS**

`halt`

Specifies that the system should be halted after shutdown. If this subcommand is not specified, the `reboot` subcommand must be used.

`reboot`

Specifies that the system should be restarted after shutdown. If this subcommand is not specified, the `halt` subcommand must be used.

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

- Access to all domains is required to run this command.
- Do not issue any commands other than `showsys` while the system is shutting down.

## EXAMPLES

The following example shuts down and then restarts the system:

```
cli% shutdownsys reboot

-----System Information-----
System Name      : maltp001
System Model     : InServ S800X
Serial Number    : 1000787

-----Partial List Of Hosts-----
Id Name          Persona      -WWN/iSCSI_Name- Port  IP_addr
0  sunx4150-01  Generic-legacy  2101001B32343495  6:5:1  n/a
1  sunx4150-02  Generic-legacy  2101001B32344695  6:5:1  n/a
2  sunx4150-03  Generic-legacy  2100001B32147595  6:5:1  n/a
3  sunx4150-04  Generic-legacy  2101001B32341C95  6:5:1  n/a
4  sunx4150-05  Generic-legacy  2101001B32342495  6:5:1  n/a
5  sunx4150-06  Generic-legacy  2100001B321A6E82  6:5:1  n/a
6  sunx4150-07  Generic-legacy  2101001B3231EE79  6:5:1  n/a
7  sunx4150-08  Generic-legacy  2101001B323BC44C  6:5:1  n/a
8  sunx4150-09  Generic-legacy  2101001B323BFF4C  6:5:1  n/a
9  sunx4150-10  Generic-legacy  2101001B323B0A4C  6:5:1  n/a

Shutting down this InServ will impact applications running on the hosts displayed
above

Do you REALLY want to REBOOT this InServ? yes or no: yes

The system will REBOOT in 15 seconds

Press the enter key to stop...

The system will REBOOT in 12 seconds

Press the enter key to stop...

The system will REBOOT in 9 seconds

Press the enter key to stop...

The system will REBOOT in 6 seconds

Press the enter key to stop...

The system will REBOOT in 3 seconds

Press the enter key to stop...

The system will REBOOT NOW !!!
```

## NOTES

- The execution of `shutdownsys` command can affect service. Hence, a confirmation is required before proceeding with this command.
- After the `shutdownsys` command is issued, there is no indication from the CLI that the shutdown is occurring. You can issue the `showsys` command ([shows](#) on page 22.171) to display the current status of the system during the initial stage of the shutdown process and after the system has fully restarted.
- If the node that was running on the system manager fails or if the system manager process exits while executing the `shutdownsys` command, the shutdown will not complete. The only safe action is to reissue the `shutdownsys` command.

# 24

## Start Commands

---

In this chapter

startcim	24.2
startld	24.3
startcopy	24.4
startcopygroup	24.5
starttask	24.7
startvv	24.8

---

**COMMAND**

`startcim`

**DESCRIPTION**

The `startcim` command starts the CIM server to service CIM requests. By default, the CIM server is not started until this command is issued.

**SYNTAX**

`startcim`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example starts the CIM server:

```
cli% startcim
CIM server will start in about 90 seconds.
```

**NOTES**

- By default, the CIM server is not started until this command is issued.
- Use `stopcim` to stop the CIM server.

---

**COMMAND**

startld

**DESCRIPTION**

The `startld` command starts data services on a Logical Disk (LD) that has not yet been started.

**SYNTAX**

`startld [option] <LD_name>`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-ovrd`

Specifies that the LD is forced to start, even if some underlying data is missing.

**SPECIFIERS**

`<LD_name>`

Specifies the LD name, using up to 31 characters.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example starts data services on LD `ld5`:

```
cli% startld ld5
```

**NOTES**

None.

---

**COMMAND**

`startrcopy`

**DESCRIPTION**

The `startrcopy` command starts the Remote Copy Service.

**SYNTAX**

`startrcopy`

**AUTHORITY**

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

None.

**SPECIFIERS**

None.

**RESTRICTIONS**

This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

**EXAMPLES**

The following example starts Remote Copy on a system:

```
cli% startrcopy
```

**NOTES**

The `startrcopy` command must be executed before any other Remote Copy command.



---

**COMMAND**

startrcopygroup

**DESCRIPTION**

The `startrcopygroup` command enables Remote Copy for the specified Remote Copy volume group.

**SYNTAX**

`startrcopygroup [options <arg>] <group_name>`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-nosync`

Prevents the initial synchronization and sets the virtual volumes to a synchronized state.

`-wait`

Specifies that the command waits until the initial synchronization is completed. The system generates an event when the synchronization is completed.

`-t <target_name>`

Only start the group on the specified target.

**SPECIFIERS**

`<group_name>`

The name of the Remote Copy volume group. The group name can be obtained using the `showrcopy` command.

**RESTRICTIONS**

This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

**EXAMPLES**

The following example starts Remote Copy for Group1:

```
cli% startrcopygroup Group1
```

## NOTES

- If a group's target has the `mirror_config` policy set and the group is a primary group, then this command is mirrored to that target and the corresponding secondary group is started. If the policy is set and the group is a secondary, then this command fails.
- If the `mirror_config` policy is not set, then the corresponding secondary group must already be started or this command fails.
- Synchronous groups are automatically synchronized when started. Asynchronous periodic volume groups are synchronized only on the first time they are started.
- You must enter this command on the secondary server before entering it on the primary if the `mirror_config` policy is not set.

---

**COMMAND**

`starttask`

**DESCRIPTION**

The `starttask` command provides users with the ability to execute commands with long running times. The commands run in the background.

**SYNTAX**

`starttask <command>`

**AUTHORITY**

Super, Service

**OPTIONS**

None.

**SPECIFIERS**

`<command>`

Specifies the command to be run by the `starttask` command.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example shows the `removevv` command to be started with the forced option:

```
cli% starttask removevv -f vv1
```

**NOTES**

- If the command to be run requires confirmation it must be forced with the `-f` option otherwise it will fail.
- The commands that can be run with the `starttask` command are: `admithw`, `creategroupsv`, `createsv`, `moverelocpd`, `removevv`, `updatevv`, `upgradecage`, `upgradepd`.

---

**COMMAND**

startvv

**DESCRIPTION**

The `startvv` command starts data services on a Virtual Volume (VV) that has not yet been started.

**SYNTAX**

`startvv [option] <VV_name>`

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

`-ovrd`

Specifies that the logical disk is forced to start, even if some underlying data is missing.

**SPECIFIERS**

`<VV_name>`

Specifies the VV name, using up to 31 characters.

**RESTRICTIONS**

Access to all domains is required to run this command.

**EXAMPLES**

The following example starts data services on virtual volume `testvv`:

```
cli% startvv testvv
```

**NOTES**

None.

# 25

## Stat Commands

---

In this chapter

statch	25.2
statcmp	25.7
statcpu	25.10
statisci	25.12
statiscisession	25.16
statld	25.18
statlink	25.23
statpd	25.25
statport	25.33
statrcopy	25.39
statvlun	25.42
statvv	25.48

---

**COMMAND**

statch

**DESCRIPTION**

The `statch` command displays chunklet statistics in a timed loop.

**SYNTAX**

`statch [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-rw`

Specifies that reads and writes are displayed separately. If this option is not used, then the total of reads plus writes is displayed.

`-d <secs>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of 2 seconds. This option and argument are not required in the command line.

`-iter <number>`

Specifies that chunklet statistics are displayed a specified number of times as indicated by the `number` argument using an integer from 1 through 2147483647.

`-begin`

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

`-idlep`

Specifies the percent of idle columns in the output.

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on the column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting `<dir>` can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by values in later columns.

`-filt <fspec>`

Specifies that statistics that fall below the threshold as specified by the `<fspec>` arguments, where `<fspec>` is `<type>`, `<op>`, `<meas>`, `<val>`, are filtered out and not displayed. The `<type>`, `<op>`, `<meas>`, and `<val>` arguments are separated with one comma.

`<type>`

The `type` argument can be specified with one of the following arguments:

`curs`

Specifies that only devices with current statistics above the threshold are displayed.

`avgs`

Specifies that only devices with average statistics above the threshold are displayed.

`maxs`

Specifies that only devices with maximum values above the threshold are displayed.

`<op>`

The `operation` argument can be specified with one of the following arguments:

`r`

Specifies that read-only statistics are displayed.

`w`

Specifies that write-only statistics are displayed.

`t|rw`

Specifies that statistics for read and write totals are displayed.

<meas>

The meas argument can be specified with one of the following arguments:

iops

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

bw

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

svct

Specifies that statistics for service time in milliseconds are displayed.

size

Specifies that statistics for I/O operations in bytes are displayed.

<val> [k|K] | [m|M] | [g|G]

Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option in use is `-filt curs,r,iops,10k` that indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

-ld <LD\_name> | -ch <chunk\_num>

-ld <LD\_name>

Specifies that statistics are restricted to chunklets from a particular logical disk.

-ch <chunk\_num>

Specifies that statistics are restricted to a particular chunklet number.

## SPECIFIERS

None.



## RESTRICTIONS

None.

## EXAMPLES

The following example displays the collection of statistics:

```
cli% statch -iter 1
12:47:54 04/06/06 r/w I/O per second KBytes per sec Svt ms IOSz KB Util %
Ldid Ldname LdCh Pdid PdCh Cur Avg Max Cur Avg Max Cur Avg Cur Avg Qlen Cur Avg
-----
total t 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
```

For the previous example, before the `statch` command was issued, the `setstatch start` command was issued for chunklets 0 and 2 on logical disk `vv0.usr.1` and for chunklets 0 and 1 on logical disk `vv0.usr.2`.

The columns in the example above are identified as follows:

- `Ldid`. The logical disk ID.
- `Ldname`. The logical disk name.
- `LdCh`. The number of LD chunklets.
- `Pdid`. The physical disk ID.
- `r/w`. The I/O type. Values can be read (r), write (w), or read and write (t).
- `I/O per second Cur`. The current number of I/O per second.
- `I/O per second Avg`. The average number of I/O per second.
- `KBytes per Max`. The maximum number of KB per second.
- `KBytes per Cur`. The current number of KB per second.
- `KBytes per Avg`. The average number of KB per second.
- `Svt ms Max`. The maximum service time in milliseconds.
- `Svt ms Cur`. The current service time in milliseconds.
- `Svt ms Avg`. The average service time in milliseconds.
- `IOSz KB Cur`. The current I/O size in KB.
- `IOSz KB Avg`. The average I/O size in KB.

- `Qlen`. The queue length.
- `Util % Cur`. The percentage of current use.
- `Util % Avg`. The percentage of average use.

## NOTES

- For this command: KB = 1000 bytes.
- If no options are used, the command defaults to show reads, writes, and totals separately for all chunklets in intervals of 2 seconds.
- The `statch` command can only be used after the `setstatch` command has been issued to enable chunklet statistics. See [page 21.72](#) for information about the `setstatch` command.

---

**COMMAND**

statcmp

**DESCRIPTION**

The `statcmp` command displays Cache Memory Page (CMP) statistics by node or by Virtual Volume (VV).

**SYNTAX**

`statcmp [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-v`

Specifies that CMP statistics by VV instead of by node are displayed. By default, all VVs are displayed unless the `-n` option is specified.

`-n <name|pattern>`

Specifies that statistics are displayed for VVs matching the specified name or pattern. This option is valid only when used with the `-v` option.

`-domain {<domainname|pattern>}...`

Shows VVs that are in domains with names that match one or more of the specified domains or patterns. If `-domain` is not specified, the VVs that are in the current domain are shown. See the `currentdomain` parameter in the `setclienv` and `showclienv` commands. This option is only valid if the `-v` option is also specified.

`-d <seconds>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that CMP statistics are displayed a specified number of times as indicated by the `num` argument using an integer from 1 through 2147483647.

**SPECIFIERS**

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays one iteration of CMP statistics for all nodes:

```
cli% statcmp -iter 1
12:14:49
```

		Current			Total		
Node	Type	Accesses	Hits	Hit%	Accesses	Hits	Hit%
0	Read	0	0	0	0	0	0
0	Write	0	0	0	0	0	0
1	Read	0	0	0	0	0	0
1	Write	0	0	0	0	0	0

  

```
Page Stats
```

Node	Free	Clean	Write1	WriteN	WrtSched	Writing	Recov	RecFlush	LockBlk
0	50225	7537	0	0	0	0	0	0	0
1	52431	5318	0	0	0	0	0	0	0

  

```
Temporary and Page Credits
```

Node	Node0	Node1	Node2	Node3	Node4	Node5	Node6	Node7
0	0	1877	---	---	---	---	---	---
1	1892	0	---	---	---	---	---	---

Press the enter key to stop...

The columns in the example above are identified as follows:

- Node. Node ID on the InServ.
- Type. Read or Write.
- Access. Number of Current and Total Read/Write I/Os.
- Hits. Number of Read/Write I/Os in which data was already in cache.
- Hit%. Hits divided accesses displayed in percentages.
- Free. Number of cache pages without valid data on them.
- Clean. Number of clean cache pages (valid data on page). A page is clean when data in cache matches data on disk.
- Write1. Number of dirty pages that have been modified exactly 1 time. A page is dirty when it has been modified in cache but not written to disk.
- WriteN. Number of dirty pages that have been modified more than 1 time.

- **WrtSched.** Number of pages scheduled to be written to disk.
- **Writing.** Number of pages being currently written by the flusher to disk.
- **RecFlush.** During node down, number of recovered pages currently being written by the flusher to disk.
- **LockBlk.** Number of pages being modified by host I/O that are temporarily blocked because page is being written to disk by the flusher.

## NOTES

- Patterns are specified as regular expressions. See `Help sub, regexpat` for details.
- Temporary and Page Credits refer to the number of credits being given by each node to other nodes in the system. The credits plus pages (free, clean, scheduled and writing) should add up to the total amount of memory in the node.

---

**COMMAND**

statcpu

**DESCRIPTION**

The statcpu command displays CPU statistics for all nodes.

**SYNTAX**

statcpu [options <arg>]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-d <secs>

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

-iter <number>

Specifies that CPU statistics are displayed a specified number of times as indicated by the number argument using an integer from 1 through 2147483647.

-t

Show only the totals for all the CPUs on each node.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays two iterations of CPU statistics for all nodes:

```
cli% statcpu -iter 2
15:11:03
node,cpu  user sys  idle intr/s ctxt/s
0,0        0   0  100
0,1        0   0  100
0,total    0   0  100      162    412

1,0        0   1   99
1,1        0   1   99
1,total    0   1   99      158    269
Press the enter key to stop...

15:11:05
node,cpu  user sys  idle intr/s ctxt/s
0,0        0   0  100
0,1       19   7   74
0,total   10   4   87      183    444

1,0        0   0  100
1,1        0   0  100
1,total    0   0  100      158    235
Press the enter key to stop...
```

## NOTES

None.

---

**COMMAND**

statiscsi

**DESCRIPTION**

The `statiscsi` command displays the iSCSI statistics.

**SYNTAX**

`statiscsi [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d <secs>`

Looping delay in seconds `<secs>`. The default is 2.

`-iter <number>`

The command stops after a user-defined `<number>` of iterations.

`-nodes <nodelist>`

List of nodes for which the ports are included.

`-slots <slotlist>`

List of PCI slots for which the ports are included.

`-ports <port_list>`

List of ports for which the ports are included. Lists are specified in a comma-separated manner such as: `-slots 0,1,2` or `-slots 0`.

`-counts`

Shows the counts. The default is to show counts/sec.

`-fullcounts`

Show the values for the full list of counters instead of the default packets and KBytes for the specified protocols. The values are shown in three columns:

- ◆ Current - Counts since the last sample.
- ◆ CmdStart - Counts since the start of the command.
- ◆ Begin - Counts since the port was reset.



This option cannot be used with the `-prot` option. If the `-fullcounts` option is not specified, the metrics from the start of the command are displayed.

`-prot <prot>[, <prot>, ...]`

Shows the statistics for the specified protocols. The available protocols are:

- ◆ Eth - Ethernet.
- ◆ IP - Internet Protocol (IP).
- ◆ TCP - Transmission Control Protocol (TCP).
- ◆ iSCSI - iSCSI.
- ◆ all - All protocols (default).

This option cannot be used with the `-fullcounts` option.

`-prev`

Shows the differences from the previous sample.

`-begin`

Shows the values from when the system was last initiated.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays basic iSCSI statistics collection:

```
cli% statiscsi
16:37:59 04/06/06 ----Receive---- ---Transmit---- -----Total-----
port  Protocol  Pkts/s KBytes/s  Pkts/s KBytes/s  Pkts/s KBytes/s Errs/s
0:4:1      Eth 15633.2   1116.7  30115.9 45164.4  45749.1 46281.1   0.0
0:4:1      IP 15632.7    522.6  30116.4 44020.8  45749.1 44543.3   0.0
0:4:1      TCP 15632.7     22.3  30116.4 43057.0  45749.1 43079.4   0.0
0:4:1      iSCSI 547.7      0.0   1477.3 50452.8   2025.0 50452.8   0.0
0:4:2      Eth 0.5        0.0    0.0     0.0      0.5     0.0       0.0
0:4:2      IP 15632.1   522.6  30115.2 44019.0  45747.2 44541.5   0.0
0:4:2      TCP 15632.1     22.3  30115.2 43055.3  45747.2 43077.6   0.0
0:4:2      iSCSI 547.7      0.0   1478.7 50509.7   2026.4 50509.7   0.0
1:3:1      Eth 11307.2    807.0  21348.2 32048.1  32655.4 32855.1   0.0
1:3:1      IP 11309.2    377.4  21345.7 31233.2  32654.9 31610.5   0.0
1:3:1      TCP 11309.2     15.5  21344.7 30548.7  32653.9 30564.1   0.0
1:3:1      iSCSI 459.3      0.0   1247.9 42630.9   1707.2 42630.9   0.0
1:3:2      Eth 0.5        0.0    0.0     0.0      0.5     0.0       0.0
1:3:2      IP 11308.5   377.3  21344.5 31231.3  32653.0 31608.7   0.0
1:3:2      TCP 11308.5     15.5  21343.5 30546.9  32652.0 30562.3   0.0
1:3:2      iSCSI 459.3      0.0   1247.8 42628.4   1707.1 42628.4   0.0
-----
Total      Eth 26941.4   1923.7  51464.1 77212.5  78405.5 79136.2   0.0
Total      IP 53882.5   1799.8 102921.7 150504.2 156804.2 152304.1   0.0
Total      TCP 53882.5    75.6 102919.7 147207.9 156802.2 147283.5   0.0
Total      iSCSI 2014.0    0.0  5451.7 186221.9   7465.7 186221.9   0.0
Press the enter key to stop...
```

The following example displays one iteration of iSCSI counts (with Ethernet and transmission control protocols):

```
cli% statiscsi -iter 1 -counts -prot Eth,TCP -begin
13:24:38 03/29/06 -----From last port reset-----
-----Receive----- ---Transmit--- -----Total-----
port  Protocol  Pkts  KBytes  Pkts  KBytes  Pkts  KBytes Errs
1:3:1      Eth 766373.0 1066175.9 394473.0 55171.4 1160846.0 1121347.2 1.0
1:3:1      TCP 722917.0 1016532.4 394441.0 30238.3 1117358.0 1046770.7 0.0
1:3:2      Eth 0.0      0.0      0.0      0.0      0.0      0.0       0.0
1:3:2      TCP 722917.0 1016532.4 394441.0 30238.3 1117358.0 1046770.7 0.0
-----
Total      Eth 766373.0 1066175.9 394473.0 55171.4 1160846.0 1121347.2 1.0
Total      TCP 1445834.0 2033064.8 788882.0 60476.6 2234716.0 2093541.4 0.0*
```

**NOTES**

The statistics for the TCP, IP, and iSCSI protocols listed are the combined values for both ports of the HBA. Each port reports the combined stats for both ports for these protocols. The total reported at the bottom for these protocols is therefore twice the actual values.

---

**COMMAND**

`statiscsisession`

**DESCRIPTION**

The `statiscsisession` command displays the iSCSI session statistics.

**SYNTAX**

`statiscsisession [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d <secs>`

Looping delay in seconds `<secs>`. The default is 2.

`-iter <number>`

The command stops after a user-defined number of iterations.

`-nodes <nodelist>`

List of nodes for which the ports are included.

`-slots <slotlist>`

List of PCI slots for which the ports are included.

`-ports <portlist>`

List of port slots for which the ports are included. Lists are specified in a comma-separated manner such as: `-slots 0,1,2` or `-slots 0`.

`-counts`

Shows the counts. The default is to show counts/sec.

By default, the differences from initiating the command are shown. The following options change that behavior:

`-prev`

Shows the differences from the previous sample.

`-begin`

Shows the values from when the system was last initiated.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example displays the session statistics for the iSCSI:

```
cli% statiscsisession
15:31:35 04/04/06 --From start of statiscsisession command--
-----PDUs/s----- --KBytes/s--- ----Errs/s----
port -----iSCI_Name----- TPGT Cmd Resp Total Tx Rx Total Digest
TimeOut
1:3:1 ign.1991-05.com.microsoft:dt-ashok-xp.hq.3pardata.com 131 1.0 1.0 2.0 4.0 0.0 4.0 0.0 0.0
-----
Total - - 1.0 1.0 2.0 4.0 0.0 4.0 0.0 0.0
Press the enter key to stop...
```

**NOTES**

None.

---

**COMMAND**

statld

**DESCRIPTION**

The `statld` command displays read/write (I/O) statistics about Logical Disks (LDs) in a timed loop.

**SYNTAX**

`statld [options <arg>] [<LD_name|pattern>...]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-vv {<VV_name|pattern>}...`

Show only LDs that are mapped to Virtual Volumes (VVs) with names matching any of names or patterns specified. Multiple volumes or patterns can be repeated using a comma-separated list (for example `-vv <VV_name>,<VV_name>...`).

`-domain {<domain_name|pattern>}...`

Shows only LDs that are in domains with names matching any of the names or specified patterns. Multiple domain names or patterns can be repeated using a comma-separated list (for example `-domain <domain_name>,<domain_name>...`).

`-rw`

Specifies that reads and writes are displayed separately. If this option is not used, then the total of reads plus writes is displayed.

`-d <sec>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that I/O statistics are displayed a specified number of times as indicated by the `number` argument using an integer from 1 through 2147483647.

**-begin**

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

**-idlep**

Specifies the percent of idle columns in the output.

**-sortcol <col> [ ,<dir>][:<col>[ ,<dir>]...]**

Sorts command output based on the column number <col>. Columns are numbered from left to right, beginning with 0. You must specify a column number. In addition, you can specify the direction of sorting <dir> as follows:

**inc**

Sort in increasing order (default).

**dec**

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:), Rows with have the same information in them as earlier columns will be sorted by the values in the later columns.

**-filt <fspec>**

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The type, operation, meas, and value arguments are separated with one comma.

**<type>**

The type argument can be specified with one of the following arguments:

**curs**

Specifies that only devices with current statistics above the threshold are displayed.

**avgs**

Specifies that only devices with average statistics above the threshold are displayed.

**maxs**

Specifies that only devices with maximum values above the threshold are displayed.

<op>

The operation argument can be specified with one of the following arguments:

`r`

Specifies that read-only statistics are displayed.

`w`

Specifies that write-only statistics are displayed.

`t|rw`

Specifies that statistics for read and write totals are displayed.

<meas>

The meas argument can be specified with one of the following arguments:

`iops`

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `val` argument.

`bw`

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

`svct`

Specifies that statistics for service time in milliseconds are displayed.

`size`

Specifies that statistics for I/O operations in bytes are displayed.

<val> [`k|K`] [`m|M`] [`g|G`]

Specifies the minimum threshold using any integer. The integer can be optionally followed with `k` or `K` to indicate a multiple of 1,000, `m` or `M` to indicate a multiple of 1,000,000, or `g` or `G` to indicate a multiple of 1,000,000,000.

An example of this option in use is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

`-ni`

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.



## SPECIFIERS

<LD\_name|pattern>...

Only statistics are displayed for the specified LD or pattern. Multiple LDs or patterns can be repeated (for example <LD\_name> <LD\_name>...).

## RESTRICTIONS

None.

## EXAMPLES

The following example displays one iteration of I/O statistics for all LDs:

```
cli% statld -iter 1
13:03:04 04/06/06 r/w I/O per second KBytes per sec Svt ms IOSz KB Util %
      Ldname      Cur  Avg  Max  Cur  Avg  Max  Cur  Avg  Cur  Avg  Qlen  Cur  Avg
      log0.0      t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      pdsld0.0    t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      admin.usr.0 t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      admin.usr.1 t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      Test.usr.0  t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      Test.usr.1  t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      log1.0      t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      pdsld1.0    t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      admin.usr.2 t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      admin.usr.3 t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      Test.usr.2  t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
      Test.usr.3  t    0    0    0    0    0    0  0.0  0.0  0.0  0.0    0    0    0
-----
      total      t    0    0          0    0          0.0  0.0  0.0  0.0    0    0    0
```

The columns in the example above are identified as follows:

- Ldname. The logical disk name.
- r/w. The I/O type. Values can be read (r), write (w), or read and write (t).
- I/O per second Cur. The current number of I/O per second.
- I/O per second Avg. The average number of I/O per second.
- KBytes per Max. The maximum number of KB per second.
- KBytes per Cur. The current number of KB per second.
- KBytes per Avg. The average number of KB per second.

- `Svt ms Max.` The maximum service time in milliseconds.
- `Svt ms Cur.` The current service time in milliseconds.
- `Svt ms Avg.` The average service time in milliseconds.
- `IOSz KB Cur.` The current I/O size in KB.
- `IOSz KB Avg.` The average I/O size in KB.
- `Qlen.` The queue length.
- `Util % Cur.` The percentage of current use.
- `Util % Avg.` The percentage of average use.

## NOTES

- For this command: KB = 1000 bytes.
- If no option is specified, the command defaults to display statistics totals.
- If the `<LD_name>` or `<pattern>` specifier is used, then LDs with names matching any of the patterns are listed, otherwise all LDs are listed. These patterns are glob-style patterns (see Help on `sub,globpat`).

---

**COMMAND**

statlink

**DESCRIPTION**

The statlink command displays statistics for link utilization for all nodes in a timed loop.

**SYNTAX**

statlink [options <arg>]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-d <sec>

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

-iter <number>

Specifies that internode link statistics are displayed a specified number of times as indicated by the number argument using an integer from 1 through 2147483647.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

## EXAMPLES

The following example displays one iteration of statistics for the internode link:

```
cli% statlink -iter 1
11:37:28 03/11/08          XCB_sent_per_second KBytes_per_second XCBSz_KB
Node      Q ToNode  Cur   Avg      Max   Cur   Avg   Max Cur   Avg
0         CM      0 15134 14911   15134 61951 61038 61951 4.1  4.1
0        PCI0     0 22336 21852   22336 79831 77128 79831 3.6  3.5
0        PCI1     1 15312 12103   19805 62106 63198 77324 4.0  4.1
0         L0      1 10339 10307   10339 69673 69384 69673 6.7  6.7

1         CM      1 16372 16189   16372 67020 66269 67020 4.1  4.1
1        PCI0     1 18384 17899   18384 75285 73305 75285 4.1  4.1
1        PCI1     1  5329  5336    5344   984   985   987 0.2  0.2
1         L3      0 10326 10295   10326 69346 69150 69346 6.7  6.7
```

## NOTES

None.

---

**COMMAND**

statpd

**DESCRIPTION**

The statpd command displays the read/write (I/O) statistics for physical disks in a timed loop.

**SYNTAX**

statpd [options]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-w <WWN>

Specifies that statistics for a particular Physical Disk (PD) identified by World Wide Names (WWNs) are displayed.

-nodes <node\_list>

Specifies that the display is limited to specified nodes and PDs connected to those nodes. The node\_list argument is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the node\_list argument is not specified, all disks on all nodes are displayed.

-slots <slot\_list>

Specifies that the display is limited to specified PCI slots and PDs connected to those PCI slots. The slot\_list argument is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the slot\_list argument is not specified, all disks on all slots are displayed.

-ports <port\_list>

Specifies that the display is limited to specified ports and PDs connected to those ports. The port\_list argument is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the port\_list argument is not specified, all disks on all ports are displayed.

-devinfo

Indicates the device disk type and speed.

`-rw`

Specifies that reads and writes are displayed separately. If this option is not used, then the total of reads plus writes is displayed.

`-d <secs>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-begin`

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

`-idlep`

Specifies the percent of idle columns in the output.

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on the column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:), Rows that have the same information in them as earlier columns will be sorted by the values in the later columns.

`-filt <fspec>`

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The type, operation, meas, and value arguments are separated with one comma.

<type>

The `type` argument can be specified with one of the following arguments:

`curs`

Specifies that only devices with current statistics above the threshold are displayed.

`avgs`

Specifies that only devices with average statistics above the threshold are displayed.

`maxs`

Specifies that only devices with maximum values above the threshold are displayed.

<op>

The `operation` argument can be specified with one of the following arguments:

`r`

Specifies that read-only statistics are displayed.

`w`

Specifies that write-only statistics are displayed.

`t|rw`

Specifies that statistics for read and write totals are displayed.

<meas>

The `meas` argument can be specified with one of the following arguments:

`iops`

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `val` argument.

`bw`

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `val` argument.

`svct`

Specifies that statistics for service time in milliseconds are displayed.

`size`

Specifies that statistics for I/O operations in bytes are displayed.

`<val> [k|K] | [m|M] | [g|G]`

Specifies the minimum threshold using any integer. The integer can be optionally followed with `k` or `K` to indicate a multiple of 1,000, `m` or `M` to indicate a multiple of 1,000,000, or `g` or `G` to indicate a multiple of 1,000,000,000.

An example of this option in use is `-filt curs,r,iops,10k` that indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

`-ni`

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

`-p <pattern>`

Specifies that statistics for a pattern of PDs are displayed. Patterns are used to filter and select the disks from which the statistics are collected. If specified multiple times, each instance of the specified pattern adds additional candidate disks matching the pattern:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (`item`). Multiple nodes are separated with a single comma (`1,2,3`). A range of nodes is separated with a hyphen (`0-7`). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (`item`). Multiple slots are separated with a single comma (`1,2,3`). A range of slots is separated with a hyphen (`0-5`). The primary path of the disks must be on the specified PCI slot(s).



`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (<item>). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0-4). The primary path of the disks must be on the specified port(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (<item>). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0-3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the `CagePos` column of `showpd` output indicating the side of the cage is omitted when using `-mg` option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0-7). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0-3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0-3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays one iteration of I/O statistics for all PDs:

```
cli% statpd -iter 1
16:03:44 04/06/06 r/w I/O per second KBytes per sec Svt ms IOSz KB Util %
Pdid Port Cur Avg Max Cur Avg Max Cur Avg Cur Avg Qlen Cur Avg
0 1:0:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
1 0:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
2 1:0:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
4 1:0:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
5 0:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
6 1:0:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
7 0:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
8 1:0:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
9 0:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
10 1:0:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
11 0:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
12 1:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
13 0:2:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
14 1:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
15 0:2:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
16 1:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
17 0:2:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
18 1:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
19 0:2:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
20 1:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
21 0:2:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
22 1:0:2 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
23 0:2:1 t 0 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
-----
total t 0 0 0 0 0 0 0.0 0.0 0.0 0.0 0 0 0
```

The columns in the example above are identified as follows:

- Pdid. The physical disk ID.
- Port. The disk's port.
- r/w. The I/O type. Values can be read (r), write (w), or read and write (t).
- I/O per second Cur. The current number of I/O per second.
- I/O per second Avg. The average number of I/O per second.

- KBytes per Max. The maximum number of KB per second.
- KBytes per Cur. The current number of KB per second.
- KBytes per Avg. The average number of KB per second.
- Svt ms Max. The maximum service time in milliseconds.
- Svt ms Cur. The current service time in milliseconds.
- Svt ms Avg. The average service time in milliseconds.
- IOSz KB Cur. The current I/O size in KB.
- IOSz KB Avg. The average I/O size in KB.
- Qlen. The queue length.
- Util % Cur. The percentage of current use.
- Util % Avg. The percentage of average use.

## NOTES

For this command: KB = 1000 bytes.

---

**COMMAND**

statport

**DESCRIPTION**

The statport command displays read/write (I/O) statistics for ports.

**SYNTAX**

statport [options]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-both|-ctl|-data

Show data transfers only (-data), control transfers only (-ctl), or both data and control transfers (-both). If no option is included on the command line, the command shows data transfers only.

-nodes <node\_list>

Specifies that the display is limited to specified nodes and Physical Disks (PDs) connected to those nodes. The node\_list is specified as a series of integers separated by commas (for example 1,2,3). The list can also consist of a single integer (for example 1). If the node list is not specified, all disks on all nodes are displayed.

-slots <slot\_list>

Specifies that the display is limited to specified PCI slots and PDs connected to those PCI slots. The slot\_list is specified as a series of integers separated by commas (for example 1,2,3). The list can also consist of a single integer (for example 1). If the slot list is not specified, all disks on all slots are displayed.

-ports <port\_list>

Specifies that the display is limited to specified ports and PDs connected to those ports. The port\_list is specified as a series of integers separated by commas (for example 1,2,3). The list can also consist of a single integer (for example 1). If the port list is not specified, all disks on all ports are displayed.

`-host|-disk|-rcfc`

Specifies to display only host ports (target ports), only disk ports (initiator ports), or only Fibre Channel Remote Copy configured ports. If no option is specified, all ports are displayed.

`-rcip`

Includes only statistics for Ethernet configured Remote Copy ports.

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Sets the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-begin`

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

`-idlep`

Specifies the percent of idle columns in the output.

`-sortcol <col> [, <dir>][:<col>[, <dir>]...]`

Sorts command output based on the column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows that have the same information in them as earlier columns will be sorted by the values in the later columns.

`-filt <fspec>`

Specifies that statistics that fall below the threshold as specified by the `<fspec>` arguments where `<fspec>` is `<type>`, `<op>`, `<meas>`, `<val>`, are filtered out and not displayed. The `-filt` option applies to data transfers only. The `<type>`, `<operation>`, `<meas>`, and `<value>` arguments are separated with one comma.

`<type>`

The `<type>` argument can be specified with one of the following arguments:

`curs`

Specifies that only devices with current statistics above the threshold are displayed.

`avgs`

Specifies that only devices with average statistics above the threshold are displayed.

`maxs`

Specifies that only devices with maximum values above the threshold are displayed.

`<op>`

The `<operation>` argument can be specified with one of the following arguments:

`r`

Specifies that read-only statistics are displayed.

`w`

Specifies that write-only statistics are displayed.

`t|rw`

Specifies that statistics for read and write totals are displayed.

`<meas>`

The `<meas>` argument can be specified with one of the following arguments:

`iops`

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

`bw`

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

`svct`

Specifies that statistics for service time in milliseconds are displayed.

`size`

Specifies that statistics for I/O operations in bytes are displayed.

`<val> [k|K] | [m|M] | [g|G]`

Specifies the minimum threshold using any integer. The integer can be optionally followed with `k` or `K` to indicate a multiple of 1,000, `m` or `M` to indicate a multiple of 1,000,000, or `g` or `G` to indicate a multiple of 1,000,000,000.

An example of this option in use is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

`-ni`

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

## SPECIFIERS

None.

## RESTRICTIONS

None.



## EXAMPLES

The following example displays one iteration of I/O statistics for all ports:

```
cli% statport -iter 1
10:38:56 09/14/09 r/w I/O per second KBytes per sec      Svt ms      IOSz KB
  Port      D/C      Cur  Avg  Max  Cur  Avg  Max  Cur  Avg  Cur  Avg Qlen
0:0:1      Data    t     2   2   2   34   34   34 13.0 13.0 17.4 17.4   0
0:0:2      Data    t     0   0   0    4    4    4  8.2  8.2  8.2  8.2   0
0:2:1      Data    t     0   0   0    0    0    0  0.0  0.0  0.0  0.0   0
0:2:2      Data    t     0   0   0    0    0    0  0.0  0.0  0.0  0.0   0
0:3:1      Data    t     0   0   0    0    0    0  0.0  0.0  0.0  0.0   0
0:3:2      Data    t     0   0   0    0    0    0  0.0  0.0  0.0  0.0   0
1:0:1      Data    t     0   0   0    4    4    4  9.7  9.7  8.2  8.2   0
1:0:2      Data    t     2   2   2   34   34   34 11.4 11.4 17.4 17.4   0
1:3:1      Data    t     0   0   0    0    0    0  0.0  0.0  0.0  0.0   0
1:3:2      Data    t     0   0   0    0    0    0  0.0  0.0  0.0  0.0   0
-----
      10      Data    t     5   5       76   76       11.6 11.6 15.6 15.6   0
```

The columns in the example above are identified as follows:

- Port. The port ID.
- D/C. The Data or Control transfers.
- r/w. The I/O type. Values can be read (r), write (w), or read and write (t).
- I/O per second Cur. The current number of I/O per second.
- I/O per second Avg. The average number of I/O per second.
- KBytes per Max. The maximum number of KB per second.
- KBytes per Cur. The current number of KB per second.
- KBytes per Avg. The average number of KB per second.
- Svt ms Max. The maximum service time in milliseconds.
- Svt ms Cur. The current service time in milliseconds.
- Svt ms Avg. The average service time in milliseconds.
- IOSz KB Cur. The current I/O size in KB.
- IOSz KB Avg. The average I/O size in KB.
- Qlen. The queue length.

- **Errs.** Number of errors on the port.
- **Drops.** Number of dropped packets.

**NOTES**

- For this command: KB = 1000 bytes.
- The `-filt` option applies only to data transfers and not control transfers.

---

**COMMAND**

statrcopy

**DESCRIPTION**

The `statrcopy` command displays statistics for Remote Copy links.

**SYNTAX**

`statrcopy [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-d <secs>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that I/O statistics are displayed a specified number of times as indicated by the `num` argument using an integer from 1 through 2147483647.

`-u k|m|g`

Displays statistics as kilobytes (k), megabytes (m), or gigabytes (g). If no unit is specified, the default is kilobytes.

`-hb`

Specifies that the heartbeat round-trip time of the links should be displayed in addition to the link throughput.

**SPECIFIERS**

None.

**RESTRICTIONS**

This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

## EXAMPLES

The following example shows statistics for sending links link0 and link1:

cli% statrcopy						
17:37:01 05/17/07						
Target	Node	Address	IPC	Total (KBytes)	-Throughput (KBytes per sec)-	
					Current	Average
amp1	0	10.100.33.96	RCs041	404761.15	4.95	4.95
amp1	1	10.101.33.96	RCs142	404661.63	3.30	3.30
amp1				809422.78	8.25	8.25
amp2	0	10.100.33.11	RCs037	86845920.00	24612.95	24612.95
amp2	1	10.101.33.11	RCs138	85299712.32	23217.30	23217.30
amp2				172145632.32	47830.26	47830.26
receive 0	receive		RCr039	983110.14	27.32	27.32
receive 1	receive		RCr140	823292.09	23.37	23.37
Receive				1806402.23	50.69	50.69
Send				172955055.10	47838.51	47838.51
Total				174761457.34	47889.20	47889.20
Press the enter key to stop...						

The columns in this example provide the following information:

- **Target.** The system name of the actual link process that is running.
- **Node.** The node that the link is running on.
- **Address.** The IP address of the remote node.
- **IPC.** The mode of the Remote Copy link (r) Receive or (s) Send.
- **Total KBytes.** The amount of data that has been transmitted since the link was started, in KB.
- **Throughput (KBs).** The throughput on the link after the last iteration.
- **Current.** The amount of current transmitted since the last statrcopy iteration, in KB.
- **Average.** The average speed of the data transmission, in KB/s.

**NOTES**

- This command repeats until directed to stop.
- Within the context of this command, KB is 1000 bytes, MB is 1000KB, and GB is 1000MB.
- The numbers displayed by this command might be somewhat less than those displayed with `statport`, as `statrcopy` output does not include TCP/IP overhead

---

**COMMAND**

statvlun

**DESCRIPTION**

The `statvlun` command displays statistics for Virtual Volumes (VVs) and Logical Unit Number (LUN) host attachments.

**SYNTAX**

`statvlun [options <arg>]`

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-domain <domain_name|pattern>...`

Shows only Virtual Volume Logical Unit Number (VLUNs) whose VVs are in domains with names that match one or more of the specified domain names or patterns. Multiple domain names or patterns can be repeated using a comma-separated list (for example `-domain <domain_name>,<domain_name>...`).

`-host <host_name|pattern>...`

Shows only VLUNs exported to the specified hosts or patterns. Multiple hosts or patterns can be repeated using a comma-separated list (for example `-host <host_name>,<host_name>...`).

`-v <VV_name|pattern>...`

Requests that only Logical Disks (LDs) mapped to VVs that match any of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list (for example `-vv <VV_name>,<VV_name>...`).

`-l <LUN|pattern>...`

Specifies that VLUNs with LUNs matching the specified LUN(s) or pattern(s) are displayed. Multiple LUNs or patterns can be repeated using a comma-separated list (for example `-l <LUN>,<LUN>...`).

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and Physical Disks (PDs) connected to those nodes. The `node_list` is specified as a series of integers separated by commas

(1, 2, 3). The list can also consist of a single integer (1). If the `node_list` is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and PDs connected to those PCI slots. The `slot_list` is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified port slots and PDs connected to those port slots. The `port_list` is specified as a series of integers separated by commas (1, 2, 3). The list can also consist of a single integer (1). If the port list is not specified, all disks on all ports are displayed.

`-lw`

Lists the host's World Wide Name (WWN) or iSCSI names. This is especially useful when multiple WWNs or iSCSI names belonging to the same host are visible on the same port.

`-domainsum`

Specifies that sums for VLUNs are grouped by domain in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a "-" host name.

`-vvsum`

Specifies that sums for VLUNs of the same VV are displayed.

`-hostsum`

Specifies that sums for VLUNs are grouped by host in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a nameless host.

`-rw`

Specifies reads and writes to be displayed separately. If this option is not used, then the total of reads plus writes is displayed.

`-d <secs>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-begin`

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

`-idlep`

Includes a percent idle columns in the output.

`-sortcol <col>[, <dir>][:<col>[, <dir>]...]`

Sorts command output based on the column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by the values in the later columns.

`-filt <fspec>`

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

`<type>`

The <type> argument can be specified with one of the following arguments:

`curs`

Specifies that only devices with current statistics above the threshold are displayed.

`avgs`

Specifies that only devices with average statistics above the threshold are displayed.



`maxs`

Specifies that only devices with maximum values above the threshold are displayed.

`<op>`

The `<operation>` argument can be specified with one of the following arguments:

`r`

Specifies that read-only statistics are displayed.

`w`

Specifies that write-only statistics are displayed.

`t|rw`

Specifies that statistics for read and write totals are displayed.

`<meas>`

The `<meas>` argument can be specified with one of the following arguments:

`iops`

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

`bw`

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

`svct`

Specifies that statistics for service time in milliseconds are displayed.

`size`

Specifies that statistics for I/O operations in bytes are displayed.

`<val> [k|K] | [m|M] | [g|G]`

Specifies the minimum threshold using any integer. The integer can be optionally followed with `k` or `K` to indicate a multiple of 1,000, `m` or `M` to indicate a multiple of 1,000,000, or `g` or `G` to indicate a multiple of 1,000,000,000.

An example of this option in use is `-filt curs,r,iops,10k` that indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

`-ni`

Specifies that statistics for only nonidle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

## SPECIFIERS

None.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays one iteration of statistics for VVs and LUN host attachments:

```
cli% statvln -iter 1
15:34:18 05/31/07 r/w I/O per second KBytes per sec Svt ms IOSz KB
Lun VVname Host Port Cur Avg Max Cur Avg Max Cur Avg Cur Avg Qlen
-----
```

total			t	0	0		0	0		0.0	0.0	0.0	0.0	0
-------	--	--	---	---	---	--	---	---	--	-----	-----	-----	-----	---

The columns in the example above are identified as follows:

- LUN. The LUN.
- VVname. The name of the virtual volume.
- Host. The host from which the VLUN is exported.
- Port. The port to which the VLUN is exported.
- r/w. The I/O type. Values can be read (r), write (w), or read and write (t).
- I/O per second Cur. The current number of I/O per second.
- I/O per second Avg. The average number of I/O per second.
- KBytes per Max. The maximum number of KB per second.
- KBytes per Cur. The current number of KB per second.
- KBytes per Avg. The average number of KB per second.
- Svt ms Max. The maximum service time in milliseconds.
- Svt ms Cur. The current service time in milliseconds.

- `Svt ms Avg.` The average service time in milliseconds.
- `IOSz KB Cur.` The current I/O size in KB.
- `IOSz KB Avg.` The average I/O size in KB.
- `Qlen.` The queue length.

## NOTES

- For this command: KB = 1000 bytes.
- A Domain column may be included by using the `setclienv` command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option or if the CLI was started with the `TPDLISTDOM` environment variable set. Please run `cli -h` and `setclienv -h` for details of the environment variables.

---

**COMMAND**

statvv

**DESCRIPTION**

The statvv command displays statistics for Virtual Volumes (VVs) in a timed loop.

**SYNTAX**

statvv [options <arg>] [<VV\_name\_or\_pattern>...]

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

-domain <domain\_name|pattern>...

Shows only the VVs that are in domains with names that match the specified domain names or patterns.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <seconds>

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col> [, <dir>][:<col>[, <dir>]...]

Sorts command output based on the column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

`inc`

Sort in increasing order (default).

`dec`

Sort in decreasing order.

Multiple columns can be specified separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by the values in the later columns.

`-filt <fspec>`

Specifies that statistics that fall below the threshold as specified by the `<fspec>` arguments where `<fspec>` is `<type>`, `<op>`, `<meas>`, `<val>`, are filtered out and not displayed. The `<type>`, `<op>`, `<meas>`, and `<val>` arguments are separated with one comma.

`<type>`

The `type` argument can be specified with one of the following arguments:

`curs`

Specifies that only devices with current statistics above the threshold are displayed.

`avgs`

Specifies that only devices with average statistics above the threshold are displayed.

`maxs`

Specifies that only devices with maximum values above the threshold are displayed.

`<op>`

The `<operation>` argument can be specified with one of the following arguments:

`r`

Specifies that read-only statistics are displayed.

`w`

Specifies that write-only statistics are displayed.

`t|rw`

Specifies that statistics for read and write totals are displayed.

<meas>

The meas argument can be specified with one of the following arguments:

iops

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the value argument.

bw

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the value argument.

svct

Specifies that statistics for service time in milliseconds are displayed.

size

Specifies that statistics for I/O operations in bytes are displayed.

<val> [k|K][m|M][g|G]

Specifies the minimum threshold using any integer. The integer can be optionally followed with k or K to indicate a multiple of 1,000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option in use is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only nonidle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

## SPECIFIERS

<VV\_name|pattern>...

Only statistics are displayed for the specified VV or pattern. Multiple volumes or patterns can be repeated (for example <VV\_name> <VV\_name> . . .). If not specified, all VVs are listed.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays I/O statistics for all VVs:

```
cli% statvv -iter 1
15:31:21 05/31/07 r/w I/O per second KBytes per sec  Svt ms IOSz KB
      VVname      Cur  Avg  Max  Cur  Avg  Max Cur  Avg Cur  Avg Qlen
      admin    t    0    0    0    0    0    0 0.0 0.0 0.0 0.0  0
      test2    t    0    0    0    0    0    0 0.0 0.0 0.0 0.0  0
-----
      total    t    0    0          0    0          0.0 0.0 0.0 0.0  0
```

The columns in the example above are identified as follows:

- VVname. The virtual volume name.
- r/w. The I/O type. Values can be read (r), write (w), or read and write (t).
- I/O per second Cur. The current number of I/O per second.
- I/O per second Avg. The average number of I/O per second.
- KBytes per Max. The maximum number of KB per second.
- KBytes per Cur. The current number of KB per second.
- KBytes per Avg. The average number of KB per second.
- Svt ms Max. The maximum service time in milliseconds.
- Svt ms Cur. The current service time in milliseconds.
- Svt ms Avg. The average service time in milliseconds.
- IOSz KB Cur. The current I/O size in KB.
- IOSz KB Avg. The average I/O size in KB.
- Qlen. The queue length.

## NOTES

- If a <VV\_name> or <pattern> are specified, then VVs with names matching any of the patterns are listed. Otherwise all VVs are listed. These patterns are glob-style patterns (see Help on sub,globpat).
- VVs may be accessed externally by hosts and internally by the prefetcher. VV data measured by this command include accesses by the prefetcher.

- In addition to external accesses by hosts, VVs may be read internally by the system read-ahead prefetcher. The `histvv` data includes read-ahead accesses from the prefetcher, which can cause the read data to appear more than seen by the hosts. Use the `histvln -vvsum` command to see data for only accesses from the host.
- For this command KB = 1000 bytes.



# 26

## Stop Commands

---

In this chapter

stopcim	<b>26.2</b>
stoprcopy	<b>26.4</b>
stoprcopygroup	<b>26.6</b>

---

**COMMAND**

stopcim

**DESCRIPTION**

The stopcim command stops the CIM server from servicing CIM requests.

**SYNTAX**

stopcim [option]

**AUTHORITY**

Super, Service



**NOTE:** You need access to all domains in order to run this command.

**OPTIONS**

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-x

Specifies that the operation terminates the server immediately without graceful shutdown notice.

**SPECIFIERS**

None.

**RESTRICTIONS**

None.

**EXAMPLES**

The following example stops the CIM server:

```
cli% stopcim
Are you sure you want to stop CIM server?
select q=quit y=yes n=no: y
CIM server stopped successfully.
```

The following example stops the CIM server immediately without graceful shutdown notice and confirmation:

```
cli% stopcim -f -x  
CIM server stopped successfully.
```

## NOTES

By default, the CIM server is not started until the `startcim` command is issued.

---

## COMMAND

`stoprcopy`

## DESCRIPTION

The `stoprcopy` command stops the Remote Copy service and optionally stops any started Remote Copy volume groups.

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## SYNTAX

`stoprcopy [options]`

## OPTIONS

`-f`

Specifies that any started copy will not ask for confirmation for the `-clear` option.

`-stopgroups`

Specifies that any started Remote Copy volume groups are stopped.

`-clear`

Specifies that configuration entries affiliated with the stopped mode are deleted.



**CAUTION:** Issuing the `stoprcopy -clear` command completely removes the Remote Copy setup and is NOT reversible.

## SPECIFIERS

None.

## RESTRICTIONS

If the `-stopgroups` option is not used, all groups must already be stopped.

## EXAMPLES

The following example disables the Remote Copy functionality of all primary Remote Copy volume groups:

```
cli% stoprcopy -stopgroups
```

## NOTES

- Unless the `-stopgroups` option is used, the command will fail if there are any Remote Copy groups that are started. If the `-clear` option has been used, the configuration is completely erased. Consequently, Remote Copy operations can not be restarted using only the `startrcopy` command. The configuration must be rebuilt. Therefore, the `-clear` option requires confirmation with the `-f` option, the `TPDFORCE` environment variable, or by interactively typing, `y`.

---

**COMMAND**

`stoprcopygroup`

**DESCRIPTION**

The `stoprcopygroup` command stops the Remote Copy functionality for the specified Remote Copy volume group.

**SYNTAX**

`stoprcopygroup [options] <group_name>`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-nosnap`

In synchronous mode, this option turns off the creation of snapshots. This is useful if `removercopygroup` is to be run to remove Remote Copy. In asynchronous periodic mode, this option deletes any current synchronization snapshots. Using this option will result in a full resync being required if the group is later restarted.

`-t <target_name>`

Stops the group on the specified target.

**SPECIFIERS**

`<group_name>`

The name of the Remote Copy volume group to stop.

**RESTRICTIONS**

This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.

**EXAMPLES**

The following example stops Remote Copy for Group1:

```
cli% stoprcopygroup Group1
```

**NOTES**

In `sync` mode, this command creates snapshots that are used for synchronizing the primary and secondary groups if `starttrcopygroup` is run later.





# 27

## Sync Command

---

In this chapter

`syncrcopy`

**27.2**

---

**COMMAND**

`syncrcopy`

**DESCRIPTION**

The `syncrcopy` command manually synchronizes Remote Copy volume groups.

**SYNTAX**

`syncrcopy [options] <group_name>`

**AUTHORITY**

Super, Edit

**OPTIONS**

`-w`

Wait for synchronization to complete before returning to a command prompt.

`-n`

Do not save resynchronization snapshot. This option is only relevant for asynchronous periodic mode volume groups.



**NOTE:** Using the `-n` option requires a full synchronization at the next sync.

`-ovrd`

Force a full synchronization of the group even if the volumes are already synchronized. This option is only relevant for synchronous mode volume groups and can be used to re-synchronize volumes that have become inconsistent.

`-t <target_name>`

Synchronize the group only to the specified target.

**SPECIFIERS**

`<group_name>`

Specifies the name of the Remote Copy volume group to be synchronized. This name can be obtained using the `showrcopy` command, as described on [page 22.145](#).

## RESTRICTIONS

- This command requires the 3PAR Remote Copy license. Contact your local service provider for more information.
- Using the `-n` option requires a full synchronization at the next sync.

## EXAMPLE

The following example specifies that Remote Copy volume group `Group1` should be synchronized with it's corresponding secondary volume group:

```
cli% syncrcopy Group1  
Synchronization request issued for group Group1
```

## NOTES

- Mode (synchronous or asynchronous periodic) is set using the `creatercopy group` command. For information about modes and creating Remote Copy volume groups, see [creatercopygroup](#) on page 11.48.
- For information about setting targets, see [creatercopytarget](#) on page 11.50.



# 28

## Tune Commands

---

In this chapter

tunealdvv	<b>28.2</b>
tunepd	<b>28.10</b>
tunetpvv	<b>28.15</b>
tunevv	<b>28.18</b>

---

## COMMAND

tunealdvv

## DESCRIPTION

This command was deprecated in the 2.2.4 release and will be changed or removed in a future release. Please use the `tunevv` command in the future. If virtual volumes were created with deprecated commands then only deprecated commands can be used to modify virtual volumes.

The `tunealdvv` command changes the layout of a virtual volume.

## SYNTAX

The `tunealdvv` command uses one of the following syntax conventions:

- `tunealdvv [options <arg>] <VV_name>`
- `tunealdvv restart [options <arg>] <VV_name>`
- `tunealdvv rollback [options <arg>] <VV_name>`

## AUTHORITY

Super, Edit<sup>1</sup>



**NOTE:** You need access to all domains in order to run this command.

## SUBCOMMANDS

restart

Restarts a `tunealdvv` command operation that was previously interrupted because of component failure or user-initiated cancellation.

rollback

Rolls back a `tunealdvv` command operation that was previously interrupted. The `canceltask` command needs to run before the rollback.

---

<sup>1</sup> Certain options require this additional privilege restriction as indicated.

## OPTIONS

The following options can be used on all commands:

`-f`

Forces the command. The command completes without prompting for confirmation.

`-waittask`

Specifies that the command will wait for any created tasks to complete.

The following options cannot be used with the `restart` or `rollback` subcommands:

`-cpg <CPG_name>`

Name of the Common Provisioning Group (CPG) to which the snapshot data space is moved. If the `-cpg` option is specified, only the snapshot data space of a volume is tuned. When snapshot data space is tuned, the only permitted option is the `-cnt` option. If this option is not specified, only the USR space of a volume is tuned.

`-templ <template_name>`

Use the options defined in template `<template_name>`. The template is created using the `createtemplate` command. Options specified in the template are read-only or read-write. The read-write options may be overridden with new options at the time of their creation, but read-only options may not be overridden at creation time. Options not explicitly specified in the template take their default values, and all of these options are either read-only or read-write (using the `-nro` or `-nrw` options of the `createtemplate` command). Sizing options from the template are ignored. This command does not change the size of the VV on which it acts.

`-t <RAID_type>`

Specifies the RAID type of the Logical Disk (LD): `r0` for RAID-0, `r1` for RAID-1, or `r5` for RAID-5, or `r6` for RAID-6. If no RAID type is specified, the default is `r1`.

`-ssz <size_number_chunklets>`

Specifies the set size in the number of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

`-rs <size>`

Specifies the number of sets in a row. The `<size>` is a positive integer. If not specified, no row limit is imposed.

`-ss <size_KB>`

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. If no value is entered, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. For RAID-6, the default is a function of the set size.

`-ha port|cage|mag`

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0.

`-ch first|last`

Specifies the chunklet location preference, either `first` (attempt to use the lowest numbered available chunklets) or `last` (attempt to use the highest numbered available chunklets). If no argument is specified, the default location is `first`.

`-p <pattern>`

Specifies a pattern for candidate disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to all Fibre Channel (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The `-devtype` pattern cannot be used to mix Near Line (NL), FC, and Solid State Drive (SSD) drives:



**NOTE:** An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

`-nd <item>`

Specifies one or more nodes. Nodes are identified by one or more integers (`<item>`). Multiple nodes are separated with a single comma (1, 2, 3). A range of nodes is separated with a hyphen (0-7). The primary path of the disks must be on the specified node(s).

`-st <item>`

Specifies one or more PCI slots. Slots are identified by one or more integers (`<item>`). Multiple slots are separated with a single comma (1, 2, 3). A range of slots is separated with a hyphen (0-7). The primary path of the disks must be on the specified PCI slot(s).



`-pt <item>`

Specifies one or more ports. Ports are identified by one or more integers (<item>). Multiple ports are separated with a single comma (1, 2, 3). A range of ports is separated with a hyphen (0–5). The primary path of the disks must be on the specified port number(s).

`-cg <item>`

Specifies one or more drive cages. Drive cages are identified by one or more integers (<item>). Multiple drive cages are separated with a single comma (1, 2, 3). A range of drive cages is separated with a hyphen (0–3). The specified drive cage(s) must contain disks.

`-mg <item>`

Specifies one or more drive magazines. The 1. or 0. displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (<item>). Multiple drive magazines are separated with a single comma (1, 2, 3). A range of drive magazines is separated with a hyphen (0–7). The specified drive magazine(s) must contain disks.

`-pn <item>`

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (<item>). Multiple disk positions are separated with a single comma (1, 2, 3). A range of disk positions is separated with a hyphen (0–3). The specified position(s) must contain disks.

`-dk <item>`

Specifies one or more physical disks. Disks are identified by one or more integers (<item>). Multiple disks are separated with a single comma (1, 2, 3). A range of disks is separated with a hyphen (0–3). Disks must match the specified ID(s).



**NOTE:** The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to limit the disks that are used to create logical disks based on the characteristics of the disk drive.

`-tc_gt <number>`

Specifies that physical disks with total chunklets greater than the number specified are selected.

`-tc_lt <number>`

Specifies that physical disks with total chunklets less than the number specified are selected.

`-fc_gt <number>`

Specifies that physical disks with free chunklets greater than the number specified are selected.

`-fc_lt <number>`

Specifies that physical disks with free chunklets less than the number specified are selected.

`-devid <model>`

Specifies that physical disks identified by their models are selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

`-devtype <type>`

Specifies that physical disks must have the specified device type (FC for Fibre Channel, NL for Nearline, or SSD for Solid State Drive) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.

`-rpm <number>`

Disks must be of the specified speed. Device speeds are shown in the `Speed` column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). It is meant as a rough estimation of the performance difference between the drive and the other drives in the system. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD drive, the number is to be treated as relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.

Disks that satisfy all of the specified characteristics are used. For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

`-wait <secs>`

If the `tunealdivv` command fails due to the lack of clean space, the `-wait` option specifies the number of seconds to wait for the system to clean the dirty space before returning.

If `-wait 0` is issued, the command returns immediately. If this option is not used, the command will keep waiting for dirty chunklets to be cleaned if enough space will be available with the dirty chunklets cleaned.

`-dr`

Specifies that the command is a dry run and that no logical disks or virtual volumes are created.

`-verbose on|off`

Specifies that verbose output is either enabled (`on`) or disabled (`off`). If not specified, verbose is disabled.

`-cnt <number_of_VV>`

Specifies the number of identical virtual volumes to tune using an integer from 1 through 999. If not specified, one virtual volume is tuned. If the `-cnt` option is specified, then `-restart` and `-rollback` options and the `restart` and `rollback` subcommands are not permitted.

`-restart`

Restart a `tunealdivv` command call that was previously interrupted because of component failure, or because of user-initiated cancellation. This option was deprecated in the 2.2.3 release and will be changed or removed in a future release.

`-rollback`

Rollback a `tunealdivv` command call that was previously interrupted. The `canceltask` command needs to run before the rollback. This option was deprecated in the 2.2.3 release and will be changed or removed in a future release.

## SPECIFIERS

`<VV_name>`

Specifies an existing virtual volume name.

## RESTRICTIONS

- Access to all domains is required to run this command.
- This command requires a Dynamic Optimization license. Contact your local 3PAR representative for information.

- This command requires sufficient available physical space, equivalent to the user size of the volume being tuned.
- When changing the layout of a virtual volume, you can optionally apply arguments originally defined for logical disk creation through the `createtemplate` command by issuing the `tunealddv -templ <template_name>` command. The `-templ` option is only valid for logical disk template object types.

## EXAMPLES

The following example shows how to use the `tunealddv` command to convert a virtual volume (`testvol`) to RAID-5:

```
cli% tunealddv -t r5 testvol
Task 1 started.
```

The following example shows how to use the `tunealddv` command to change the availability level of a virtual volume (`testvol`) to `cage`. Note that because the default availability parameter setting (`-ha`) for `tunealddv` is `cage`, it is not necessary to explicitly specify `cage`-level availability when issuing this command.

```
cli% tunealddv testvol
Task 2 started.
```

The following example shows how to use the `tunealddv` command to add a disk filter specifying that the logical disks supporting virtual volume `testvol` must use chunklets on physical disks 20 and 31 only.

```
cli% tunealddv -p -dk 20,31 testvol
Task 3 started.
```

The following examples shows how to start, cancel, and then restart a `tunealddv` operation on VV `testvol`:

```
cli% tunealddv -t r5 -ha mag testvol
Task 1 started.
```

```
cli% canceltask 1
Are you sure you want to cancel task 1?
select q=quit y=yes n=no: y
```

```
cli% tunealdevv -restart testvol
Task 2 started.
```

## NOTES

- This command was deprecated in the 2.2.4 release and will be changed or removed in a future release. Please use the `tunevv` command in the future.
- Automatic chunklet selection for relocations will prefer selection of chunklets on physical disks with the same device type as the source. Use the `-p devtype` option to override this default. Use the `showpd` command to see the device types of physical disks in the system.
- When canceling a `tunealdevv` task, the `canceltask` command can return before a cancellation is completed. Thus, resources reserved for the task cannot be immediately available. This can prevent actions like restarting the canceled `tunealdevv` task. Use the `waittask` command in your scripts to ensure orderly completion of the `tunealdevv` cancellation before taking other actions. See [waittask](#) on page 31.2 for details about using the `waittask` command.
- This command is only used for the non-provisioned VV. The `tunevv` command should be used instead to create a VV that has its user space provisioned from a Common Provision Group.

---

## COMMAND

tunepd

## DESCRIPTION

The `tunepd` command identifies physical disks with high service times and optionally executes load balancing.

## SYNTAX

The syntax for the `tunepd` command can be one of the following:

- `tunepd [options <arg>] maxsvct <msecs>|highest`
- `tunepd [options <arg>] avgsvct <msecs>|highest`

## AUTHORITY

Super, Edit



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS



**NOTE:** If the `-nodes`, `-slots`, `-ports`, or `-vv` options are not specified, all physical disks are included.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (1, 2, 3). The node list can also consist of a single integer (1). If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (1, 2, 3). A list can also consist of a single integer (1). If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified port slots and physical disks connected to those port slots. The port list is specified as a series of integers separated by commas (1, 2, 3). A list can also consist of a single integer (1). If the port list is not specified, all disks on all ports are displayed.

`-vv <VW_name>`

Specifies that the physical disks used by the indicated virtual volume name are included for statistic sampling.

`-d <seconds>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to 30 seconds.

`-iter <number>`

Specifies that I/O statistics are sampled a specified number of times as indicated by the `number` argument using an integer greater than 0. If 0 is specified, I/O statistics are looped indefinitely. If this option is not specified, the command defaults to 1 iteration.

`-freq <minutes>`

Specifies the interval, in minutes, that the command enters standby mode between iterations using an integer greater than 0. If this option is not specified, the number of iterations is looped indefinitely.

`-vvlayout`

Specifies that the layout of the virtual volume is displayed. If this option is not specified, the layout of the virtual volume is not displayed.

`-portstat`

Specifies that statistics for all disk ports in the system are displayed. If this option is not specified, statistics for ports are not displayed.

`-pdstat`

Specifies that statistics for all physical disks, rather than only those with high service times, are displayed. If this option is not specified, statistics for all disks are not displayed.

`-chstat`

Specifies that chunklet statistics are displayed. If not specified, chunklet statistics are not displayed. If this option is used with the `-movech` option, either `on` or `force` must be specified. See [Restrictions](#) on page 28.12 for conditions on using the `-chstat` option.

`-maxpd <number>`

Specifies that only the indicated number of physical disks with high service times are displayed. If this option is not specified, 10 PDs are displayed.

`-movech auto|manual`

Specifies that if any disks with unbalanced loads are detected that chunklets are moved from those disks for load balancing.

`auto`

Specifies that the system chooses source and destination chunklets. If not specified, you are prompted for selecting the source and destination chunklets.

`manual`

Specifies that the source and destination chunklets are manually entered.

## SPECIFIERS

`maxsvct <msec>|highest`

Specifies that either the maximum service time threshold (`<msec>`) that is used to discover overutilized physical disks, or the physical disks that have the highest maximum service times (`highest`). If a threshold is specified, then any disk whose maximum service time exceeds the specified threshold is considered a candidate for load balancing.

`avgsvct <msec>|highest`

Specifies that either the average service time threshold (`<msec>`) that is used to discover overutilized physical disks, or the physical disks that have the highest average service time (`highest`). If a threshold is specified, any disk whose average service time exceeds the specified threshold is considered a candidate for load balancing.

## RESTRICTIONS

- Access to all domains is required to run this command.
- This command requires a System Tuner license. Contact your local 3PAR representative for information.
- The system does not allow multiple `tunepd` processes with the `-chstat on` or `-chstat force` options to run concurrently.
- If one or more chunklet statistics collection processes are running, the `tunepd` command exits to prevent possible `tunepd -chstat on` processes running simultaneously.



- Use the `-chstat force` option only when you have verified that no other `tunepd` processes are running.

## EXAMPLES

In the following example, PDs with average service times exceeding 50 milliseconds are identified and their chunklets automatically relocated to rebalance the PDs' load.

```
cli% tunepd -vvlayout -chstat -movech auto avgsvct 50
Collecting I/O statistic for physical disks (PDs) ...
```

PdId	Pos	APort	BPort	Iops	Kbps	Svct(ms)	IOSz(KB)
29	1:3:3	0:4:1	1:5:1*	231.9	8668.8	87.0	37.4

The following physical disks were identified as candidates for load balancing: 29

Layout of related virtual volumes (on candidate physical disks).

VvId	VvName	LdId	LdName	Avail	RAID	PdId
0	admin	5	admin.usr.2	cage	1	0,2,4,6,8,10,12,14,17,19 23,25,27,29
1	vv0	8	vv0.usr.2	cage	1	0,2,4,6,8,10,12,14,17,19 23,25,27,29
2	vv1	11	vv1.usr.1	mag	5	2,6,19,29

Enable statistic collection for chunklets on PD 29  
Collecting statistic for chunklets of candidate PDs...

Disable statistic collection for chunklets on PD 29

Statistic of chunklets of candidate PDs:

PdId	PdCh	LdId	LdName	LdCh	Iops	Kbps	Svct(ms)	IOSz(KB)
29	20	8	vv0.usr.2	29	22.8	842.7	109.9	37.0
29	19	8	vv0.usr.2	17	22.0	842.7	113.7	38.2
29	18	8	vv0.usr.2	5	7.8	282.3	117.0	36.3
29	26	11	vv1.usr.1	23	42.0	1561.6	77.0	37.2
29	25	11	vv1.usr.1	19	37.8	1401.4	70.2	37.1
29	24	11	vv1.usr.1	15	39.5	1454.1	73.3	36.8
29	23	11	vv1.usr.1	11	35.3	1333.6	73.1	37.7
29	22	11	vv1.usr.1	7	26.5	1018.1	76.8	38.4

Statistic of Logical disks of candidate PDs:

PdId	LdId	LdName	Iops	Kbps	Svct(ms)	IOSz(KB)
29	8	vv0.usr.2	52.6	1967.8	112.5	37.4
29	11	vv1.usr.1	181.1	6768.8	74.0	37.4

Statistic of candidate physical disks:

PdId	Pos	APort	BPort	Iops	Kbps	Svct(ms)	IOSz(KB)
29	1:3:3	0:4:1	1:5:1*	233.7	8736.6	82.7	37.4

```
Physical Disk 29: The following chunklets have been marked for moving: 26
```

```
Pass 1: Dryrun moving chunklets:
```

Move	Status	qset	grow	qcon	qlocal	qlost
29:26-10:16	source and destination are valid	mag	0	2	true	false

```
Pass 2:
```

Move	Status	qset	grow	qcon	qlocal	qlost
29:26-10:16	ch was moved	mag	0	2	true	false

Note that physical disk 29 is identified as matching the search criterion of average service times greater than 50 milliseconds and its chunklets are relocated.

## NOTES

- When the `-movech auto` or `-movech manual` option is specified the system only identifies (auto mode) or recommends (manual mode) source chunklets for which there are destination chunklets retaining the availability of the source chunklets' logical disks.
- Separate instances of the `tunepd` command can be issued to identify load balancing candidates and relocate chunklets, or execute both tasks with one instance of the `tunepd` command.
- Previous options `-n <node_list>`, `-s <slot_list>`, and `-p <port_list>` have been deprecated in 2.2.3 and replaced with `-nodes <node_list>`, `-slots <slot_list>`, and `-ports <port_list>`.
- If the `-nodes`, `-slots`, `-ports` or `-vv` options are not specified, all physical disks are included.

---

## COMMAND

tunetpvv

## DESCRIPTION

This command was deprecated in the 2.2.3 release and will be changed or removed in a future release. Please use the `tunevv` command in the future. If Virtual Volumes (VVs) were created with deprecated commands then only deprecated commands can be used to modify VVs.

The `tunetpvv` command allows the RAID and Availability characteristics of an exiting Thin Provisioned Virtual Volume (TPVV) to be dynamically modified.

## SYNTAX

- `tunetpvv usr_cpg <CPG> [options] <VV_name>`
- `tunetpvv snp_cpg <CPG> [options] <VV_name>`
- `tunetpvv [options] <VV_name>`

## AUTHORITY

Super, Edit

## SUBCOMMANDS

The following subcommands are only used for tuning a multi-space TPVV:

`usr_cpg <CPG>`

Specifies the name of the Common Provisioning Group (CPG) to which the user space of the volume is moved.

`snp_cpg <CPG>`

Specifies the name of the CPG to which the snapshot space of the volume is moved.

## OPTIONS

`-f`

Forces the command. The command completes without prompting for confirmation.

`-waittask`

Indicates to wait for created tasks to complete.

`-dr`

Specifies that the command is a dry run and that no logical disks or Virtual Volumes (VVs) are actually tuned.

`-cnt <nb_of_VV>`

Specifies the number of identical VVs to tune using an integer from 1 through 999. The default will tune one VV.

The following option is only used for tuning a legacy single-space TPVV:

`-cpg <CPG_name>`

Specifies the CPG to which the snapshot data space is moved. This option is not valid with either subcommands.

This option is deprecated and will be removed in a future release.

## SPECIFIERS

`<VV_name>`

Specifies an existing VV name.

## RESTRICTIONS

- This command requires a Dynamic Optimization license. Contact your local 3PAR representative for information.
- This command requires sufficient available physical space with the same owner or backup node as the space it is being copied from.

## EXAMPLES

The following example shows how to use the `tunetpvpv` command to change the RAID type of a TPVV (`testvol`) from `cpg1`, which is originally `cpg1`, which is of RAID 1.

```
cli% createcpg -t r5 cpg2

cli% tunetpvpv -cpg cpg2 testvol

Are you sure you want to tune VV 'testvol' ?

select y=yes n=no: y
Task 1 started.
```

The following example shows how to use the `tunetpvv` command to increase the availability level of a TPVV (`testvol`). Note that there is mag availability because `testvol` was originally on `cpg1`, which is created with option (`-ha mag`). To increase cage availability, create a CPG with cage availability and tune `testvol` to the associated CPG.

```
cli% createcpg cpg3
cli% tunetpvv -cpg cpg3 testvol
Are you sure you want to tune VV 'testvol'?
select y=yes n=no: y
Task2 started.
```

The following example shows how to tune multiple VVs together. `Testvol.0`, `testvol.1` and `testvol.2` are the three VVs to be tuned:

```
cli% tunetpvv -cnt 3 testvol
Are you sure you want to tune VV 'testvol'?
select y=yes n=no: y
Task 3 started.
```

## NOTES

- When canceling a `tunetpvv` command task, the `canceltask` command may return before a cancellation is complete. Thus, resources reserved for the task may not be immediately available. This will prevent actions like attempting another tune on the same VV. Scripts should use the `waittask` command in order to ensure orderly completion of the `tunetpvv` command cancellation prior to taking other actions.

---

## COMMAND

tunevv

## DESCRIPTION

The `tunevv` command is used to change the properties of a virtual volume that was created with either the `createvv` or `createtpvv` command by associating it with a different Common Provisioning Group.

## SYNTAX

The `tunevv` command uses one of the following syntax conventions:

- `tunevv usr_cpg <CPG> [options <arg>] <VV_name>`
- `tunevv snp_cpg <CPG> [options <arg>] <VV_name>`
- `tunevv restart [options <arg>] <VV_name>`
- `tunevv rollback [options <arg>] <VV_name>`

## AUTHORITY

Super, Edit<sup>1</sup>

## SUBCOMMANDS

`usr_cpg <CPG>`

Moves the logical disks being used for user space to the specified Common Provisioning Group.

`snp_cpg <CPG>`

Moves the LDs being used for snapshot space to the specified Common Provisioning Group. This option cannot be used on Thinly Provisioned Virtual Volumes.

`restart`

Restarts a `tunevv` command call that was previously interrupted because of component failure, or because of user initiated cancellation. This cannot be used on Thinly Provisioned Virtual Volumes.

---

<sup>1</sup> Certain options require this additional privilege restriction as indicated.

## rollback

Returns to a previously issued `tunevv` operation call that was interrupted. The `canceltask` command needs to run before the rollback. This cannot be used on Thinly Provisioned Virtual Volumes.

## OPTIONS

`-f`

Forces the command. The command completes without prompting for confirmation.

`-waittask`

Specifies that the command will wait for any created tasks to complete.

`-dr`

Specifies that the command is a dry run and that no logical disks or virtual volumes are actually tuned.

The following option is used when the `<usr_cpg>` and `<snp_cpg>` subcommands are specified:

`-cnt <count>`

Specifies the number of identical virtual volumes to tune using an integer from 1 through 999. For Thinly Provisioned Virtual Volumes, this must be an integer between 1 and 64. If not specified, one virtual volume is tuned. If the `-cnt` option is specified, then the subcommands, `restart` and `rollback` are not permitted.

## SPECIFIERS

`<VV_name>`

Specifies the name of the existing virtual volume.

## RESTRICTIONS

This command requires a Dynamic Optimization license. Contact your local 3PAR representative for information.

## EXAMPLES

In the following example, the LDs used for user space are moved to CPG `cpg_sn1.0_p` for VV `nf_st_tp_22.0`:

```
cli% tunevv usr_cpg cpg_sn1.0_p nf_st_tp_22.0
Task 999 started
```

## NOTES

When canceling a `tunevv` command task, the `canceltask` command can return before a cancellation is completed. Therefore, resources reserved for the task might not be immediately available. This can prevent actions such as restarting the canceled `tunevv` command task. Scripts should use the `waittask` command in order to ensure orderly completion of the `tunevv` command cancellation before taking other actions. See [waittask](#) on page 31.2 for details about using the `waittask` command.



# 29

## Update Commands

---

In this chapter

updatesnapspace	<b>29.2</b>
updatevv	<b>29.4</b>

---

## COMMAND

updatesnapspace

## DESCRIPTION

The `updatesnapspace` command starts a non-cancelable task to update the snapshot space usage accounting. The snapshot space usage displayed by `showvv -snapspace` is not necessarily the current usage and the `TimeCalculated` column will show when it was last calculated. This command causes the system to start calculating current snapshot space usage. If one or more VV names or patterns are specified, only the specified VVs will be updated. If none are specified, all VVs will be updated.

This command will return immediately, displaying the task ID associated with the update. To wait for completion, use the `waittask` command.

## SYNTAX

`updatesnapspace [<VV_name>...|<pattern>...]`

## AUTHORITY

Super, Edit

## OPTIONS

None.

## SPECIFIERS

`<VV_name>...`

Specifies the VV name to update. This specifier can be repeated to display the task ID about multiple VVs. This specifier is not required on the command line. If not specified, all VVs in the system are updated.

`<pattern>...`

Specifies that the VVs matching the specified glob-style pattern is updated. This specifier can be repeated. This specifier is not required. See [Glob-Style Pattern](#) on page 2.4 for more information.

## RESTRICTIONS

None.

## EXAMPLES

The following example displays the actual snapshot space used by all VVs:

```
cli% updatesnapspace
Task 2 has been started to calculate actual space usage.
```

Issuing a `waittask` command displays the information about the task in process or if it has been processed:

```
cli% waittask 2
Task 2 done
```

The following example displays the task information using `-d <task_id>`:

```
cli% showtask -d 2
Id          Type          Name Status Phase Step -----StartTime-----
-----FinishTime-----
  2 snapspace_accounting ss_accounting   Done   0/0   0/0 Thu Sep 22 18:43:35 PDT
2005 Thu Sep 22 18:43:35 PDT 2005

Detailed status:
{Thu Sep 22 18:43:35 PDT 2005} Created      task.
{Thu Sep 22 18:43:35 PDT 2005} Started      snapshot usage data collection proces
s for VVs
{Thu Sep 22 18:43:35 PDT 2005} Updated      snapshot usage data for VV emaildb
{Thu Sep 22 18:43:35 PDT 2005} Updated      snapshot usage data for VV www
{Thu Sep 22 18:43:35 PDT 2005} Updated      snapshot usage data for VV snapname
{Thu Sep 22 18:43:35 PDT 2005} Finished     snapshot usage data collection process
```

## NOTES

- If one or more VV names or patterns are specified, only the specified VVs are updated. If VV names are not specified, all VVs are updated.
- To check snapshot space usage, use the `showvv -s` command.
- See [showtask](#) on page 22.180 for more information.
- See [waittask](#) on page 31.2 for more information.

---

## COMMAND

updatevv

## DESCRIPTION

The `updatevv` command updates a snapshot Virtual Volume (VV) with a new snapshot.

## SYNTAX

`updatevv [options] <VV_name | VV_set>...`

## AUTHORITY

Super, Edit, Browse

## OPTIONS

`-ro`

Specifies that if the specified VV (<VV\_name>) is a read/write snapshot the snapshot's read-only parent volume is also updated with a new snapshot if the parent volume is not a member of a virtual volume set. If this option is not specified, the VV, as indicated with the <VV\_name> specifier, is replaced by a new snapshot. See *Notes* for additional information.

`-anyid`

Specifies that any VV ID can be used for the new snapshot(s).

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

## SPECIFIERS

`<VV_name | VV_set>...`

Specifies the name(s) of the snapshot virtual volume(s) or virtual volume set(s) to be updated. They must be of same type (read-only or read/write). When a <VV\_set> or multiple <VV\_name>'s are specified group-consistent snapshots are taken. The virtual volume set name must start with `set :`. See `creategroupsv` for more information.

## RESTRICTIONS

For each snapshot VV name specified, the `updatevv` command first verifies in your Access Control List (ACL) that you have permission to run the `updatevv` command on the specified VVs. For information on viewing and setting your ACL, see [showuseracl](#) on page 22.190 and [setuseracl](#) on page 21.86.

## EXAMPLES

The following example displays the snapshot update of snapshot VV `avvro`:

```
cli% updatevv -f avvro  
Updating VV avvro
```

## NOTES

- The `updatevv` command attempts to use the same VV IDs for new snapshots as the snapshots that they replace so that the new VV have the same World Wide Names (WWNs) as the originals (this might be important for hosts using the VV WWN to identify the VV). However, between the removal of the original snapshot and the creation of the new snapshot it is possible that another VV might have been created with the same ID causing the creation of the new snapshot to fail. If the `-anyid` option is specified, the command then creates the snapshot with any available ID.

Because new VVs, by default, are assigned the lowest ID available, VVs that are updated with the `updatevv` command should be assigned large IDs (using the `-i` option to the `createalddv` command) to reduce the likelihood that their IDs are taken during the `updatevv` command operation.

- After the `updatevv` command is executed, all VLUNs associated with the specified VV name are removed. The command then updates the snapshots as follows:
  - ◆ If the `-ro` option is not specified, or if the specified VV name is a read-only snapshot, the command removes the snapshot, and creates a new snapshot of the same name and of the same parent.
  - ◆ If the `-ro` option is specified and the specified VV name is a read/write snapshot, the VV is removed and its read-only parent is replaced by a new read-only snapshot of the same name. Then a new read/write snapshot is created of the new read-only snapshot. If the read-only parent has multiple read/write snapshots, the `updatevv` operation will fail because the read-only parent volume cannot be removed.

The command then re-creates all the VLUNs associated with the specified VV name.

- Running concurrent `updatevv` sessions for VV sets or list of VVs which have common VVs could have unpredictable results.



# 30

## Upgrade Commands

---

In this chapter

upgradecage	<b>30.2</b>
upgradepd	<b>30.4</b>

---

## COMMAND

upgradecage

## DESCRIPTION

The upgradecage command downloads new firmware into the specified cage.

## SYNTAX

The syntax for the upgradecage command can be one of the following:

- upgradecage [options <arg>] <cagename>...
- upgradecage [options <arg>] -a

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

-a

Specifies that all drive cages are upgraded with new firmware (<cagename> should not be specified).

-noA

Do not upgrade the A loop. This option is only valid for DC1 cages.

-noB

Do not upgrade the B loop. This option is only valid for DC1 cages.

-skiptest

Skips the 10 second per PD diagnostic test normally completed after each cage upgrade.

-model <model>

Only upgrade cages of specified model. Use showcage command to find out the correct model for the cages. The model can be DC1, DC2, DC3, or DC4.



`-file <fpname>`

Use firmware in file <fpname> , where <fpname> is the full path name of the file on the node. The default is to use the firmware that is marked current in the firmware database (see the `showfirmwaredb` command).

`-minlevel <revlevel>`

Minimum firmware level that the cage firmware must be at for the upgrade to be allowed. This option is deprecated and will be removed in a future release.

`-maxlevel <revlevel>`

Maximum firmware level that the cage firmware must be at for the upgrade to be allowed. This option is deprecated and will be removed in a future release.

## SPECIFIERS

`<cagename>`

Specifies the name displayed in the Name column using the `showcage` command.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays a cage with firmware level of 4.41 upgrading to 4.42:

```
cli% upgradecage -a

Upgrading cage cage0 cpuA from rev 1.46 to revision in file /opt/tpd/fw/cage/dc2/lbod_fw.bin-2.05.
Upgrading cage cage1 cpuA from rev 1.46 to revision in file /opt/tpd/fw/cage/dc2/lbod_fw.bin-2.05.
Upgrading cage cage2 cpuA from rev 1.46 to revision in file /opt/tpd/fw/cage/dc2/lbod_fw.bin-2.05.
Upgrading cage cage3 cpuA from rev 03 to revision in file /opt/tpd/fw/cage/dc3/dc3_fw.bin-04.
Skipping cage cage4 cpuA already up to date at rev 2.05
Skipping cage cage5 cpuA & cpuB already up to date at rev 04
```

## NOTES

- Before executing the `upgradecage` command, issue the `showcage` command to obtain the names of the drive cages in the system.
- When the `upgradecage` command is issued, the drive cage becomes temporarily unavailable. The cage automatically restarts following the firmware update, and may briefly disrupt the loops during the restart. It is recommended `showcage` be used approximately one minute after the update completes to ensure both loops to the cage are available again.

---

## COMMAND

upgradepd

## DESCRIPTION

The `upgradepd` command upgrades the Physical Disk (PD) firmware.

## SYNTAX

```
upgradepd [-f] [-skiptest] {-a | -w <WWN>... | <PD_ID>...}
```

## AUTHORITY

Super, Service



**NOTE:** You need access to all domains in order to run this command.

## OPTIONS

`-f`

Upgrades the PD firmware without requiring confirmation.

`-skiptest`

Skips the 10 second diagnostic test normally completed after each PD upgrade.

`-a`

Specifies that all PDs with valid IDs and whose firmware is not current are upgraded. If this option is not specified, then either the `-w` option or `PD_ID` specifier must be issued on the command line.

`-w <WWN>...`

Specifies that the firmware of either one or more PDs, identified by their WWNs, is upgraded. If this option is not specified, then either the `-a` option or `PD_ID` specifier must be issued on the command line.

## SPECIFIERS

`<PD_ID>...`

Specifies that the firmware of either one or more PDs identified by their IDs (`PD_ID`) is upgraded. If this specifier is not used, then the `-a` option or `-w` option must be issued on the command line.

## RESTRICTIONS

Access to all domains is required to run this command.

## EXAMPLES

The following example displays the upgrade of all PDs:

```
cli% upgradepd -a
```

## NOTES

- PDs can be upgraded while I/O is occurring.
- If a disk with RAID-0 chunklets is upgraded, I/O to those chunklets at the time of the upgrade results in data loss. This is not applicable to other RAID types.
- Specify the `-w` option when upgrading unadmitted PDs.



# 31

## Wait Command

---

In this chapter

`waittask`

**31.2**

---

**COMMAND**

`waittask`

**DESCRIPTION**

The `waittask` command asks the CLI to wait for a task to complete before proceeding. The command automatically notifies you when the specified task is finished.

**SYNTAX**

The syntax for the `waittask` command can be one of the following:

```
waittask -v <task_ID>
```

```
waittask [<task_ID>...]
```

**AUTHORITY**

Super, Service, Edit, Browse

**OPTIONS**

`-v <task_ID>`

Displays the detailed status of the task specified by `<task_ID>` as it executes. When the task completes, this command exits.

`-q`

Quiet; do not report the end state of the tasks, only wait for them to exit.

**SPECIFIERS**

`[<task_ID>...]`

Indicates one or more tasks to wait for using their task IDs. When no task IDs are specified, the command waits for all tasks to complete.

**RESTRICTIONS**

None.

## EXAMPLES

The following example shows how to wait for a task using the task ID. When successful, the command returns only after the task is complete.

```
cli% waittask 1  
Task 1 done
```

## NOTES

- See the *InForm OS CLI Administrator's Manual* for additional information and examples regarding task management and task management commands.
- This command returns an error if any of the tasks it is waiting for fail.





# Index

---

... (ellipsis), multiuse option specifier 2.2

() (parentheses), grouped elements 2.2

[] (brackets) optional elements 2.2

| (vertical bar), **or** option specifier 2.2

## A

access

    user information about 22.186

active VLUN 22.193

add

    virtual volume to group 5.10

add command

    See *also* specific commands (add\*)

addsnmpmgr 4.2

admin volume 22.173

admit command

    See *also* specific commands (admit\*)

admithw 5.2

admitpd 5.4

admitrcopylink 5.6

admitrcopytarget 5.8

admitrcopyvv 5.10

alerts

    removing from system 19.3

    service, registered SNMP software for 22.155

    system

        removing preregistered SNMP software 19.24

        status of 22.4

        status of, setting 21.3

        using addsnmpmgr for 4.2

API (CIM) 24.2, 26.2

arg (argument) specific option variable 2.2

argument (arg). See -arg (argument) specific option variable

audience for this reference 1.2

authentication

    parameters 21.5

## B

backing up virtual volumes 11.58

battery

    setting information about 21.12

    status of 22.8

brackets ( [] ), optional elements 2.2

**C**

cages, new firmware, downloading into 30.2

cancel command

See *also* specific commands (cancel\*)

canceltask 6.2

CBIOS version 3.10

chapters, overview of 1.4

characters, length limits 2.3

check commands

See *also* specific commands (check\*)

checkhealth 7.2

checkId 7.2, 7.4

checkpassword 7.6

checkpd 7.8

checkport 7.11

checkvv 7.13

chunklets

lists of

moving 17.2

spares, removing from 19.28

moving from physical disk to physical disk 17.12

moving to spare 17.18

spares 22.164

allocating resources as 11.56

moving to 17.18

removing from lists of 19.28

statistics collection mode, setting 21.72, 21.74

statistics, timed loop display 25.2

CIM API 26.2

CLI (command line interface)

syntax for 2.2

cli command 8.1

clusters, shutting down 23.4

cmd (command) operations 2.2

syntax for 2.3

command (cmd). See cmd (command)

command line interface (CLI). See CLI (command line interface)

commands 7.1

startcopygroup 26.7

stopcopy 22.144

See individual commands by name

Common Provisioning Group (CPG), length limit 2.3

common provisioning groups

creating 11.22

displaying 22.26

modifying 21.20

removing 19.5

compact commands

See *also* specific commands (compact\*)

compactcpg 9.2

compactId 9.4

configuration details, display 22.143

control commands 10.1

See *also* specific commands (control\*)

controliscsiport 10.2

controlmag 10.5

controlpd 10.7

controlport 10.9

copying virtual volumes 11.90

CPU speed 3.10

CPU statistics 25.10

create

group 11.48

link to remote storage server 5.6

create commands

See *also* specific commands (create\*)

createald 11.3

createaldvv 11.9

createavv 11.18

createcpg 11.22, 11.30

createdomain 11.30

creategroupsv 11.34

creategroupvvcopy 11.37

createhost 11.40

createhostset 11.43

createld 11.45

creatercopygroup 11.48

creatercopytarget 11.50

createsched 11.52

createspare 11.56 to 11.58

createsv 11.58 to 11.60

createtemplate 11.61

createtpvv 11.71

createuser 11.77

createvlun 11.79

createvv 11.83

createvvcopy 11.90

createvvset 11.94

creating domains 11.30

creating new system hosts 11.40

creating new users 11.77

creating provisioned virtual volumes 11.83

creating templates 11.61

## D

data services

starting on logical disk 24.3

virtual volume, starting 24.7, 24.8

databases, current firmware levels,  
displaying 22.45

date and time

for system nodes 22.33

setting nodes 21.27

definitions, removing 12.2

deleting VLUNs 19.39

designate site as primary or secondary 24.4

disabling 26.4

disks

logical

chunklets, displaying information  
about 22.66

creating 11.3, 11.9

creating virtual volumes in 11.83

data services, starting on 24.3

I/O statistics 25.18

information about, displaying 22.58

system pool, removing from 19.16

virtual volumes

region mapping information 22.215

physical

allocatable, marking as 21.53

chunklets, moving to another 17.12

creating definitions of 5.4

definitions, removing 12.2

firmware, upgrading 30.4

I/O statistics 25.25

information about, displaying 22.91

moving to spare 17.6, 17.9

spinning up or down 10.7

- virtual volumes, mapping information 22.115
- replacing 10.5
- dismiss command 12.1
  - See *also* specific commands (dismiss\*)
- dismisspd 12.2
- dismissrcopylink 12.3
- dismissrcopytarget 12.5
- dismissrcopyvv 12.6
- display
  - configuration details 22.143
  - statistics for group 25.39
- displaying 22.107
- displaying administration network interface
  - status and configuration 22.75
- documentation, related 1.3
- domains
  - creating 11.30
- drive cages
  - identifying 16.2
  - information about, displaying 22.17
- drive magazines
  - taking on- or off- loop 10.5

## E

- EEPROM log information 22.38
- enable Remote Copy 24.4
- enabling physical disks 5.4
- event logs
  - removing from system 19.10
  - system, displaying of 22.41

## F

- Fibre Channel ports
  - arbitrated loops 22.132

- controlling 10.9
- firmware
  - levels of 22.45
  - new, downloading into cages 30.2
  - physical disks, upgrading 30.4
  - specific cages, downloading into 30.2
- freespace 13.2

## G

- group
  - add virtual volume to 5.10
  - create 11.48
  - display statistics for 25.39
  - remove 19.18
  - remove virtual volume from 12.6
- grow commands
  - See *also* specific commands (grow\*)
- growaldv 14.2
- growavv 14.10
- growth increment (CPG) 21.21
- growth limit (CPG) 21.21, 22.30
- growth warning (CPG) 21.21
- growtpvv 14.12
- growvv 14.14

## H

- histch 15.2
- histld 15.6
- histograms 15.2
  - access times, logical disks 15.6
  - commands
    - See *also* specific commands (hist\*)
  - service times
    - displaying 15.2

- for ports 15.18
  - physical disks, of 15.10
  - VLUNs 15.23
- histpd 15.10
- histport 15.18
- histvlun 15.23
- histvv 15.28
- host name, length limit 2.3
- host sees VLUN template 11.79
- hosts
  - creating new system 11.40
  - removing 19.12
  - system configuration, displaying 22.47, 22.56
  - VLUN template conflicts and 11.42

## I

- I/O statistics
  - logical disks 25.18
  - physical disks 25.25
  - ports 25.33
  - virtual volumes 25.48
- identify primary and secondary storage servers 11.50
- identifying drive cages 16.2
- interfaces, user 1.2
- IP address, specifying SNMP additions 4.3

## K

- kernel build text 3.10
- kernel version number 3.10

## L

- length limits 2.3
- link

- create to remote storage server 5.6
  - remove 12.3
- link utilization statistics 25.23
- LIP 10.10
- locate command
  - See also* specific commands (locate\*)
- locatecage 16.2
- locatesys 16.4
- logical disk name, length limit 2.3
- logical disks
  - access times, histogram of 15.6
  - check data on 7.4
  - creating 11.3, 11.9
  - creating virtual volumes in 11.83
  - data services, starting on 24.3
  - displaying mapping information 22.14
  - I/O statistics 25.18
  - information about, displaying 22.58
  - system pool, removing from 19.16
- logs, event
  - removing from the system 19.10
  - system, displaying of 22.41
- loop initialization 10.10
- LUN host attachments, virtual volumes, statistics for 25.42

## M

- mapping, virtual volumes
  - on physical disks 22.115
  - region information on 22.215
- matched set VLUN template 11.79
- move commands
  - See also* specific commands (move\*)
- movech 17.2

movechtodomain 17.18

movechtospare 17.6

movepdtospare 17.9

moverelocpd 17.12

moving chunklet lists 17.2

moving domains 17.18

moving physical disks to spare 17.6, 17.9

multiuse option specifier. See . . . (ellipsis)  
multiuse option specifier

## N

Nearline drives 11.27, 14.8, 15.15, 21.25, 21.26, 22.63,  
22.95, 22.105, 25.30, 28.6

new users, creating 11.77

nodenames 3.10

nodes

- CPU statistics for 25.10
- date and time, setting 21.27
- EEPROM log information 22.38
- system, shutting down 23.2

notes, explanation of 1.6

## O

operating environment status 22.87

optional elements. See brackets ( [ ] )

options

- command line element 2.2
- 2.2

overview of chapters 1.4

## P

parentheses (), grouped elements 2.2

passwords

- setting 21.51

## SNMP

- removing access 19.26
- updating access 21.68

persistent repository 22.173

physical disks 12.2, 17.6, 17.9

- allocatable, marking as 21.53
- displaying mapping information 22.14
- firmware, upgrading 30.4
- I/O statistics 25.25
- information about, displaying 22.91
- moving chunklets to another 17.12
- moving to spare 17.6, 17.9
- spinning up or down 10.7
- virtual volumes, mapping information 22.115

port presents VLUN template 11.79

## ports

Fibre Channel

- arbitrated loops 22.132
- controlling 10.9
- histograms of service times for 15.18
- I/O statistics 25.33

PR (persistent repository) 22.173

preserved data 22.107

primary site

- designate 24.4

primary storage server

- identify 11.50

promotesv 18.2

promotevvcopy 18.4

## R

raw space alert 21.75

Registered State Change Notification  
(RSCN) 10.10, 22.124

related documents 1.3

Remote Copy 26.4

enable 24.4

ports 10.9

stop 26.6

remove

group 19.18

link 12.3

target designation 19.20

virtual volume from group 12.6

remove commands

See *also* specific commands (remove\*)

removealert 19.3

removecpg 19.5

removedomain 19.7, 19.8

removedomainset 19.8

removeeventlog 19.10

removehost 19.12, 19.14

removehostset 19.14

removeld 19.16

removercopygroup 19.18

removercopytarget 19.20

removesched 19.22

removesnmpmgr 19.24

removesnmppw 19.26

removespare 19.28

removesshkey 19.30, 19.46

removetask 19.31

removetemplate 19.33

removeuser 19.35

removeuserconn 19.37

removevln 19.39

removevv 19.43

removing hosts 19.12

removing users 19.35

removing virtual volumes 19.43

replacing drive magazines or disk drives 10.5

revision history RH.1

RSCN 10.10, 22.124

rules, syntax 2.3

## S

SCSI-3 reservation types 22.152

secondary site

designate 24.4

secondary storage server

identify 11.50

service commands

See *also* specific command (servicemag\*)

service times

chunklet timed loops, histograms of 15.2

physical disks, histograms of 15.10

servicecage 20.2

servicehost 20.5

servicemag 20.8

set commands

See *also* specific commands (set\*)

setalert 21.3

setauthparam 21.5

setbattery 21.12

setcage 21.14

setcim 21.16

setclienv 21.18

setcpg 21.20  
setdate 21.27  
setdomain 21.31, 21.33  
setdomainset 21.33  
sethost 21.35  
sethostset 21.39  
setlicense 21.39, 21.40  
setnet 21.42  
setnode 21.47  
setntp 2.5 to 2.6, 21.49  
setpassword 21.51  
setpd 21.53  
setrcopygroup 21.55  
setrcopytarget 21.62  
setsched 21.66  
setsnmppw 21.68  
setsshkey 21.70  
setstatch 21.72  
setstatpdch 21.74  
setsys 21.75  
setsysmgr 21.79  
settemplate 21.82  
setting the administration network interface configuration 21.42  
setting the system node properties 21.47  
setting the system NTP server 21.49  
setuser 21.84, 21.94  
setuseracl 21.86  
setvv 21.88  
setvvset 21.94  
show commands  
    *See also specific commands (show\*)*  
showalert 22.4  
showauthparam 22.6  
showbattery 22.8  
showblock 22.14  
showcage 22.17  
showcim 22.23  
showclienv 22.25  
showcpg 22.26  
showdate 22.33  
showdomain 22.34  
showdomainset 22.36  
showeeprom 22.36, 22.38  
showeventlog 22.41  
showfirmwaredb 22.45  
showhost 22.47  
showhostset 22.52  
showiscsisession 22.56  
showld 22.58  
showldch 22.66  
showldmap 22.71  
showlicense 22.73  
shownet 22.75  
shownode 22.77  
shownodeenv 22.87  
showpatch 22.89  
showpd 22.91  
showpdata 22.107  
showpdch 22.108



showpdvv 22.115

showport 22.119

showportarp 22.130

showportdev 22.132

showportisns 22.135

showportlesb 22.137

showrcopy 22.143

showrctransport 22.148

showrsv 22.151

showsched 22.153

showsnmpmgr 22.155

showsnmpppw 22.157

showspace 22.159

showspare 22.164

showsys 22.167, 22.169

showsysmgr 22.174

showtarget 22.177

showtask 22.178

showtemplate 22.182

showtoc 22.183

showtocgen 22.185

showuser 22.186

showuseracl 22.188

showuserconn 22.189

showversion 22.191

showvlun 22.193

showvv 22.199

showvvmap 22.215, 22.222

showvvpd 22.217

showvvset 22.222

shutdown commands

*See also* specific commands (shutdown\*)

shutdownnode 23.2

shutdownsys 23.4

shutting down system nodes 23.2

snapshot volumes (copies), creating 11.58

SNMP access

removing passwords for 19.26

updating 21.68

SNMP software

for alerts 22.155

removing preregistered 19.24

spares, allocating chunklet resources as 11.56

spec (specifier) required command line element,  
used with 2.2

specifier (spec). *See* spec (specifier) required  
command line element, used with

start commands

*See also* specific commands (start\*)

startcim 24.2

startld 24.3

startrcopy 24.4

startcopygroup 24.5

startcopygroup command 26.7

starttask 24.7

startup date for system manager, setting 21.79

start-up state information, system  
manager 22.174

startvv 24.7, 24.8

stat commands

*See also* specific commands (stat\*)

statch 25.2

- statcmp 25.7
  - statcpu 25.10
  - statiscsi 25.12
  - statiscsisession 25.16
  - statistics
    - cache memory page 25.7
    - chunklet collection mode 21.72, 21.74
    - chunklets, timed loop display 25.2
    - CPU, for all nodes 25.10
    - I/O
      - logical disk 25.18
      - physical disks 25.25
      - virtual volumes 25.48
    - I/O port 25.33
    - link utilization 25.23
    - virtual volumes
      - I/O 25.48
      - LUN host attachments 25.42
  - statistics, display for group 25.39
  - statld 25.18
  - statlink 25.23
  - statpd 25.25
  - statport 25.33
  - statrcopy 25.39
  - status of system alerts 22.4
  - statvlun 25.42, 25.47
  - statvv 25.48
  - stop commands
    - See also* specific commands (stop\*)
  - stopcim 26.2
  - stoping Remote Copy 26.6
  - stoprcopy 26.4
  - stoprcopy command 22.144
  - stoprcopygroup 26.6
  - storage server, information about 22.191
  - subcmd (subcommand) operations 2.2
  - subcommand (subcmd). *See* subcmd (subcommand)
  - sync command 27.1 to 27.3
    - See also* specific commands (sync\*)
  - syncrcopy 27.2 to 27.3
  - syntax, rules for 2.3
  - system alerts
    - removing 19.3
    - setting 21.3, 21.5
    - status of 22.4
  - system configuration, displaying for hosts 22.47, 22.56
  - system information 3.10
  - system manager
    - startup date, setting 21.79
    - start-up state information 22.174
  - system nodes
    - date and time 22.33
    - shutting down 23.2
  - system resources, summary table of 22.183
  - system table of contents, summary of 22.183
- T**
- table of contents generation number 22.185
  - target designation, remove 19.20
  - targets, unrecognized 22.177
  - temperature status 22.87
  - template name, length limit 2.3
  - template VLUN 22.193

**templates**

- creating 11.61

Thinly Provisioned Virtual Volume (TPVV), length limit 2.3

**time and date**

- for system nodes 22.33

- setting nodes 21.27

**traps**

- removing preregistered SNMP software 19.24

- using addsnmpmgr for 4.2

tunealdivv 28.2

tunepd 28.10

tunetpvv 28.15

tunevv 28.18

**U**

updatesnapspace 29.2

updatevv 29.4

**upgrade commands**

- See *also* specific commands (upgrade\*)

upgradecage 30.2

upgradedpd 30.4

upgrades, firmware levels 22.45

upgrading physical disk firmware 30.4

user interfaces 1.2

**users**

- changing passwords 21.51

- information about 22.186

- new, creating 11.77

- removing 19.35

**V****validity checks**

- logical disk data 7.4

- of virtual volumes 7.13

VCN 10.10, 22.124

**verifying**

- controlport command success 10.18

- physical disk admittance 5.5

- physical disk status 21.54

vertical bar (|), **or** option specifier 2.2

virtual copy names, length limit 2.3

virtual volume names, length limit 2.3

**virtual volumes**

- backing up 11.58

- copy (snapshot) of 11.58

- copying to another virtual volume 11.90

- creating 11.9, 11.18

- creating from logical disk 11.83

- creating VLUN of 11.79

- data services, starting on 24.7, 24.8

- displaying mapping information 22.14

- histograms of service time 15.28

- I/O statistics 25.48

- information about, displaying distribution of 3.31, 22.217

- information about, displaying of 22.199

- LUN statistics 25.48

- physical disk mapping information 22.115

- region mapping information 22.215

- removing 19.43

- statistics

- I/O 25.48

- LUN host attachments 25.42

- validity checks of 7.13

VLUN change notification 10.10

VLUN Change Notification (VCN) 22.124

VLUN templates

host sees 11.79  
matched set 11.79  
port presents 11.79

## VLUNs

active 22.193  
deleting 19.39  
information about 22.193  
of virtual volumes, creating 11.79  
service time, histograms of 15.23

template 22.193  
template conflicts in 11.42  
voltage status 22.87

## W

wait command  
    *See also* specific commands (wait\*)  
waittask 31.2

# Revision History

---

Release Level	Revision Summary
320-200166 Rev A October 2009	First release of this manual to support initial release of the 2.3.1 InForm OS.
320-200166 Rev B March 2010	Second release of this manual to support the 2.3.1 MU1 InForm OS release.

