FreeNAS® 11.2-RELEASE User Guide

December 2018 Edition

FreeNAS[®] is © 2011-2018 iXsystems FreeNAS[®] and the FreeNAS[®] logo are registered trademarks of iXsystems FreeBSD[®] is a registered trademark of the FreeBSD Foundation Written by users of the FreeNAS[®] network-attached storage operating system. Version 11.2 Copyright © 2011-2018 iXsystems (https://www.ixsystems.com/)

CONTENTS

	Welcome	
1	Introduction 1.1 New Features in 11.2 1.2 Path and Name Lengths 1.3 Hardware Recommendations 1.3.1 RAM 1.3.2 The Operating System Device 1.3.3 Storage Disks and Controllers 1.3.4 Network Interfaces 1.4 Getting Started with ZFS	12 13 14 14 15 16
2	Installing and Upgrading 2.1 Getting FreeNAS [®] 2.2 Preparing the Media 2.2.1 On FreeBSD or Linux 2.2.2 On Windows 2.2.3 macOS 2.4 Installation 2.4 Installation Troubleshooting 2.5 Upgrading 2.5.1 Caveats 2.5.2 Initial Preparation 2.5.3 Upgrading From the GUI 2.5.4 Upgrading From the GUI 2.5.5 If Something Goes Wrong 2.5.6 Upgrading a ZFS Pool 2.6.1 Virtualization 2.6.2 VMware ESXi	18 19 19 20 28 29 29 29 30 33 33 35 36 37
3	Booting 3.1 Obtaining an IP Address 3.2 Logging In 3.3 Initial Configuration	56
4	Account 4.1 Groups	
5	System 5.1 Information	

	5.3	Boot
		5.3.1 Mirroring the Boot Device
	5.4	Advanced
		5.4.1 Autotune
		5.4.2 Self-Encrypting Drives
	5.5	Email
	5.6	System Dataset
	5.7	Tunables
	5.8	Update
		5.8.1 Preparing for Updates
		5.8.2 Updates and Trains
		5.8.3 Checking for Updates
		5.8.4 Applying Updates
		5.8.5 Manual Updates
	5.9	Cloud Credentials
	5.10	Alerts
	5.11	Alert Services
		5.11.1 How it Works
	5.12	CAs
	5.13	Certificates
	5.14	Support
	Task	
	6.1	Cloud Sync
		6.1.1 Cloud Sync Example
	6.2	Cron Jobs
	6.3	Init/Shutdown Scripts
	6.4	Rsync Tasks
		6.4.1 Rsync Module Mode
		6.4.2 Rsync over SSH Mode
	6.5	6.4.2 Rsync over SSH Mode
7		S.M.A.R.T. Tests
7	Net	S.M.A.R.T. Tests
	Netv 7.1	S.M.A.R.T. Tests 110 work 113 Global Configuration 113
	Netv 7.1 7.2	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115
	Netv 7.1 7.2 7.3	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117
	Netv 7.1 7.2	S.M.A.R.T. Tests
	Netv 7.1 7.2 7.3	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi
	Netv 7.1 7.2 7.3 7.4	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120
	Netv 7.1 7.2 7.3 7.4	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123
	Netv 7.1 7.2 7.3 7.4 7.5 7.6	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123
	Netv 7.1 7.2 7.3 7.4	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123
	Netv 7.1 7.2 7.3 7.4 7.5 7.6	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 123 Static Routes 123 VLANs 123
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 age 125 Volumes 125
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 age 125 Volumes 125 8.1.1 Volume Manager 125
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 8.1.1 Volume Manager 125 8.1.1 Encryption 127
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 123 Static Routes 123 123 vLANs 123 125 8.1.1 Volume Manager 125 8.1.1 Encryption 127 8.1.2 Encryption Performance 128
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 8.1.1 Volume Manager 125 8.1.1 Encryption Performance 127 8.1.1.3 Manual Setup 128
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 117 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 static Routes 125 8.1.1 Volume Manager 125 8.1.1.1 Encryption 127 8.1.1.2 Encryption 127 8.1.1.3 Manual Setup 128 8.1.1.4 Extending a ZFS Volume 129
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 static Routes 125 8.1.1 Volume Manager 125 8.1.1.1 Encryption 127 8.1.1.2 Encryption 127 8.1.1.3 Manual Setup 128 8.1.1.4 Extending a ZFS Volume 129 8.1.2 Change Permissions 130
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 117 Yolumes 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 123 Static Routes 123 123 ValANs 123 123 age 125 8.1.1 Volume Manager 125 8.1.1.1 Encryption Performance 128 8.1.1.2 Encryption Performance 128 8.1.1.2 Encryption a ZFS Volume 129 128 130 8.1.2 Change Permissions 130 8.1.3 130 8.1.3 Create Dataset 132 132
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 113 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 123 Static Routes 123 123 VLANs 123 123 age 125 8.1.1 Volumes 125 8.1.1.1 Encryption 127 8.1.1.2 Encryption Performance 128 8.1.1.2 Encryption Performance 128 128 129 8.1.2 Change Permissions 130 130 8.1.3 Grade Dataset 130 130 8.1.3 Deduplication 132 134
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 123 Static Routes 123 123 VLANs 123 123 age 125 8.1.1.1 Encryption Volumes 125 8.1.1.1 Encryption 8.1.1.2 Encryption Performance 128 8.1.1.3 Manual Setup 128 8.1.1.4 Extending a ZFS Volume 129 8.1.2 Change Permissions 130 8.1.3 Deduplication 134 8.1.3.1 Deduplication 134 8.1.3.2 Compression 135
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 age 125 8.1.1 Notume Manager 125 8.1.1 Encryption 126 8.1.1.2 Encryption Performance 128 8.1.1.4 Extending a ZFS Volume 129 8.1.2 Change Permissions 130 8.1.3 Deduplication 132 8.1.3 Deduplication 134 8.1.3 Deduplication 134 8.1.3 Compression 135 8.1.4 Create zvol 135
	Netv 7.1 7.2 7.3 7.4 7.5 7.6 7.7 Stor	S.M.A.R.T. Tests 110 work 113 Global Configuration 113 Interfaces 115 IPMI 117 Link Aggregations 119 7.4.1 LACP, MPIO, NFS, and ESXi 119 7.4.2 Creating a Link Aggregation 120 Network Summary 123 Static Routes 123 VLANs 123 age 125 8.1.1 Encryption 8.1.1 Encryption 8.1.1.1 Encryption Performance 8.1.1.3 Maual Setup 8.1.1.4 Extending a ZFS Volume 8.1.2 Change Permissions 8.1.3 Deduplication 8.1.3 Deduplication 8.1.4 Create zvol 8.1.5 Import Disk

		8.1.7 View Disks
		8.1.8 Volumes
		8.1.8.1 Managing Encrypted Volumes
		8.1.8.2 Additional Controls for Encrypted Volumes
		8.1.9 View Multipaths
		8.1.10 Replacing a Failed Drive
		8.1.10.1 Replacing an Encrypted Drive
		8.1.10.2 Removing a Log or Cache Device
		8.1.11 Replacing Drives to Grow a ZFS Pool
		8.1.12 Hot Spares
	8.2	Periodic Snapshot Tasks
	8.3	Replication Tasks
	0.5	8.3.1 Examples: Common Configuration
		8.3.1.1 <i>Alpha</i> (Source)
		8.3.1.2 <i>Beta</i> (Destination)
		8.3.2 Example: FreeNAS [®] to FreeNAS [®] Semi-Automatic Setup
		8.3.3 Example: FreeNAS [®] to FreeNAS [®] Dedicated User Replication
		8.3.4 Example: FreeNAS [®] to FreeNAS [®] or Other Systems, Manual Setup
		8.3.4.1 Encryption Keys
		8.3.5 Replication Options
		8.3.6 Replication Encryption
		8.3.7 Limiting Replication Times
		8.3.8 Troubleshooting Replication
		8.3.8.1 SSH
		8.3.8.2 Compression
		8.3.8.3 Manual Testing
	8.4	Resilver Priority
	8.5	Scrubs
	86	Spanshots 165
	8.6	Snapshots
		8.6.1 Browsing a snapshot collection
	8.6 8.7	
0	8.7	8.6.1 Browsing a snapshot collection
9	8.7 Dire	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169
9	8.7	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169
9	8.7 Dire	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173
9	8.7 Dire	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173
9	8.7 Dire	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173
9	8.7 Dire 9.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173
9	8.7 Dire 9.1 9.2	8.6.1Browsing a snapshot collection167VMware-Snapshot169ctory Services169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176
9	8.7 Dire 9.1 9.2 9.3 9.4	8.6.1Browsing a snapshot collection167VMware-Snapshot167ctory Services169Active Directory1699.1.1Troubleshooting Tips9.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177
9	 8.7 Diree 9.1 9.2 9.3 9.4 9.5 	8.6.1Browsing a snapshot collection167VMware-Snapshot167ctory Services169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177Kerberos Keytabs178
9	8.7 Dire 9.1 9.2 9.3 9.4	8.6.1Browsing a snapshot collection167VMware-Snapshot167ctory Services169Active Directory1699.1.1Troubleshooting Tips9.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6	8.6.1Browsing a snapshot collection167VMware-Snapshot169ctory Services169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177Kerberos Keytabs178Kerberos Settings179
	 8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 174 NIS 176 176 Kerberos Realms 177 Kerberos Settings 178 Kerberos Settings 179 Ting 180
	 8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 	8.6.1Browsing a snapshot collection167VMware-Snapshot167ctory Services169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177Kerberos Settings178Kerberos Settings179ing180Apple (AFP) Shares181
	 8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 174 NIS 176 176 Kerberos Realms 177 Kerberos Settings 178 Kerberos Settings 179 ing 180 Apple (AFP) Shares 181 10.1.1 Creating AFP Guest Shares 183
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 177 Kerberos Keytabs 177 Nig 178 Kerberos Settings 178 10.1.1 Creating AFP Guest Shares 181 10.1.2 Creating Authenticated and Time Machine Shares 184
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1Browsing a snapshot collection167VMware-Snapshot169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177Kerberos Settings178Kerberos Settings179ing180Apple (AFP) Shares18110.1.1Creating Athenticated and Time Machine Shares184Unix (NFS) Shares184Unix (NFS) Shares188
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1Browsing a snapshot collection167VMware-Snapshot169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177Kerberos Realms178Kerberos Settings179ing180Apple (AFP) Shares18110.1.1Creating AtPP Guest Shares18310.1.2Creating Athenticated and Time Machine Shares184Unix (NFS) Shares18810.2.1Example Configuration190
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 177 Kerberos Realms 177 Kerberos Settings 178 10.1.1 Creating AFP Guest Shares 180 10.1.2 Creating Authenticated and Time Machine Shares 184 Unix (NFS) Shares 184 10.2.1 Example Configuration 190 10.2.2 Connecting to the Share 191
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1Browsing a snapshot collection167VMware-Snapshot169Active Directory1699.1.1Troubleshooting Tips1739.1.2If the System Does not Join the Domain173LDAP174NIS176Kerberos Realms177Kerberos Realms178Kerberos Settings179ing180Apple (AFP) Shares18110.1.1Creating AtPP Guest Shares18310.1.2Creating Athenticated and Time Machine Shares184Unix (NFS) Shares18810.2.1Example Configuration190
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 177 Kerberos Realms 177 Kerberos Settings 178 10.1.1 Creating AFP Guest Shares 180 10.1.2 Creating Authenticated and Time Machine Shares 184 Unix (NFS) Shares 184 10.2.1 Example Configuration 190 10.2.2 Connecting to the Share 191
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 176 Kerberos Realms 177 Kerberos Settings 178 Napple (AFP) Shares 180 10.1.1 Creating AFP Guest Shares 183 10.1.2 Creating Atthenticated and Time Machine Shares 184 Unix (NFS) Shares 188 10.2.1 Example Configuration 190 10.2.2.2 From Microsoft 191
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 174 NIS 176 174 Kerberos Realms 177 Kerberos Settings 177 ing 180 Apple (AFP) Shares 181 10.1.1 Creating Authenticated and Time Machine Shares 183 10.2.2 Cronnecting to the Share 191 10.2.2.3 From Microsoft 191 10.2.2.3 From macOS 192
	8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 175 Kerberos Realms 176 Kerberos Keytabs 176 Kerberos Settings 177 10.11 Creating Authenticated and Time Machine Shares 181 10.1.2 Creating Authenticated and Time Machine Shares 188 10.2.1 Example Configuration 190 10.2.2.2 From Microsoft 191 10.2.2.3 From macOS 192 10.2.3 Troubleshooting NFS 193
	 8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1 10.2 10.3 	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 177 Kerberos Keytabs 178 Kerberos Settings 179 ting 180 Apple (AFP) Shares 181 10.1.1 Creating Atthenticated and Time Machine Shares 183 10.2.2 Connecting to the Share 191 10.2.2.1 From BSD or Linux 191 10.2.2.3 From Microsoft 191 10.2.2.3 From MacOS 192 10.2.3 Troubleshooting NFS 193 WebDAV Shares 193 194
	 8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1 10.2 10.3 	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 167 ctory Services 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 177 Kerberos Keytabs 177 Kerberos Settings 177 ing 180 Apple (AFP) Shares 181 10.1.1 Creating AtPP Guest Shares 183 10.1.2 Creating AtPP Guest Shares 184 Unix (NFS) Shares 188 182 10.2.2 Connecting to the Share 191 10.2.2.1 From Microsoft 191 10.2.2.2 100 192 10.2.3 Troubleshooting NFS 193 WebDAV Shares 194 193 Windows (SMB) Shares 194
	 8.7 Dire 9.1 9.2 9.3 9.4 9.5 9.6 Shar 10.1 10.2 10.3 	8.6.1 Browsing a snapshot collection 167 VMware-Snapshot 169 Active Directory 169 9.1.1 Troubleshooting Tips 173 9.1.2 If the System Does not Join the Domain 173 LDAP 174 NIS 176 Kerberos Realms 177 Kerberos Keytabs 178 Kerberos Settings 179 ting 180 Apple (AFP) Shares 181 10.1.1 Creating Atthenticated and Time Machine Shares 183 10.2.2 Connecting to the Share 191 10.2.2.1 From BSD or Linux 191 10.2.2.3 From Microsoft 191 10.2.2.3 From MacOS 192 10.2.3 Troubleshooting NFS 193 WebDAV Shares 193 194

10.4.3 Configuring Shadow Copies	
10.5 Block (iSCSI)	
10.5.1 Target Global Configuration	
10.5.2 Portals	
10.5.4 Authorized Accesses	
10.5.5 Targets	
10.5.6 Extents	
10.5.7 Target/Extents	
10.5.8 Connecting to iSCSI	
10.5.9 Growing LUNs	
10.5.9.1 Zvol Based LUN	
10.5.9.2 File Extent Based LUN	. 218
11 Services	219
11.1 Control Services	
11.2 AFP	
11.2.1 Troubleshooting AFP	
11.3 Domain Controller	
11.3.1 Samba Domain Controller Backup	
11.4 Dynamic DNS	
11.5 FTP	
11.5.1 Anonymous FTP	
11.5.2 FTP in chroot	
11.5.3 Encrypting FTP	
11.6 iSCSI	
11.7 LLDP	
11.8 Netdata	
11.9 NFS	
11.10Rsync	. 234
11.10.1 Configure Rsyncd	
11.10.2 Rsync Modules	
11.1153	
11.12S.M.A.R.T.	
11.13SMB	
11.13.1 Troubleshooting SMB	
11.15SSH	
11.15.1 SCP Only	
11.15.2 Troubleshooting SSH	
11.16TFTP	
11.17UPS	. 248
11.17.1 Multiple Computers with One UPS	
11.18WebDAV	. 252
12 Plugins	254
12.1 Installed Plugins	
12.2 Deleting Plugins	
13 Jails	256
13.1 Jails Configuration	
13.2 Managing Jails	
13.2.1 Accessing a Jail Using SSH 13.2.2 Add Storage	
13.3 Starting Installed Software	
	. 204

14 Virtual Machines

 14.1 Creating VMs. 14.2 Adding Devices to a VM. 14.2.1 Network Interfaces. 14.2.2 Disk Devices. 14.2.3 Raw Files. 14.2.4 CD-ROM Devices 14.2.5 VNC Interface 14.2.6 Virtual Serial Ports 14.3 Running VMs. 14.4 Deleting VMs. 14.5 Docker VM 14.5.1 Docker VM Requirements 14.5.2 Create the Docker VM 14.5.3 Start the Docker VM 14.5.4 SSH into the Docker VM 14.5.5 Installing and Configuring the Rancher Server 	. 267 . 268 . 269 . 269 . 270 . 270 . 272 . 273 . 273 . 273 . 273 . 273 . 278 . 278 . 278
15 Reporting	279
16 Wizard	281
17 Display System Processes	288
18 Shell	289
19 Log Out	291
20 Reboot	292
21 Shutdown	293
22 Support Icon	294
23 Guide	295
24 Alert	296
25 Support Resources 25.1 Website and Social Media 25.2 Forums 25.3 IRC 25.4 Videos 25.5 Professional Support	. 298 . 299 . 299
26 Command Line Utilities 26.1 lperf. 26.2 Netperf 26.3 lOzone 26.4 arcstat 26.5 tw_cli 26.6 MegaCli 26.7 freenas-debug 26.8 tmux 26.9 Dmidecode 26.10Midnight Commander	. 303 . 304 . 306 . 311 . 312 . 313 . 313 . 314
27.1 Translation	

	VAAI 29.1 VAAI for iSCSI	324 . 324
	Using the API 30.1 A Simple API Example 30.2 A More Complex Example	
Ind	dex	329

Welcome

This Guide covers the installation and use of FreeNAS[®] 11.2.

The FreeNAS[®] User Guide is a work in progress and relies on the contributions of many individuals. If you are interested in helping us to improve the Guide, read the instructions in the README (https://github.com/freenas/freenas-docs/blob/master/README.md). IRC Freenode users are welcome to join the *#freenas* channel where you will find other FreeNAS[®] users.

The FreeNAS[®] User Guide is freely available for sharing and redistribution under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/3.0/). This means that you have permission to copy, distribute, translate, and adapt the work as long as you attribute iXsystems as the original source of the Guide.

FreeNAS[®] and the FreeNAS[®] logo are registered trademarks of iXsystems.

Active Directory[®] is a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries.

Apple, Mac and Mac OS are trademarks of Apple Inc., registered in the U.S. and other countries.

Broadcom is a trademark of Broadcom Corporation.

Chelsio[®] is a registered trademark of Chelsio Communications.

Cisco[®] is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

Django[®] is a registered trademark of Django Software Foundation.

Facebook[®] is a registered trademark of Facebook Inc.

FreeBSD[®] and the FreeBSD[®] logo are registered trademarks of the FreeBSD Foundation[®].

Intel, the Intel logo, Pentium Inside, and Pentium are trademarks of Intel Corporation in the U.S. and/or other countries.

LinkedIn[®] is a registered trademark of LinkedIn Corporation.

Linux[®] is a registered trademark of Linus Torvalds.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

Twitter is a trademark of Twitter, Inc. in the United States and other countries.

UNIX[®] is a registered trademark of The Open Group.

VirtualBox[®] is a registered trademark of Oracle.

VMware[®] is a registered trademark of VMware, Inc.

Wikipedia[®] is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.

Windows[®] is a registered trademark of Microsoft Corporation in the United States and other countries.

Typographic Conventions

Typographic Conventions

The FreeNAS[®] 11.2 User Guide uses these typographic conventions:

Item	Visual Example
Graphical elements: buttons, icons, fields, columns, and boxes	Click the <i>Import CA</i> button.
Menu selections	Select System \rightarrow Information.
Commands	Use the scp command.
File names and volume and dataset names	Locate the /etc/rc.conf file.
Keyboard keys	Press the Enter key.
Important points	This is important.
Values entered into fields, or device names	Enter <i>127.0.0.1</i> in the address field.

Table 1: Text Format Examples

INTRODUCTION

FreeNAS[®] is an embedded open source network-attached storage (NAS) operating system based on FreeBSD and released under a 2-clause BSD license (https://opensource.org/licenses/BSD-2-Clause). A NAS has an operating system optimized for file storage and sharing.

FreeNAS[®] provides a browser-based, graphical configuration interface. The built-in networking protocols provide storage access to multiple operating systems. A plugin system is provided for extending the built-in features by installing additional software.

1.1 New Features in 11.2

FreeNAS[®] 11.2 is a feature release, which includes several new significant features, many improvements and bug fixes to existing features, and version updates to the operating system, base applications, and drivers. Users are encouraged to *Update* (page 79) to this release in order to take advantage of these improvements and bug fixes.

These major features are new in this version:

- The login screen defaults to the new, Angular-based UI. Users who wish to continue to use the classic UI can select "Legacy UI" in the login screen.
- Beginning with this release, the screenshots that appear in the published version of the Guide (http://doc.freenas.org/11.2/freenas.html) and in the *Guide* option within the new UI are for the new UI. However, users who click the *Guide* option while logged into the classic UI will continue to see screenshots for the old UI. The availability of both versions of the Guide is to assist users as they become familiar with the new UI during the transition period before the classic UI is deprecated in a future release.
- The rewrite from the old API to the new middlewared continues. Once the rewrite is complete, api.freenas.org (http://api.freenas.org/) will be deprecated and replaced by the new API documentation. In the mean time, to see the API documentation for the new middleware, log into the new UI, click on the URL for the FreeNAS system in your browser's location bar, and add :api/docs to the end of that URL.
- The boot loader has changed from GRUB to the native FreeBSD boot loader. This should resolve several issues that some users experienced with GRUB. GRUB was introduced as a temporary solution until the FreeBSD boot loader had full support for boot environments, which it now has.
- The *Plugins* (page 254) and *Jails* (page 256) backend has switched from warden to iocage and warden will no longer receive bug fixes. The new UI will automatically use iocage to create and manage *Plugins* (page 254) and *Jails* (page 256). Users are encouraged to recreate any existing *Plugins* (page 254) and *Jails* (page 256) using the new UI to ensure that they are running the latest supported application versions.
- Plugins (page 254) have switched to FreeBSD 11.2-RELEASE and all Plugins have been rebuilt for this version.
- *Virtual Machines* (page 266) are more crash-resistant. When a guest is started, the amount of available memory is checked and an initialization error will occur if there is insufficient system resources. When a guest is stopped, its resources are returned to the system. In addition, the UEFI boot menu fix allows Linux kernels 4.15 and higher to boot properly.
- *Cloud Sync* (page 95) provides configuration options to encrypt data before it is transmitted and to keep it in the encrypted format while stored on the cloud. The filenames can also be encrypted.

• Preliminary support has been added for Self-Encrypting Drives (page 74) (SEDs).

This software has been added or updated:

- The base operating system is the STABLE branch of FreeBSD 11.2 (https://www.freebsd.org/releases/11.2R/announce.html), which brings in many updated drivers and bug fixes. This branch has been patched to include the FreeBSD security advisories up to FreeBSD-SA-18:13.nfs (https://www.freebsd.org/security/advisories/FreeBSD-SA-18:13.nfs.asc).
- OpenZFS is up-to-date with Illumos and slightly ahead due to support for sorted scrubs which were ported from ZFS on Linux. Notable improvements include channel programs, data disk removal, more resilient volume import, the ability to import a pool with missing vdevs, pool checkpoints, improved compressed ARC performance, and ZIL batching. As part of this change, the default ZFS indirect block size is reduced to 32 KiB from 128 KiB. Note that many of these improvements need further testing so have not yet been integrated into the UI.
- The IPsec kernel module has been added. It can be manually loaded with kldload ipsec.
- Support for eMMC flash storage has been added.
- The em (https://www.freebsd.org/cgi/man.cgi?query=em&apropos=0&sektion=4), igb (https://www.freebsd.org/cgi/man.cgi?query=igb&apropos=0&sektion=4), ixgbe (https://www.freebsd.org/cgi/man.cgi?query=ixl&apropos=0&sektion=4) Intel drivers have been patched to resolve a performance degradation issue that occurs when the MTU is set to 9000 (9k jumbo clusters). Before configuring 9k jumbo clusters for cxgbe (https://www.freebsd.org/cgi/man.cgi?query=cxgb&apropos=0&sektion=4) create a *Tunables* (page 77) with a *Variable* of *hw.cxgbe.largest_rx_cluster*, a *Type* of *Loader*, and a *Value* of 4096. The cxgb (https://www.freebsd.org/cgi/man.cgi?query=cxgb&apropos=0&sektion=4) driver does not support jumbo clusters and should not use an MTU greater than 4096.
- The bnxt (https://www.freebsd.org/cgi/man.cgi?query=bnxt) driver has been added which provides support for Broadcom NetXtreme-C and NetXtreme-E Ethernet drivers.
- The vt terminal (https://www.freebsd.org/cgi/man.cgi?query=vt&sektion=4&manpath=FreeBSD+11.2-RELEASE+and+Ports) is now used by default and the syscons terminal is removed from the kernel.
- ncdu (https://dev.yorhel.nl/ncdu) has been added to the base system. This CLI utility can be used to analyze disk usage from the console or an SSH session.
- drm-next-kmod (https://www.freshports.org/graphics/drm-next-kmod/) has been added to the base system, adding support for UTF-8 fonts to the console for Intel graphic cards.
- Samba 4.7 has been patched to address the latest round of security vulnerabilities (https://www.samba.org/samba/latest_news.html#4.9.3).
- Netatalk has been updated to the 3.1.12 development version which addresses known issues with Time Machine timeouts.
- rsync has been updated to version 3.1.3 (https://download.samba.org/pub/rsync/src/rsync-3.1.3-NEWS).
- rclone has been updated to version 1.44 (https://rclone.org/changelog/#v1-44-2018-10-15).
- Minio has been updated to version 2018-04-04T05 (https://github.com/minio/minio/releases/tag/RELEASE.2018-04-04T05-20-54Z).
- Netdata as been updated to version 1.10.1 (https://github.com/firehol/netdata/releases/tag/v1.10.0).
- iocage has been synced with upstream as of October 3, providing many bug fixes and improved IPv6 support.
- RancherOS has been updated to version 1.4.2 (https://github.com/rancher/os/releases/tag/v1.4.2).
- zsh (http://www.zsh.org/) is the root shell for new installations. Upgrades will continue to use the csh shell as the default root shell.
- xattr (https://github.com/xattr/xattr) has been added to the base system and can be used to modify file extended attributes from the command line. Type <code>xattr -h</code> to view the available options.
- convmv (https://www.j3e.de/linux/convmv/man/) has been added to the base system and can be used to convert the encoding of filenames from the command line. Type convmv to view the available options.
- The cloneacl CLI utility has been added. It can be used to quickly clone a complex ACL recursively to or from an existing share. Type cloneacl for usage instructions.

- These switches have been added to *freenas-debug* (page 313): –M for dumping SATADOM info and –Z to delete old debug information. The –G switch has been removed as the system no longer uses GRUB. The –J switch has been removed and the –j switch has been reworked to show iocage jail information instead of Warden.
- These switches have been added to *arcstat* (page 306): -a for displaying all available statistics and -p for displaying raw numbers without suffixes.

These screen options have changed:

- The ATA Security User, SED Password, and Reset SED Password fields have been added to System \rightarrow Advanced.
- The *Enable screen saver* field has been removed from *System* \rightarrow *Advanced*.
- The Enable automatic upload of kernel crash dumps and daily telemetry checkbox has been removed from System \rightarrow Advanced.
- *Alerts* has been added to *System* and can be used to list the available alert conditions and to configure the notification frequency on a per-alert basis.
- These *Cloud Credentials* (page 82) have been added to *System* → *Cloud Credentials*: Amazon Cloud Drive, Box, Dropbox, FTP, Google Drive, HTTP, Hubic, Mega, Microsoft OneDrive, pCloud, SFTP, WebDAV, and Yandex.
- The *Team Drive ID* field has been added to *System* \rightarrow *Cloud Credentials* \rightarrow *Add* form when *Google Drive* is the *Provider*.
- The Endpoint URL has been added to System -> Cloud Credentials -> Add Cloud Credential but only appears when Amazon S3 is selected as the Provider. This can be used to configure a connection to another S3-compatible service, such as Wasabi.
- Drive Account Type and Drive ID has been added to System -> Cloud Credentials -> Add Cloud Credential. These fields appear when Microsoft OneDrive is selected as the Provider.
- The Automatically check for new updates option in System \rightarrow Update has been renamed to Check for Updates Daily and Download if Available.
- The Remote encryption, Filename encryption, Encryption password, and Encryption salt fields have been added to Tasks \rightarrow Cloud Sync Tasks \rightarrow Add Cloud Sync.
- The Exec field has been added to Storage \rightarrow Volumes \rightarrow Create Dataset \rightarrow Advanced Mode.
- The Password for SED column has been added to Storage \rightarrow Volumes \rightarrow View Disks.
- The MSDOSFS locale drop-down menu has been added to Storage \rightarrow Import Disk.
- The User Base and Group Base fields have been removed from Directory Services \rightarrow Active Directory \rightarrow Advanced Mode.
- The Enable home directories, Home directories, Home share name, and Home Share Time Machine fields have been removed from Services → AFP and the Time Machine Quota field has been removed from Sharing → Apple (AFP) Shares. These fields have been replaced by Sharing → Apple (AFP) Shares → Use as home share.
- The Umask field in Services \rightarrow TFTP has changed to a File Permissions selector.
- Disk temperature graphs have been added to *Reporting* \rightarrow *Disk*.

1.2 Path and Name Lengths

Names of files, directories, and devices are subject to some limits imposed by the FreeBSD operating system. The limits shown here are for names using plain-text characters that each occupy one byte of space. Some UTF-8 characters take more than a single byte of space, and using those characters reduces these limits proportionally. System overhead can also reduce the length of these limits by one or more bytes.

Туре	Maximum	Description
	Length	
File Paths	1024 bytes	Total file path length (<i>PATH_MAX</i>). The full path includes directory separa- tor slash characters, subdirectory names, and the name of the file itself. For example, the path /mnt/tank/mydataset/mydirectory/myfile. txt is 42 bytes long. Using very long file or directory names can be problematic. A complete path with long directory and file names can exceed the 1024-byte limit, preventing direct access to that file until the directory names or filename are shortened or the file is moved into a directory with a shorter total path length.
File and Directory Names	255 bytes	Individual directory or file name length (<i>NAME_MAX</i>).
Mounted Filesystem Paths	88 bytes	Mounted filesystem path length (<i>MNAMELEN</i>). Longer paths can prevent a device from being mounted or data from being accessible.
Device Filesystem Paths	63 bytes	devfs(8) (https://www.freebsd.org/cgi/man.cgi?query=devfs) device path lengths (<i>SPECNAMELEN</i>). Longer paths can prevent a device from being cre- ated.

Table 1.1: Path and Name Lengths

Note: 88 bytes is equal to 88 ASCII characters. The number of characters will vary when using Unicode.

Warning: If the mounted path length for a snapshot exceeds 88 bytes the data in the snapshot will be safe but inaccessible. When the mounted path length of the snapshot is less than the 88 byte limit, the data will be accessible again.

The 88 byte limit affects automatic and manual snapshot mounts in slightly different ways:

- Automatic mount: ZFS temporarily mounts a snapshot whenever a user attempts to view or search the files within the snapshot. The mountpoint used will be in the hidden directory .zfs/snapshot/name within the same ZFS dataset. For example, the snapshot mypool/dataset/snap1@snap2 is mounted at /mnt/mypool/dataset/.zfs/snapshot/ snap2/. If the length of this path exceeds 88 bytes the snapshot will not be automatically mounted by ZFS and the snapshot contents will not be visible or searchable. This can be resolved by renaming the ZFS pool or dataset containing the snapshot to shorter names (mypool or dataset), or by shortening the second part of the snapshot name (snap2), so that the total mounted path length does not exceed 88 bytes. ZFS will automatically perform any necessary unmount or remount of the file system as part of the rename operation. After renaming, the snapshot data will be visible and searchable again.
- Manual mount: If the same example snapshot is mounted manually from the CLI, using mount -t zfs mypool/ dataset/snap1@snap2 /mnt/mymountpoint the path /mnt/mountpoint/ must not exceed 88 bytes, but the length of the snapshot name will be *irrelevant*. When renaming a manual mountpoint, any object mounted on the mountpoint must be manually unmounted (using the umount command in the CLI) before renaming the mountpoint and can be remounted afterwards.

Note: A snapshot that cannot be mounted automatically by ZFS, can still be mounted manually from the CLI using a shorter mountpoint path. This makes it possible to mount and access snapshots that cannot be accessed automatically in other ways, such as from the GUI or from features such as "File History" or "Versions".

1.3 Hardware Recommendations

FreeNAS[®] 11.2 is based on FreeBSD 11.2 and supports the same hardware found in the FreeBSD Hardware Compatibility List (https://www.freebsd.org/releases/11.2R/hardware.html). Supported processors are listed in section 2.1 amd64 (https://www.freebsd.org/releases/11.2R/hardware.html#proc). FreeNAS[®] is only available for 64-bit processors. This architecture is called *amd64* by AMD and *Intel 64* by Intel.

Note: FreeNAS[®] boots from a GPT partition. This means that the system BIOS must be able to boot using either the legacy BIOS firmware interface or EFI.

Actual hardware requirements vary depending on the usage of the FreeNAS[®] system. This section provides some starter guidelines. The FreeNAS[®] Hardware Forum (https://forums.freenas.org/index.php?forums/hardware.18/) has performance tips from FreeNAS[®] users and is a place to post questions regarding the hardware best suited to meet specific requirements. Hardware Recommendations (https://forums.freenas.org/index.php?resources/hardware-recommendations-guide.12/) gives detailed recommendations for system components, with the FreeNAS[®] Quick Hardware Guide (https://forums.freenas.org/index.php?resources/freenas%C2%AE-quick-hardware-guide.7/) providing short lists of components for various configurations. Building, Burn-In, and Testing your FreeNAS[®] system (https://forums.freenas.org/index.php?threads/building-burn-in-and-testing-your-freenas-system.17750/) has detailed instructions on testing new hardware.

1.3.1 RAM

The best way to get the most out of a FreeNAS[®] system is to install as much RAM as possible. More RAM allows ZFS to provide better performance. The FreeNAS® Forums (https://forums.freenas.org/index.php) provide anecdotal evidence from users on how much performance can be gained by adding more RAM.

General guidelines for RAM:

• A minimum of 8 GiB of RAM is required.

Additional features require additional RAM, and large amounts of storage require more RAM for cache. An old, somewhat overstated guideline is 1 GiB of RAM per terabyte of disk capacity.

- To use Active Directory with many users, add an additional 2 GiB of RAM for the winbind internal cache.
- For iSCSI, install at least 16 GiB of RAM if performance is not critical, or at least 32 GiB of RAM if good performance is a requirement.
- *Jails* (page 256) are very memory-efficient, but can still use memory that would otherwise be available for ZFS. If the system will be running many jails, or a few resource-intensive jails, adding 1 to 4 additional gigabytes of RAM can be helpful. This memory is shared by the host and will be used for ZFS when not being used by jails.
- *Virtual Machines* (page 266) require additional RAM beyond any amounts listed here. Memory used by virtual machines is not available to the host while the VM is running, and is not included in the amounts described above. For example, a system that will be running two VMs that each need 1 GiB of RAM requires an additional 2 GiB of RAM.
- When installing FreeNAS[®] on a headless system, disable the shared memory settings for the video card in the BIOS.
- For ZFS deduplication, ensure the system has at least 5 GiB of RAM per terabyte of storage to be deduplicated.

If the hardware supports it, install ECC RAM. While more expensive, ECC RAM is highly recommended as it prevents in-flight corruption of data before the error-correcting properties of ZFS come into play, thus providing consistency for the checksumming and parity calculations performed by ZFS. If your data is important, use ECC RAM. This Case Study (http://research.cs.wisc.edu/adsl/Publications/zfs-corruption-fast10.pdf) describes the risks associated with memory corruption.

Do not use FreeNAS[®] to store data without at least 8 GiB of RAM. Many users expect FreeNAS[®] to function with less memory, just at reduced performance. The bottom line is that these minimums are based on feedback from many users. Requests for help in the forums or IRC are sometimes ignored when the installed system does not have at least 8 GiB of RAM because of the abundance of information that FreeNAS[®] may not behave properly with less memory.

1.3.2 The Operating System Device

The FreeNAS[®] operating system is installed to at least one device that is separate from the storage disks. The device can be a SSD, USB memory stick, or DOM (Disk on Module). Installation to a hard drive is discouraged as that drive is then not

available for data storage.

Note: To write the installation file to a USB stick, **two** USB ports are needed, each with an inserted USB device. One USB stick contains the installer, while the other USB stick is the destination for the FreeNAS[®] installation. Be careful to select the correct USB device for the FreeNAS[®] installation. FreeNAS[®] cannot be installed onto the same device that contains the installer. After installation, remove the installer USB stick. It might also be necessary to adjust the BIOS configuration to boot from the new FreeNAS[®] boot device.

When determining the type and size of the target device where FreeNAS® is to be installed, keep these points in mind:

- The absolute *bare minimum* size is 8 GiB. That does not provide much room. The *recommended* minimum is 16 GiB. This provides room for the operating system and several boot environments created by updates. More space provides room for more boot environments and 32 GiB or more is preferred.
- SSDs (Solid State Disks) are fast and reliable, and make very good FreeNAS[®] operating system devices. Their one disadvantage is that they require a disk connection which might be needed for storage disks.

Even a relatively large SSD (120 or 128 GiB) is useful as a boot device. While it might appear that the unused space is wasted, that space is instead used internally by the SSD for wear leveling. This makes the SSD last longer and provides greater reliability.

- When planning to add your own boot environments, budget about 1 GiB of storage per boot environment. Consider deleting older boot environments after making sure they are no longer needed. Boot environments can be created and deleted using *System* → *Boot*.
- Use quality, name-brand USB sticks, as ZFS will quickly reveal errors on cheap, poorly-made sticks.
- For a more reliable boot disk, use two identical devices and select them both during the installation. This will create a mirrored boot device.

Note: Current versions of FreeNAS[®] run directly from the operating system device. Early versions of FreeNAS[®] ran from RAM, but that has not been the case for years.

1.3.3 Storage Disks and Controllers

The Disk section (https://www.freebsd.org/releases/11.2R/hardware.html#disk) of the FreeBSD Hardware List lists the supported disk controllers. In addition, support for 3ware 6 Gbps RAID controllers has been added along with the CLI utility tw_cli for managing 3ware RAID controllers.

FreeNAS[®] supports hot pluggable drives. Using this feature requires enabling AHCI in the BIOS.

Reliable disk alerting and immediate reporting of a failed drive can be obtained by using an HBA such as an Broadcom MegaRAID controller or a 3Ware twa-compatible controller.

Note: Upgrading the firmware of Broadcom SAS HBAs to the latest version is recommended.

Some Highpoint RAID controllers do not support pass-through of S.M.A.R.T. data or other disk information, potentially including disk serial numbers. It is best to use a different disk controller with FreeNAS[®].

Note: The system is configured to prefer the mrsas(4) (https://www.freebsd.org/cgi/man.cgi?query=mrsas) driver for controller cards like the Dell PERC H330 and H730 which are supported by several drivers. Although not recommended, the mfi(4) (https://www.freebsd.org/cgi/man.cgi?query=mfi) driver can be used instead by removing the loader *Tunable* (page 77): hw.mfi.mrsas_enable or setting the *Value* to 0.

Suggestions for testing disks before adding them to a RAID array can be found in this forum post (https://forums.freenas.org/index.php?threads/checking-new-hdds-in-raid.12082/#post-55936). Additionally, badblocks (https://linux.die.net/man/8/badblocks) is installed with FreeNAS[®] for testing disks.

If the budget allows optimization of the disk subsystem, consider the read/write needs and RAID requirements:

• For steady, non-contiguous writes, use disks with low seek times. Examples are 10K or 15K SAS drives which cost about \$1/GiB. An example configuration would be six 600 GiB 15K SAS drives in a RAID 10 which would yield 1.8 TiB of usable space, or eight 600 GiB 15K SAS drives in a RAID 10 which would yield 2.4 TiB of usable space.

For ZFS, Disk Space Requirements for ZFS Storage Pools (https://docs.oracle.com/cd/E19253-01/819-5461/6n7ht6r12/index.html) recommends a minimum of 16 GiB of disk space. FreeNAS[®] allocates 2 GiB of swap space on each drive. Combined with ZFS space requirements, this means that **it is not possible to format drives smaller than 3 GiB**. Drives larger than 3 GiB but smaller than the minimum recommended capacity might be usable but lose a significant portion of storage space to swap allocation. For example, a 4 GiB drive only has 2 GiB of available space after swap allocation.

New ZFS user who are purchasing hardware should read through ZFS Storage Pools Recommendations (https://web.archive.org/web/20161028084224/http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide#ZFS_Sto first.

ZFS *vdevs*, groups of disks that act like a single device, can be created using disks of different sizes. However, the capacity available on each disk is limited to the same capacity as the smallest disk in the group. For example, a vdev with one 2 TiB and two 4 TiB disks will only be able to use 2 TiB of space on each disk. In general, use disks that are the same size for the best space usage and performance.

The ZFS Drive Size and Cost Comparison spreadsheet (https://forums.freenas.org/index.php?threads/zfs-drive-size-and-cost-comparison-spreadsheet.38092/) is available to compare usable space provided by different quantities and sizes of disks.

1.3.4 Network Interfaces

The Ethernet section (https://www.freebsd.org/releases/11.2R/hardware.html#ethernet) of the FreeBSD Hardware Notes indicates which interfaces are supported by each driver. While many interfaces are supported, FreeNAS[®] users have seen the best performance from Intel and Chelsio interfaces, so consider these brands when purchasing a new NIC. Realtek cards often perform poorly under CPU load as interfaces with these chipsets do not provide their own processors.

At a minimum, a GigE interface is recommended. While GigE interfaces and switches are affordable for home use, modern disks can easily saturate their 110 MiB/s throughput. For higher network throughput, multiple GigE cards can be bonded together using the LACP type of *Link Aggregations* (page 119). The Ethernet switch must support LACP, which means a more expensive managed switch is required.

When network performance is a requirement and there is some money to spend, use 10 GigE interfaces and a managed switch. Managed switches with support for LACP and jumbo frames are preferred, as both can be used to increase network throughput. Refer to the 10 Gig Networking Primer (https://forums.freenas.org/index.php?threads/10-gig-networking-primer.25749/) for more information.

Note: At present, these are not supported: InfiniBand, FibreChannel over Ethernet, or wireless interfaces.

Both hardware and the type of shares can affect network performance. On the same hardware, SMB is slower than FTP or NFS because Samba is single-threaded (https://www.samba.org/samba/docs/old/Samba3-Developers-Guide/architecture.html). So a fast CPU can help with SMB performance.

Wake on LAN (WOL) support depends on the FreeBSD driver for the interface. If the driver supports WOL, it can be enabled using ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig). To determine if WOL is supported on a particular interface, use the interface name with the following command. In this example, the capabilities line indicates that WOL is supported for the *re0* interface:

ifconfig -m re0

If WOL support is shown but not working for a particular interface, create a bug report using the instructions in *Support* (page 92).

1.4 Getting Started with ZFS

Readers new to ZFS should take a moment to read the ZFS Primer (page 320).

INSTALLING AND UPGRADING

The FreeNAS[®] operating system must be installed on a separate device from the drives which hold the storage data. In other words, with only one disk drive, the FreeNAS[®] graphical interface is available, but there is no place to store any data. And storing data is, after all, the whole point of a NAS system. Home users experimenting with FreeNAS[®] can install FreeNAS[®] on an inexpensive USB thumb drive and use the computer disks for storage.

This section describes:

- Getting FreeNAS® (page 18)
- Preparing the Media (page 18)
- Performing the Installation (page 20)
- Installation Troubleshooting (page 28)
- Upgrading (page 29)
- Virtualization (page 36)

2.1 Getting FreeNAS®

The latest STABLE version of FreeNAS[®] 11.2 can be downloaded from https://download.freenas.org/latest/.

Note: FreeNAS[®] requires 64-bit hardware.

The download page contains an *.iso* file. This is a bootable installer that can be written to either a CD or USB flash as described in *Preparing the Media* (page 18).

The .iso file has an associated sha256.txt file which is used to verify the integrity of the downloaded file. The command to verify the checksum varies by operating system:

- on a BSD system use the command sha256 name_of_file
- on a Linux system use the command sha256sum name_of_file
- on a Mac system use the command shasum -a 256 name_of_file
- Windows or Mac users can install additional utilities like HashCalc (http://www.slavasoft.com/hashcalc/) or HashTab (http://implbits.com/products/hashtab/)

The value produced by running the command must match the value shown in the sha256.txt file. Checksum values that do not match indicate a corrupted installer file that should not be used.

2.2 Preparing the Media

The FreeNAS[®] installer can run from either a CD or a USB memory stick.

A CD burning utility is needed to write the .iso file to a CD.

The .iso file can also be written to a USB memory stick. The method used to write the file depends on the operating system. Examples for several common operating systems are shown below.

Note: To install from a USB stick to another USB stick, **two** USB ports are needed, each with an inserted USB device. One USB stick contains the installer. The other USB stick is the destination for the FreeNAS[®] installation. Take care to select the correct USB device for the FreeNAS[®] installation. It is **not** possible to install FreeNAS[®] onto the same USB stick containing the installer. After installation, remove the installer USB stick. It might also be necessary to adjust the BIOS configuration to boot from the new FreeNAS[®] USB stick.

Ensure the boot device order in the BIOS is set to boot from the device containing the FreeNAS[®] installer media, then boot the system to start the installation.

2.2.1 On FreeBSD or Linux

On a FreeBSD or Linux system, the dd command is used to write the .iso file to an inserted USB thumb drive.

Warning: The dd command is very powerful and can destroy any existing data on the specified device. Make **absolutely sure** of the device name to write to and do not mistype the device name when using dd! The use of this command can be avoided by writing the .iso file to a CD instead.

This example demonstrates writing the image to the first USB device connected to a FreeBSD system. This first device usually reports as /dev/da0. Replace FreeNAS-RELEASE.iso with the filename of the downloaded FreeNAS[®] ISO file. Replace /dev/da0 with the device name of the device to write.

```
dd if=FreeNAS-RELEASE.iso of=/dev/da0 bs=64k
6117+0 records in
6117+0 records out
400883712 bytes transferred in 88.706398 secs (4519220 bytes/sec)
```

When using the dd command:

- if= refers to the input file, or the name of the file to write to the device.
- of= refers to the output file; in this case, the device name of the flash card or removable USB drive. Note that USB device numbers are dynamic, and the target device might be *da1* or *da2* or another name depending on which devices are attached. Before attaching the target USB drive, use ls /dev/da*. Then attach the target USB drive, wait ten seconds, and run ls /dev/da* again to see the new device name and number of the target USB drive. On Linux, use /dev/sdX, where X refers to the letter of the USB device.
- **bs=** refers to the block size, the amount of data to write at a time. The larger 64K block size shown here helps speed up writes to the USB drive.

2.2.2 On Windows

Microsoft provides the USB/DVD Download Tool to create a USB bootable image from an .iso file. Follow these instructions (https://www.microsoft.com/en-us/download/windows-usb-dvd-download-tool), but enter the name of the downloaded . iso into the SOURCE FILE box.

Image Writer (https://launchpad.net/win32-image-writer/) and Rufus (http://rufus.akeo.ie/) are alternate programs for writing images to USB sticks on a computer running Windows. When using Rufus, check *Create a bootable disk using* and select *DD Image* from the drop-down menu.

2.2.3 macOS

Insert the USB thumb drive. In the Finder, go to *Applications* \rightarrow *Utilities* \rightarrow *Disk Utility*. Unmount any mounted partitions on the USB thumb drive. Check that the USB thumb drive has only one partition, or partition table errors will be shown on boot.

If needed, use Disk Utility to set up one partition on the USB drive. Selecting *Free space* when creating the partition works fine.

Determine the device name of the inserted USB thumb drive. From TERMINAL, navigate to the Desktop, then type this command:

disku /dev/	til list disk0		
#:	TYPE NAME	SIZE	IDENTIFIER
0:	GUID_partition_scheme	*500.1 GB	disk0
1:	EFI	209.7 MB	disk0s1
2:	Apple_HFS Macintosh HD	499.2 GB	disk0s2
3:	Apple_Boot Recovery HD	650.0 MB	disk0s3
/dev/	disk1		
#:	TYPE NAME	SIZE	IDENTIFIER
0:	FDisk_partition_scheme	*8.0 GB	disk1
1:	DOS_FAT_32 UNTITLED	8.0 GB	disk1s1

This shows which devices are available to the system. Locate the target USB stick and record the path. To determine the correct path for the USB stick, remove the device, run the command again, and compare the difference. Once sure of the device name, navigate to the Desktop from TERMINAL, unmount the USB stick, and use the dd command to write the image to the USB stick. In this example, the USB thumb drive is /dev/disk1. It is first unmounted. The dd command is used to write the image to the faster "raw" version of the device (note the extra r in /dev/rdisk1).

When running these commands, replace FreeNAS-RELEASE.iso with the name of the FreeNAS[®] ISO. Replace /dev/rdisk1 with the correct path to the USB thumb drive:

```
diskutil unmountDisk /dev/disk1
Unmount of all volumes on disk1 was successful
dd if=FreeNAS-RELEASE.iso of=/dev/rdisk1 bs=64k
```

Note: If the error "Resource busy" is shown when the dd command is run, go to Applications \rightarrow Utilities \rightarrow Disk Utility, find the USB thumb drive, and click on its partitions to make sure all of them are unmounted. If a "Permission denied" is shown, use sudo for elevated rights: sudo dd if=FreeNAS-RELEASE.iso of=/dev/rdisk1 bs=64k. This will prompt for the password.

The dd command can take some minutes to complete. Wait until the prompt returns and a message is displayed with information about how long it took to write the image to the USB drive.

2.3 Performing the Installation

With the installation media inserted, boot the system from that media.

The FreeNAS[®] installer boot menu is displayed as is shown in Figure 2.1.

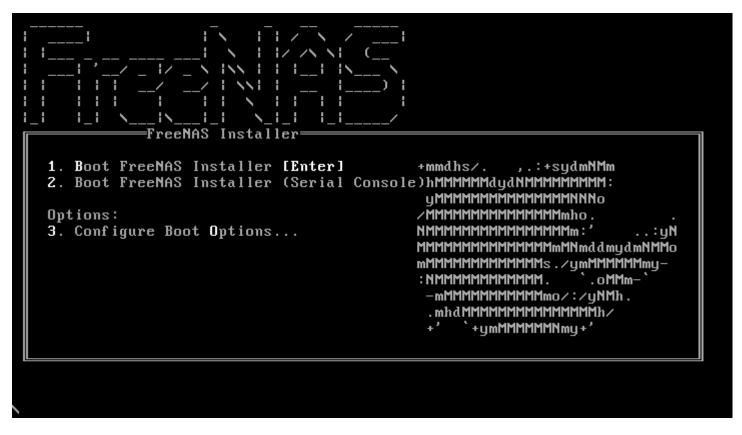


Fig. 2.1: Installer Boot Menu

The FreeNAS[®] installer automatically boots into the default option after ten seconds. If needed, choose another boot option by pressing the Spacebar to stop the timer and then enter the number of the desired option.

Tip: The *Serial Console* option is useful on systems which do not have a keyboard or monitor, but are accessed through a serial port, *Serial over LAN*, or *IPMI* (page 117).

Note: If the installer does not boot, verify that the installation device is listed first in the boot order in the BIOS. When booting from a CD, some motherboards may require connecting the CD device to SATA0 (the first connector) to boot from CD. If the installer stalls during bootup, double-check the SHA256 hash of the .iso file. If the hash does not match, re-download the file. If the hash is correct, burn the CD again at a lower speed or write the file to a different USB stick.

Once the installer has finished booting, the installer menu is displayed as shown in Figure 2.2.

FreeNAS 11.2-RELEASE Console Setup
Install/Upgrade 2 Shell 3 Reboot System 4 Shutdown System
<pre> Cancel> </pre>

Fig. 2.2: Installer Menu

Press Enter to select the default option, *1 Install/Upgrade*. The next menu, shown in Figure 2.3, lists all available drives. This includes any inserted USB thumb drives, which have names beginning with *da*.

Note: A minimum of 8 GiB of RAM is required and the installer will present a warning message if less than 8 GiB is detected.

In this example, the user is performing a test installation using VirtualBox and has created a 16 GiB virtual disk to hold the operating system.

for in: +		elect a drive with the spacebar).	
i.	[] ada0	VBOX HARDDISK 16.0 Gib VBOX HARDDISK 20.0 Gib	
	[] ada2	VBOX HARDDISK 8.0 GIB	
-	[] ada3	VBOX HARDDISK 9.0 GiB	
1	[] ada4	VBOX HARDDISK 11.0 GiB	
i	[] ada5	VBOX HARDDISK 14.0 GIB	
1	[] ada6	VBOX HARDDISK 8.0 GiB	
i		VBOX HARDDISK 10.0 Gib	
1	[] ada8	VBOX HARDDISK 8.0 GiB	11
ł	[] <mark>ada</mark> 9	VBOX HARDDISK 11.0 GiB	
•		<u>90</u> %	+
	< <u>0</u> K	> <cancel></cancel>	++

Fig. 2.3: Selecting the Install Drive

Use the arrow keys to highlight the destination USB drive, SSD, DOM (Disk on Module), or virtual disk. Press the spacebar to select it. To mirror the boot device, move to the second device and press spacebar to select it also. After making these selections, press Enter. The warning shown in Figure 2.4 is displayed, a reminder not to install the operating system on a drive that is meant for storage. Press Enter to continue on to the screen shown in Figure 2.6.

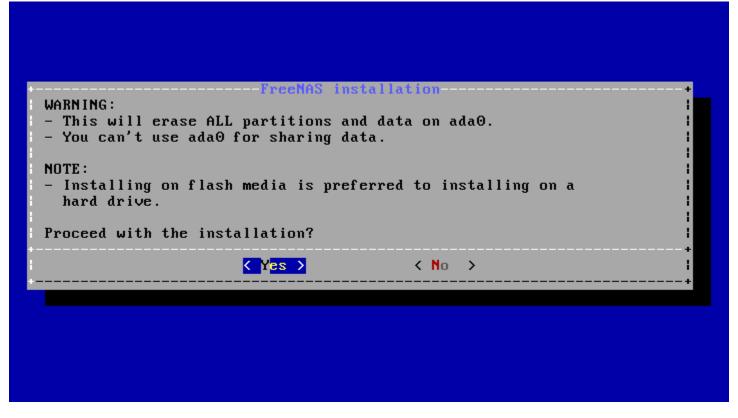


Fig. 2.4: Installation Warning

Note: A minimum of 8 GiB of space on the boot device is required. However, 32 GiB is recommended to provide room for future additions and boot environments. When using mirrored boot devices, it is best to use devices of the same size. If the device sizes are different, the mirror is limited to the size of the smaller device.

The installer recognizes existing installations of previous versions of FreeNAS[®]. When an existing installation is present, the menu shown in Figure 2.5 is displayed. To overwrite an existing installation, use the arrows to move to *Fresh Install* and press Enter twice to continue to the screen shown in Figure 2.6.

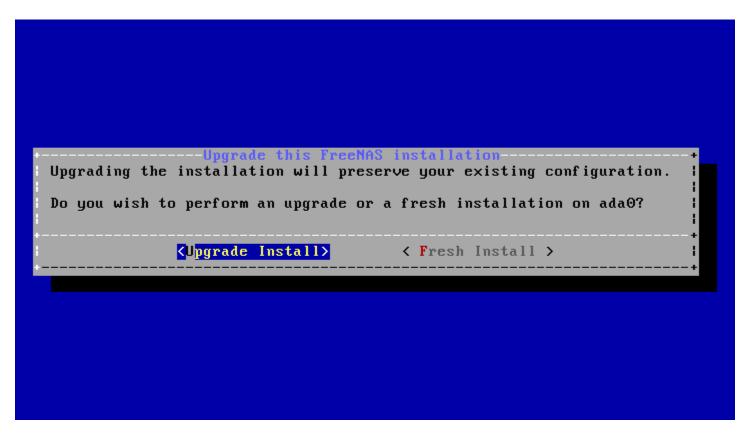


Fig. 2.5: Performing a Fresh Install

The screen shown in Figure 2.6 prompts for the *root* password which is used to log in to the administrative graphical interface.



Fig. 2.6: Set the Root Password

Setting a password is mandatory and the password cannot be blank. Since this password provides access to the administrative GUI, it should be hard to guess. Enter the password, press the down arrow key, and confirm the password. Then press Enter to continue with the installation. Choosing *Cancel* skips setting a root password during the installation, but the administrative GUI will require setting a root password when logging in for the first time.

Note: For security reasons, the SSH service and *root* SSH logins are disabled by default. Unless these are set, the only way to access a shell as *root* is to gain physical access to the console menu or to access the web shell within the administrative GUI. This means that the FreeNAS[®] system should be kept physically secure and that the administrative GUI should be behind a properly configured firewall and protected by a secure password.

FreeNAS[®] can be configured to boot with the standard BIOS boot mechanism or UEFI booting as shown Figure 2.7. BIOS booting is recommended for legacy and enterprise hardware. UEFI is used on newer consumer motherboards.

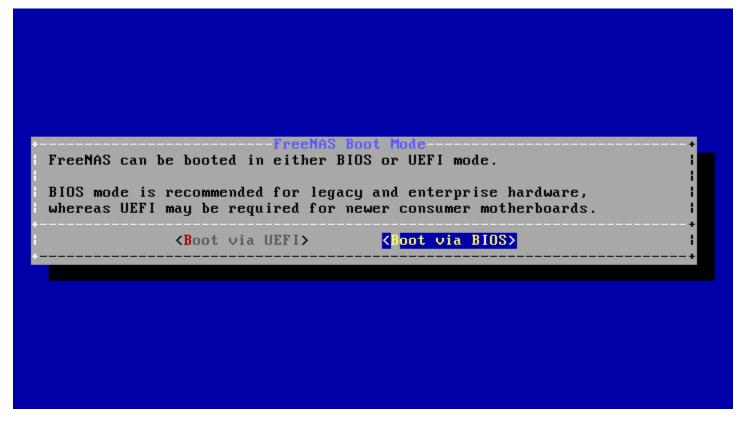


Fig. 2.7: Choose UEFI or BIOS Booting

Note: Most UEFI systems can also boot in BIOS mode if CSM (Compatibility Support Module) is enabled in the UEFI setup screens.

The message in Figure 2.8 is shown after the installation is complete.

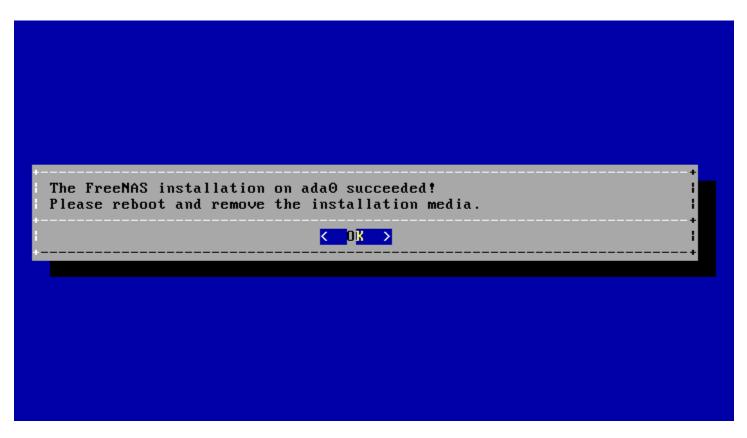


Fig. 2.8: Installation Complete

Press Enter to return to *Installer Menu* (page 22). Highlight *3 Reboot System* and press Enter. If booting from CD, remove the CDROM. As the system reboots, make sure that the device where FreeNAS[®] was installed is listed as the first boot entry in the BIOS so the system will boot from it.

FreeNAS[®] boots into the *Console Setup* menu described in *Booting* (page 54) after waiting five seconds in the *boot menu* (page 34). Press the *Spacebar* to stop the timer and use the boot menu.

2.4 Installation Troubleshooting

If the system does not boot into FreeNAS[®], there are several things that can be checked to resolve the situation.

Check the system BIOS and see if there is an option to change the USB emulation from CD/DVD/floppy to hard drive. If it still will not boot, check to see if the card/drive is UDMA compliant.

If the system BIOS does not support EFI with BIOS emulation, see if it has an option to boot using legacy BIOS mode.

When the system starts to boot but hangs with this repeated error message:

run_interrupt_driven_hooks: still waiting after 60 seconds for xpt_config

Go into the system BIOS and look for an onboard device configuration for a 1394 Controller. If present, disable that device and try booting again.

If the system starts to boot but hangs at a *mountroot*> prompt, follow the instructions in Workaround/Semi-Fix for Mountroot Issues with 9.3 (https://forums.freenas.org/index.php?threads/workaround-semi-fix-for-mountroot-issues-with-9-3.26071/).

If the burned image fails to boot and the image was burned using a Windows system, wipe the USB stick before trying a second burn using a utility such as Active@ KillDisk (http://how-to-erase-hard-drive.com/). Otherwise, the second burn attempt will fail as Windows does not understand the partition which was written from the image file. Be very careful to specify the correct USB stick when using a wipe utility!

2.5 Upgrading

FreeNAS[®] provides flexibility for keeping the operating system up-to-date:

- 1. Upgrades to major releases, for example from version 9.3 to 9.10, can still be performed using either an ISO or the graphical administrative interface. Unless the Release Notes for the new major release indicate that the current version requires an ISO upgrade, either upgrade method can be used.
- 2. Minor releases have been replaced with signed updates. This means that it is not necessary to wait for a minor release to update the system with a system update or newer versions of drivers and features. It is also no longer necessary to manually download an upgrade file and its associated checksum to update the system.
- 3. The updater automatically creates a boot environment, making updates a low-risk operation. Boot environments provide the option to return to the previous version of the operating system by rebooting the system and selecting the previous boot environment from the boot menu.

This section describes how to perform an upgrade from an earlier version of FreeNAS[®] to 11.2. After 11.2 has been installed, use the instructions in *Update* (page 79) to keep the system updated.

2.5.1 Caveats

Be aware of these caveats **before** attempting an upgrade to 11.2:

- Warning: upgrading the ZFS pool can make it impossible to go back to a previous version. For this reason, the update process does not automatically upgrade the ZFS pool, though the *Alert* (page 296) system shows when newer feature flags are available for a pool. Unless a new feature flag is needed, it is safe to leave the pool at the current version and uncheck the alert. If the pool is upgraded, it will not be possible to boot into a previous version that does not support the newer feature flags.
- The *Wizard* (page 281) does not recognize an encrypted ZFS pool. If the ZFS pool is GELI-encrypted and the *Wizard* (page 281) starts after the upgrade, cancel the *Wizard* (page 281) and use the instructions in *Importing an Encrypted Pool* (page 138) to import the encrypted volume. The *Wizard* (page 281) can be run afterward for post-configuration. It will then recognize that the volume has been imported and not prompt to reformat the disks.
- Upgrading the firmware of Broadcom SAS HBAs to the latest version is recommended.
- If upgrading from 9.3.x, please read the FAQ: Updating from 9.3 to 9.10 (https://forums.freenas.org/index.php?threads/faq-updating-from-9-3-to-9-10.54260/) first.
- Upgrades from FreeNAS[®] 0.7x are not supported. The system has no way to import configuration settings from 0.7x versions of FreeNAS[®]. The configuration must be manually recreated. If supported, the FreeNAS[®] 0.7x volumes or disks must be manually imported.
- **Upgrades on 32-bit hardware are not supported.** However, if the system is currently running a 32-bit version of FreeNAS[®] **and** the hardware supports 64-bit, the system can be upgraded. Any archived reporting graphs will be lost during the upgrade.
- **UFS is not supported.** If the data currently resides on **one** UFS-formatted disk, create a ZFS volume using **other** disks after the upgrade, then use the instructions in *Import Disk* (page 137) to mount the UFS-formatted disk and copy the data to the ZFS volume. With only one disk, back up its data to another system or media before the upgrade, format the disk as ZFS after the upgrade, then restore the backup. If the data currently resides on a UFS RAID of disks, it is not possible to directly import that data to the ZFS volume. Instead, back up the data before the upgrade, create a ZFS volume after the upgrade, then restore the data from the backup.
- The VMware Tools VMXNET3 drivers are no longer supported. Configure and use the vmx(4) (https://www.freebsd.org/cgi/man.cgi?query=vmx) driver instead.

2.5.2 Initial Preparation

Before upgrading the operating system, perform the following steps:

1. Back up the FreeNAS[®] configuration in System \rightarrow General \rightarrow Save Config.

- 2. If any volumes are encrypted, **remember** to set the passphrase and download a copy of the encryption key and the latest recovery key. After the upgrade is complete, use the instructions in *Importing an Encrypted Pool* (page 138) to import the encrypted volume.
- 3. Warn users that the FreeNAS[®] shares will be unavailable during the upgrade; scheduling the upgrade for a time that will least impact users is recommended.
- 4. Stop all services in Services \rightarrow Control Services.

2.5.3 Upgrading Using the ISO

To perform an upgrade using this method, download (http://download.freenas.org/latest/) the .iso to the computer that will be used to prepare the installation media. Burn the downloaded .iso file to a CD or USB thumb drive using the instructions in *Preparing the Media* (page 18).

Insert the prepared media into the system and boot from it. The installer waits ten seconds in the *installer boot menu* (page 21) before booting the default option. If needed, press the Spacebar to stop the timer and choose another boot option. After the media finishes booting into the installation menu, press Enter to select the default option of *1 Install/Upgrade*. The installer presents a screen showing all available drives.

Warning: All drives are shown, including boot drives and storage drives. Only choose boot drives when upgrading. Choosing the wrong drives to upgrade or install will cause loss of data. If unsure about which drives contain the FreeNAS[®] operating system, reboot and remove the install media. In the FreeNAS[®] GUI, use *System* \rightarrow *Boot* to identify the boot drives. More than one drive is shown when a mirror has been used.

Move to the drive where FreeNAS[®] is installed and press the Spacebar to mark it with a star. If a mirror has been used for the operating system, mark all of the drives where the FreeNAS[®] operating system is installed. Press Enter when done.

The installer recognizes earlier versions of FreeNAS[®] installed on the boot drive or drives and presents the message shown in Figure 2.9.



Fig. 2.9: Upgrading a FreeNAS[®] Installation

Note: If *Fresh Install* is chosen, the backup of the configuration data must be restored using *System* \rightarrow *General* \rightarrow *Upload Config* after booting into the new operating system.

To perform an upgrade, press Enter to accept the default of *Upgrade Install*. The installer recommends installing the operating system on a disk not used for storage.

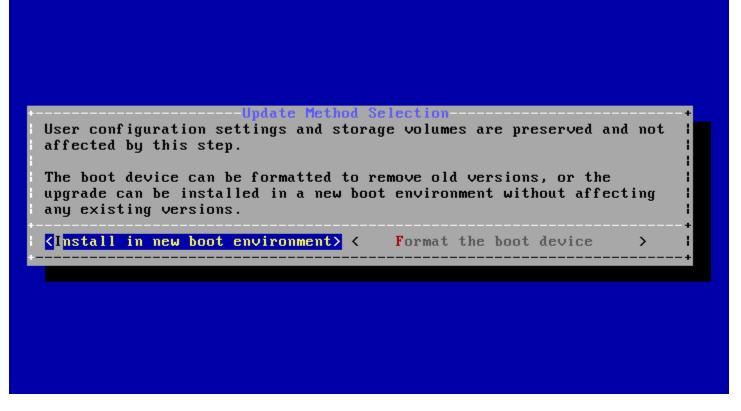


Fig. 2.10: Install in New Boot Environment or Format

The updated system can be installed in a new boot environment, or the entire boot device can be formatted to start fresh. Installing into a new boot environment preserves the old code, allowing a roll-back to previous versions if necessary. Formatting the boot device is usually not necessary but can reclaim space. User data and settings are preserved when installing to a new boot environment and also when formatting the boot device. Move the highlight to one of the options and press Enter to start the upgrade.

The installer unpacks the new image and displays the menu shown in Figure 2.11. The database file that is preserved and migrated contains the FreeNAS[®] configuration settings.

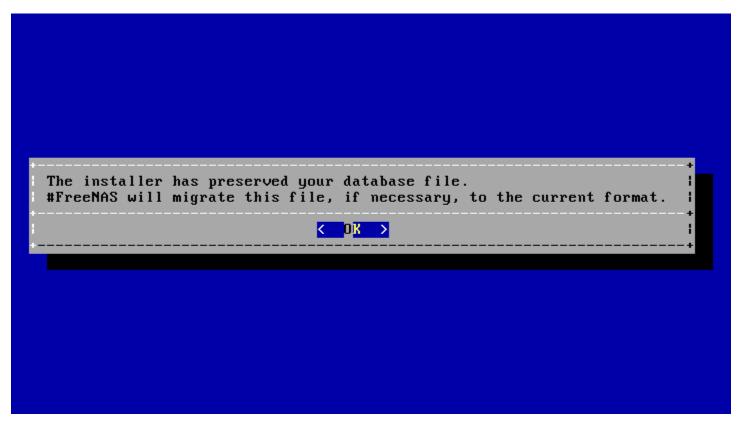


Fig. 2.11: Preserve and Migrate Settings

Press Enter. FreeNAS[®] indicates that the upgrade is complete and a reboot is required. Press *OK*, highlight *3 Reboot System*, then press Enter to reboot the system. If the upgrade installer was booted from CD, remove the CD.

During the reboot there can be a conversion of the previous configuration database to the new version of the database. This happens during the "Applying database schema changes" line in the reboot cycle. This conversion can take a long time to finish, sometimes fifteen minutes or more, and can cause the system to reboot again. The system will start normally afterwards. If database errors are shown but the graphical administrative interface is accessible, go to *Settings* \rightarrow *General* and use the *Upload Config* button to upload the configuration that was saved before starting the upgrade.

2.5.4 Upgrading From the GUI

To perform an upgrade using this method, go to *System* \rightarrow *Update*.

After the update is complete, the connection will be lost temporarily as the FreeNAS[®] system reboots into the new version of the operating system. The FreeNAS[®] system will normally receive the same IP address from the DHCP server. Refresh the browser after a moment to see if the system is accessible.

2.5.5 If Something Goes Wrong

If an update fails, an alert is issued and the details are written to /data/update.failed.

To return to a previous version of the operating system, physical or IPMI access to the FreeNAS[®] console is needed. Reboot the system and watch for the boot menu:



Fig. 2.12: Boot Menu

FreeNAS[®] waits five seconds before booting into the default boot environment. Press the Spacebar to stop the automatic boot timer. Press 4 to display the available boot environments and press 3 as needed to scroll through multiple pages.



Fig. 2.13: Boot Environments

In the example shown in Figure 2.13, the first entry in *Boot Environments* is 11.2-MASTER-201807250900. This is the current version of the operating system, after the update was applied. Since it is the first entry, it is the default selection.

The next entry is Initial-Install. This is the original boot environment created when FreeNAS[®] was first installed. Since there are no other entries between the initial installation and the first entry, only one update has been applied to this system since its initial installation.

To boot into another version of the operating system, enter the number of the boot environment to set it as *Active*. Press Backspace to return to the *Boot Menu* (page 34) and press Enter to boot into the chosen *Active* boot environment.

If a boot device fails and the system no longer boots, don't panic. The data is still on the disks and there is still a copy of the saved configuration. The system can be recovered with a few steps:

- 1. Perform a fresh installation on a new boot device.
- 2. Import the volumes in Storage \rightarrow Auto Import Volume.
- 3. Restore the configuration in System \rightarrow General \rightarrow Upload Config.

Note: It is not possible to restore a saved configuration that is newer than the installed version. For example, if a reboot into an older version of the operating system is performed, a configuration that was created in a later version cannot be restored.

2.5.6 Upgrading a ZFS Pool

In FreeNAS[®], ZFS pools can be upgraded from the graphical administrative interface.

Before upgrading an existing ZFS pool, be aware of these caveats first:

- the pool upgrade is a one-way street, meaning that if you change your mind you cannot go back to an earlier ZFS version or downgrade to an earlier version of the software that does not support those feature flags.
- before performing any operation that may affect the data on a storage disk, always back up all data first and verify the integrity of the backup. While it is unlikely that the pool upgrade will affect the data, it is always better to be safe than sorry.
- upgrading a ZFS pool is **optional**. Do not upgrade the pool if the the possibility of reverting to an earlier version of FreeNAS[®] or repurposing the disks in another operating system that supports ZFS is desired. It is not necessary to upgrade the pool unless the end user has a specific need for the newer ZFS feature flags. If a pool is upgraded to the latest feature flags, it will not be possible to import that pool into another operating system that does not yet support those feature flags.

To perform the ZFS pool upgrade, go to *Storage* \rightarrow *Volumes* \rightarrow *View Volumes* and highlight the volume (ZFS pool) to upgrade. Click the *Upgrade* button as shown in Figure 2.14.

Note: If the Upgrade button does not appear, the pool is already at the latest feature flags and does not need to be upgraded.

Volume Man	ager Import Disk	Import Vol	ume View Disks		
Name	Used	Available	Compression	Compression Ratio	Status
∡ volume1	3.1 MiB (0%)	23.9 GiB	-	-	HEALTHY
volume	1 1.3 MiB (0%)	15.4 GiB	lz4	2.26x	-
	e you sure you wa n is irreversible!	nt to upgrade	syour pool?		

Fig. 2.14: Upgrading a ZFS Pool

The warning serves as a reminder that a pool upgrade is not reversible. Click *OK* to proceed with the upgrade.

The upgrade itself only takes a few seconds and is non-disruptive. It is not necessary to stop any sharing services to upgrade the pool. However, it is best to upgrade when the pool is not being heavily used. The upgrade process will suspend I/O for a short period, but is nearly instantaneous on a quiet pool.

2.6 Virtualization

FreeNAS[®] can be run inside a virtual environment for development, experimentation, and educational purposes. Note that running FreeNAS[®] in production as a virtual machine is not recommended (https://forums.freenas.org/index.php?threads/please-do-not-run-freenas-in-production-as-a-virtualmachine.12484/). Before using FreeNAS within a virtual environment for the first time, read this post (https://forums.freenas.org/index.php?threads/absolutely-must-virtualize-freenas-a-guide-to-not-completely-losing-yourdata.12714/) as it contains useful guidelines for minimizing the risk of losing data.

To install or run FreeNAS[®] within a virtual environment, create a virtual machine that meets these minimum requirements:

- at least 8192 MB (8 GiB) base memory size
- a virtual disk at least 8 GiB in size to hold the operating system and boot environments
- at least one additional virtual disk at least 4 GiB in size to be used as data storage
- a bridged network adapter

This section demonstrates how to create and access a virtual machine within VirtualBox and VMware ESXi environments.

2.6.1 VirtualBox

VirtualBox <https://www.virtualbox.org/>'___ is an open source virtualization program originally created by Sun Microsystems. VirtualBox runs on Windows, BSD, Linux, Macintosh, and OpenSolaris. It can be configured to use a downloaded FreeNAS:sup:(® .iso file, and makes a good testing environment for practicing configurations or learning how to use the features provided by FreeNAS[®].

To create the virtual machine, start VirtualBox and click the *New* button, shown in Figure 2.15, to start the new virtual machine wizard.

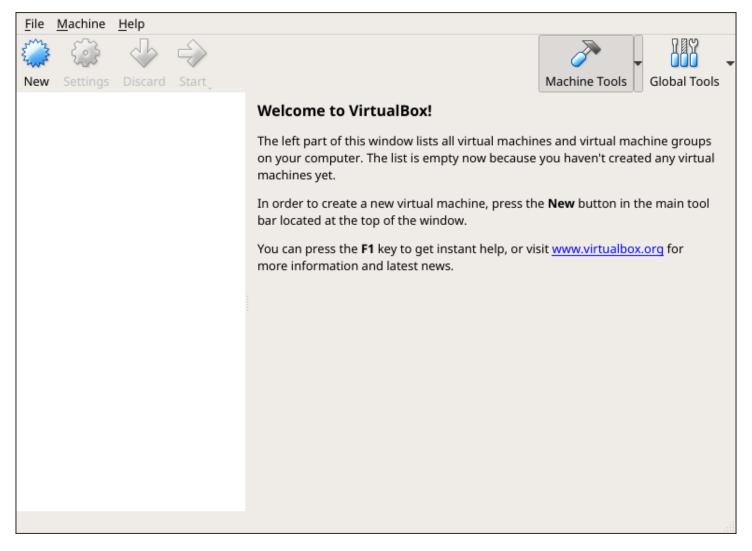


Fig. 2.15: Initial VirtualBox Screen

Click the *Next* button to see the screen in Figure 2.16. Enter a name for the virtual machine, click the *Operating System* drop-down menu and select BSD, and select *FreeBSD (64-bit)* from the *Version* dropdown.

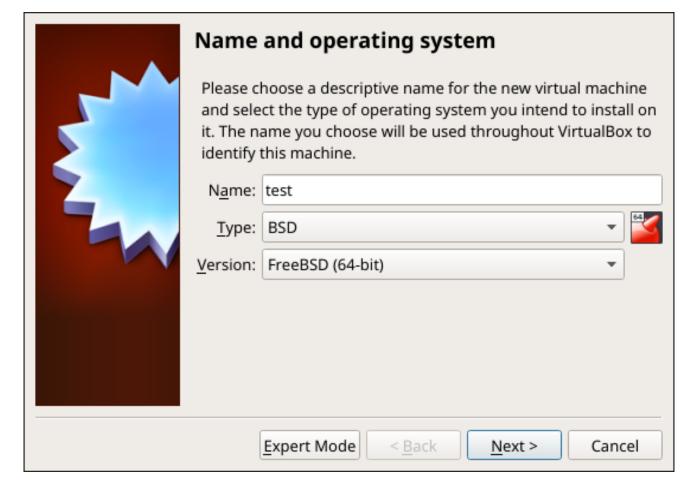


Fig. 2.16: Type in a Name and Select the Operating System for the New Virtual Machine

Click *Next* to see the screen in Figure 2.17. The base memory size must be changed to **at least 8192 MB**. When finished, click *Next* to see the screen in Figure 2.18.

Memory size				
Select the amount of memory (RAM) in megabytes to be allocated to the virtual machine.				
The recommended memory size is 1024 MB.				
	0		8192 🗘 MB	
4 MB		16384 MB		
	< <u>B</u> ack	<u>N</u> ext >	Cancel	

Fig. 2.17: Select the Amount of Memory Reserved for the Virtual Machine

	Hard disk		
	If you wish you can add a virtual hard disk to the new machine. You can either create a new hard disk file or select one from the list or from another location using the folder icon.		
	If you need a more complex storage set-up you can skip this step and make the changes to the machine settings once the machine is created.		
	The recommended size of the hard disk is 16.00 GB .		
	O <u>D</u> o not add a virtual hard disk		
	• <u>Create a virtual hard disk now</u>		
	O Use an existing virtual hard disk file		
	Empty 🗸 🖉		
	< <u>B</u> ack Create Cancel		

Fig. 2.18: Select Existing or Create a New Virtual Hard Drive

Click Create to launch the Create Virtual Hard Drive Wizard shown in Figure 2.19.

Hard disk file type		
Please choose the type of file that you would like to use for the new virtual hard disk. If you do not need to use it with other virtualization software you can leave this setting unchanged.		
• VDI (VirtualBox Disk Image)		
O VHD (Virtual Hard Disk)		
 VMDK (Virtual Machine Disk) 		
Expert Mode < Back Next > Cancel		

Fig. 2.19: Create New Virtual Hard Drive Wizard

Select *VDI* and click the *Next* button to see the screen in Figure 2.20.

	Storage on physical hard disk	
	Please choose whether the new virtual hard disk file should grow as it is used (dynamically allocated) or if it should be created at its maximum size (fixed size).	
	A dynamically allocated hard disk file will only use space on your physical hard disk as it fills up (up to a maximum fixed size), although it will not shrink again automatically when space on it is freed.	
	A fixed size hard disk file may take longer to create on some systems but is often faster to use.	
	• <u>Dynamically allocated</u>	
	○ <u>F</u> ixed size	
	< <u>B</u> ack <u>N</u> ext > Cancel	

Fig. 2.20: Select Storage Type for Virtual Disk

Choose either *Dynamically allocated* or *Fixed-size* storage. The first option uses disk space as needed until it reaches the maximum size that is set in the next screen. The second option creates a disk the full amount of disk space, whether it is used or not. Choose the first option to conserve disk space; otherwise, choose the second option, as it allows VirtualBox to run slightly faster. After selecting *Next*, the screen in Figure 2.21 is shown.

	File location and size	e	
	2 .	new virtual hard disk file into the box below o ect a different folder to create the file in.	or
	test		
		hard disk in megabytes. This size is the limit o a virtual machine will be able to store on the	on
			GB
	4.00 MB	2.00 TB	
		< <u>B</u> ack Create Cance	

Fig. 2.21: Select File Name and Size of Virtual Disk

This screen is used to set the size (or upper limit) of the virtual disk. **Set the default size to a minimum of 8 GiB**. Use the folder icon to browse to a directory on disk with sufficient space to hold the virtual disk files. Remember that there will be a system disk of at least 8 GiB and at least one data storage disk of at least 4 GiB.

Use the *Back* button to return to a previous screen if any values need to be modified. After making a selection and pressing *Create*, the new VM is created. The new virtual machine is listed in the left frame, as shown in the example in Figure 2.22. Open the *Machine Tools* drop-down menu and select *Details* to see extra information about the VM.

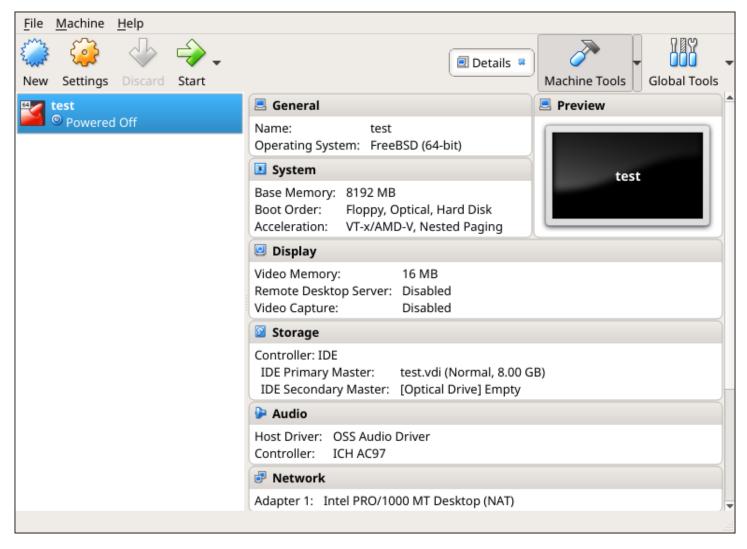


Fig. 2.22: The New Virtual Machine

Create the virtual disks to be used for storage. Highlight the VM and click *Settings* to open the menu. Click the *Storage* option in the left frame to access the storage screen seen in Figure 2.23.

	General	Storage	
1	System	Storage Devices Attributes	
	Display	Controller: IDE	
9	Storage		
	Audio		•
7	Network	Empty Use Host I/O Cac	he
	Serial Ports		
Ø	USB		
	Shared Folders		
•	User Interface		
		🍐 🗞 📴	
		G Add Optical Drive	Cancel
		🙆 Add Hard Disk	

Fig. 2.23: Storage Settings of the Virtual Machine

Click the *Add Attachment* button, select *Add Hard Disk* from the pop-up menu, then click the *Create new disk* button. This launches the *Create Virtual Hard Disk* Wizard seen in Figure 2.19 and 2.20.

This disk will be used for storage, so create a size appropriate to your needs, making sure that it is **at least 4 GiB**. To practice with RAID configurations, create as many virtual disks as needed. Two disks can be created on each IDE controller. For additional disks, click the *Add Controller* button to create another controller for attaching additional disks.

Create a device for the installation media. Highlight the word "Empty", then click the CD icon as shown in Figure 2.24.

	General	Storage				
		Storage Devices	Attributes			
		😂 Controller: IDE	Optical <u>D</u> rive:	IDE Secondary Ma: 👻	0	
9		🛛 🖸 test.vdi		Live CD/DVD	🛛 C	hoose Virtual Optical Disk File
	Audio	NewVirtualDisk1.vdi	Information		н	ost Drive ASUS DRW-24B1ST j (cd0)
	Network	Empty	Type:		F	reeNAS-11.2-INTERNAL24.iso
	Serial Ports	NewVirtualDisk2.vdi	Size:	-	🕲 R	emove Disk from Virtual Drive
	USB		Location:			
	Shared Folders		Attached to:			
	User Interface					
		la 🕹 🖾 🗮				
			(OK Cancel		

Fig. 2.24: Configuring ISO Installation Media

Click Choose Virtual Optical disk file... to browse to the location of the .iso file. If the .iso was burned to CD, select the detected Host Drive.

Depending on the extensions available in the host CPU, it might not be possible to boot the VM from an .iso. If "your CPU does not support long mode" is shown when trying to boot the .iso, the host CPU either does not have the required extension or AMD-V/VT-x is disabled in the system BIOS.

Note: If there is a kernel panic when booting into the ISO, stop the virtual machine. Then, go to *System* and check the box *Enable IO APIC*.

To configure the network adapter, go to *Settings* \rightarrow *Network* \rightarrow *Adapter* 1. In the *Attached to* drop-down menu select *Bridged Adapter*, then choose the name of the physical interface from the *Name* drop-down menu. In the example shown in Figure 2.25, the Intel Pro/1000 Ethernet card is attached to the network and has a device name of *em0*.

	General	Network
	System	Adapter 1 Adapter 2 Adapter 2 Adapter 4
	Display	Adapter <u>1</u> Adapter <u>2</u> Adapter <u>3</u> Adapter <u>4</u>
9	Storage	✓ Enable Network Adapter
	Audio	Attached to: Bridged Adapter 💌
₽	Network	Name: re0 •
٨	Serial Ports	▶ A <u>d</u> vanced
Ø	USB	_
	Shared Folders	
-	User Interface	
-		OK Cancel

Fig. 2.25: Configuring a Bridged Adapter in VirtualBox

After configuration is complete, click the *Start* arrow and install FreeNAS[®] as described in *Performing the Installation* (page 20). Once FreeNAS[®] is installed, press F12 when the VM starts to boot to access the boot menu. Select the primary hard disk as the boot option. To permanently boot from disk, remove the *Optical* device in *Storage* or uncheck *Optical* in the *Boot Order* section of *System*.

2.6.2 VMware ESXi

Before using ESXi, read this post (https://forums.freenas.org/index.php?threads/sync-writes-or-why-is-my-esxi-nfs-so-slowand-why-is-iscsi-faster.12506/) for an explanation of why iSCSI will be faster than NFS.

ESXi is a bare-metal hypervisor architecture created by VMware Inc. Commercial and free versions of the VMware vSphere Hypervisor operating system (ESXi) are available from the VMware website (https://www.vmware.com/products/esxi-andesx.html). After the operating system is installed on supported hardware, use a web browser to connect to its IP address. The welcome screen provides a link to download the VMware vSphere client which is used to create and manage virtual machines.

Once the VMware vSphere client is installed, use it to connect to the ESXi server. To create a new virtual machine, click *File* \rightarrow *New* \rightarrow *Virtual Machine*. The New Virtual Machine Wizard will launch as shown in Figure 2.26.

Create New Virtual Machine	
Configuration Select the configuration fo	virtual Machine Version: 8
Configuration Name and Location Storage Guest Operating System Network Create a Disk Ready to Complete	Configuration
Help	< Back Next > Cancel

Fig. 2.26: New Virtual Machine Wizard

Click *Next* and enter a name for the virtual machine. Click *Next* and highlight a datastore. An example is shown in Figure 2.27. Click *Next*. In the screen shown in Figure 2.28, click *Other*, then select a FreeBSD 64-bit architecture.

itorage Select a destination stora	age for the virtual machine	: files		Virtual Mach	nine Version:
Configuration	Select a destination st	orage for the virtua	al machine files:		
lame and Location itorage	VM Storage Profile:		-		
est Operating System	Name	Drive Type	Capacity Provisioned	Free Type	Thin Pro
iork te a Disk	datastore1	Non-SSD	227.25 GB 752.00 MB	226.52 GB VMFS3	Support
Complete	FN1-DS1	SSD	1023.75 GB 65.11 GB	1009.59 G VMFS5	Supporte
	FN1-NF51	Unknown	2.78 TB 96.00 KB 1.33 TB 121.32 GB	2.78 TB NFS 1.30 TB NFS	Supporte Supporte
	✓ Disable Storage [DRS for this virtual I	III machine		4
	Select a datastore:				
	Name	Drive Type	Capacity Provisioned	Free Type	Thin Prov
	Name	Drive Type	Capacity Provisioned	Free Type	Thin Prov
		Drive Type		Free Type	

Fig. 2.27: Select Datastore

🕝 Create New Virtual Machine		×	3
Guest Operating System Specify the guest operatin	g system to use with this virtual machine	Virtual Machine Versior	n: 8
Configuration Name and Location Storage Guest Operating System Network Create a Disk Ready to Complete	Guest Operating System:	propriate defaults for	
Help	< Back Ne	ext > Cancel	

Fig. 2.28: Select Operating System

Click *Next* and create a virtual disk file of **8 GiB** to hold the FreeNAS[®] operating system, as shown in Figure 2.29.

🕜 Create New Virtual Machine	1		
Create a Disk Specify the virtual disk size	and provisioning policy		Virtual Machine Version: 8
Configuration Name and Location Storage Guest Operating System Network Create a Disk Ready to Complete	Datastore: Available space (GB): Virtual disk size: Thick Provision Lazy Z Thick Provision Eager Thin Provision		
Help		< Back Next	t > Cancel

Fig. 2.29: Create Disk for the Operating System

Click *Next* and *Finish*. The new virtual machine is listed in the left frame. Right-click the virtual machine and select *Edit Settings* to access the screen shown in Figure 2.30.

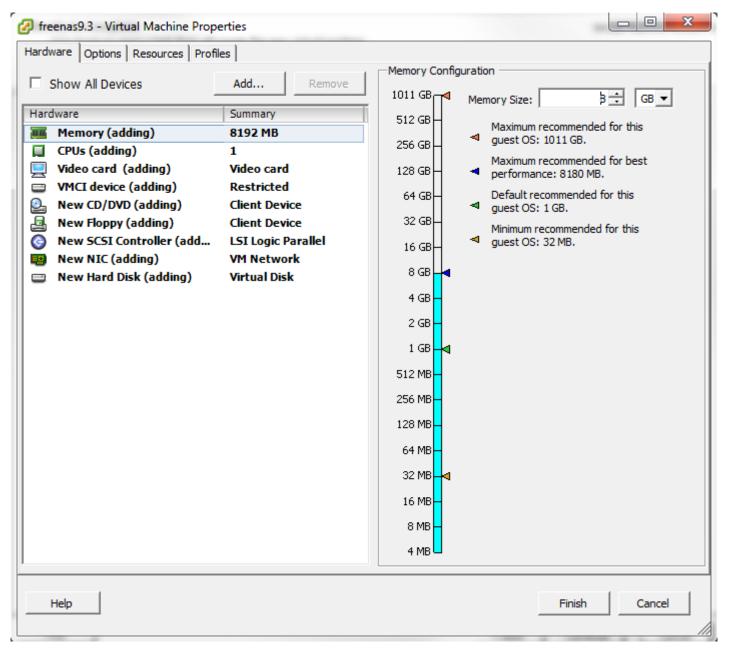


Fig. 2.30: Virtual Machine Settings

Increase the Memory Configuration to at least 8192 MB.

To create a storage disk, click *Hard disk* $1 \rightarrow Add$. In the *Device Type* menu, highlight *Hard Disk* and click *Next*. Select *Create a new virtual disk* and click *Next*. In the screen shown in Figure 2.31, select the size of the disk. To dynamically allocate space as needed, check the box *Allocate and commit space on demand (Thin Provisioning)*. Click *Next*, then *Next*, then *Finish* to create the disk. Repeat to create the amount of storage disks needed to meet your requirements.

Device Type Select a Disk Create a Disk Advanced Options Ready to Complete	Capacity Disk Size: 20 GB Disk Provisioning Thick Provision Lazy Zeroed Thick Provision Eager Zeroed Thin Provision Location Specify a datastore or datastore duster: Browse
--	--

Fig. 2.31: Creating a Storage Disk

For ESX 5.0, Workstation 8.0, or Fusion 4.0 or higher, additional configuration is needed so that the virtual HPET setting does not prevent the virtual machine from booting.

If ESX is running, while in *Edit Settings*, click *Options* \rightarrow *Advanced* \rightarrow *General* \rightarrow *Configuration Parameters*. Change *hpet0.present* from *true* to *false*, then click *OK* twice to save the setting.

For Workstation or Player, while in *Edit Settings*, click *Options* \rightarrow *Advanced* \rightarrow *File Locations*. Locate the path for the Configuration file named filename.vmx. Open that file in a text editor, change *hpet0.present* from *true* to *false*, and save the change.

CHAPTER THREE

BOOTING

The Console Setup menu, shown in Figure 3.1, appears at the end of the boot process. If the FreeNAS[®] system has a keyboard and monitor, this Console Setup menu can be used to administer the system.

Note: When connecting to the FreeNAS[®] system with SSH or the web *Shell* (page 289), the Console Setup menu is not shown by default. It can be started by the *root* user or another user with root permissions by typing /etc/netcli.

The Console Setup menu can be disabled by unchecking *Enable Console Menu* in *System* \rightarrow *Advanced*.

Console setup
1) Configure Noticesh Interforme
1) Configure Network Interfaces
2) Configure Link Aggregation
 Configure VLAN Interface
4) Configure Default Route
5) Configure Static Routes
6) Configure DNS
7) Reset Root Password
8) Reset Configuration to Defaults
9) Shell
10) Reboot
11) Shut Down
The web user interface is at:
http://10.0.0.102
Enter an option from 1-11:

Fig. 3.1: Console Setup Menu

The menu provides these options:

1) Configure Network Interfaces provides a configuration wizard to set up the system's network interfaces.

2) Configure Link Aggregation is for creating or deleting link aggregations.

3) Configure VLAN Interface is used to create or delete VLAN interfaces.

4) Configure Default Route is used to set the IPv4 or IPv6 default gateway. When prompted, enter the IP address of the default gateway.

5) Configure Static Routes prompts for the destination network and gateway IP address. Re-enter this option for each static route needed.

6) Configure DNS prompts for the name of the DNS domain and the IP address of the first DNS server. When adding multiple DNS servers, press Enter to enter the next one. Press Enter twice to leave this option.

7) Reset Root Password is used to reset a lost or forgotten root password. Select this option and follow the prompts to set the password.

8) Reset Configuration to Defaults **Caution**! This option deletes *all* of the configuration settings made in the administrative GUI and is used to reset a FreeNAS[®] system back to defaults. **Before selecting this option, make a full backup of all data and make sure all encryption keys and passphrases are known!** After this option is selected, the configuration is reset to defaults and the system reboots. *Storage* \rightarrow *Pools* \rightarrow *Import Pool* can be used to re-import pools.

9) Shell starts a shell for running FreeBSD commands. To leave the shell, type exit.

10) Reboot reboots the system.

11) Shut Down shuts down the system.

Note: The numbering and quantity of options on this menu can change due to software updates, service agreements, or other factors. Please carefully check the menu before selecting an option, and keep this in mind when writing local procedures.

3.1 Obtaining an IP Address

During boot, FreeNAS[®] automatically attempts to connect to a DHCP server from all live network interfaces. If it successfully receives an IP address, the address is displayed so it can be used to access the graphical user interface. The example in Figure 3.1 shows a FreeNAS[®] system that is accessible at *http://192.168.1.119*.

Some FreeNAS[®] systems are set up without a monitor, making it challenging to determine which IP address has been assigned. On networks that support Multicast DNS (mDNS), the hostname and domain can be entered into the address bar of a browser. By default, this value is *freenas.local*.

If the FreeNAS[®] server is not connected to a network with a DHCP server, use the console network configuration menu to manually configure the interface as shown here. In this example, the FreeNAS[®] system has one network interface, *em0*.

```
Enter an option from 1-12: 1
1) em0
Select an interface (q to quit): 1
Reset network configuration (y/n) n
Configure interface for DHCP? (y/n) n
Configure IPv4? (y/n) y
Interface name:
                   (press enter, the name can be blank)
Several input formats are supported
Example 1 CIDR Notation:
   192.168.1.1/24
Example 2 IP and Netmask separate:
   IP: 192.168.1.1
   Netmask: 255.255.255.0, or /24 or 24
IPv4 Address: 192.168.1.108/24
Saving interface configuration: Ok
Configure IPv6? (y/n) n
Restarting network: ok
. . .
The web user interface is at
http://192.168.1.108
```

After the system has an IP address, enter that address into a graphical web browser from a computer connected to the same network as the FreeNAS[®] system.

3.2 Logging In

By default, the login screen shown in Figure 3.2 prompts to log into the new UI.

E.	FreeNAS
Username * root	
Password *	
	LOG IN
LEGAC	Y WEB INTERFACE
FreeNAS ® (© 2018 - iXsystems, Inc.

Fig. 3.2: Enter the Root Password

To instead log into the legacy web interface, click *LEGACY WEB INTERFACE*. A prompt appears to confirm the choice.

Enter the password for the root user that was chosen during the installation. There is a prompt to set a root password if this was not set during the installation. The administrative GUI is displayed as shown in Figure 3.3.

FreeNAS																syst	
count System Tasks Network	Storage	Directory	Sharing	Services	Plugins	Jails	VMs	Reporting	T Wizard						Sup	Guide	Ю ок
and all collapse all	System																
Account	Information	Seneral	Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentia	ls Update	Alerts	Alert Services	CAs	Certificates	Support			
🙀 System 🔯 Tasks																	
Network	System Inf	formation	1														
Storage	Hostname	freenas.loc	al Edit														
Directory Service																	
Sharing Services	Build	FreeNAS-1	1.2-BETA2														
Plugins	Platform	AMD Ryzer	n 5 1600 Six-	Core Processo	or.												
Jails	Memory	16295MB															
VMs																	
Reporting	System Time	Mon, 30 Ju	1 2018 14:21:	28 -0400													
Wizard	Uptime	2:21PM up	4:30, 0 users	3													
Display System Processes	Load Average	1.32, 1.25,	1.25														
Shell																	
Keboot																	
Shutdown																	
			Fi	g. 3.3:	Free	eNAS [®]	Grap	hical C	onfigu	ratio	n Menu	L					

Note: The rest of this Guide describes the legacy UI. To access the Guide for the new UI, log into the new UI and click *Guide* or access it online at doc.freenas.org/11.2/freenas.html.

If the FreeNAS[®] system does not respond to the IP address or mDNS name entered in a browser:

- If proxy settings are enabled in the browser configuration, disable them and try connecting again.
- If the page does not load, check whether the FreeNAS[®] system's IP address responds to a ping from another computer on the same network. If the FreeNAS[®] IP address is in a private IP address range, it can only be accessed from within that private network.
- If the user interface loads but is unresponsive or seems to be missing menu items, try a different web browser. IE9 has known issues and does not display the graphical administrative interface correctly if compatibility mode is turned on. Firefox (https://www.mozilla.org/en-US/firefox/all/) is recommended.
- If *An error occurred!* messages are shown when attempting to configure an item in the GUI, make sure that the browser is set to allow cookies from the FreeNAS[®] system.

This blog post (http://fortysomethinggeek.blogspot.com/2012/10/ipad-iphone-connect-with-freenas-or-any.html) describes some applications which can be used to access the FreeNAS[®] system from an iPad or iPhone.

3.3 Initial Configuration

The first time the FreeNAS[®] GUI is accessed, the *Wizard* (page 281) starts automatically to help configure the FreeNAS[®] device quickly and easily.

ACCOUNT

The Account Configuration section of the web interface describes how to manually create and manage users and groups. This section contains these entries:

- Groups (page 58): used to manage UNIX-style groups on the FreeNAS[®] system.
- Users (page 60): used to manage UNIX-style accounts on the FreeNAS® system.

Each entry is described in more detail in this section.

4.1 Groups

The Groups interface provides management of UNIX-style groups on the FreeNAS[®] system.

Note: It is unnecessary to recreate the network users or groups when a directory service is running on the same network. Instead, import the existing account information into FreeNAS[®]. Refer to *Directory Services* (page 169) for details.

This section describes how to create a group and assign user accounts to it. The next section, *Users* (page 60), describes creating user accounts.

Click *Groups* \rightarrow *View Groups* to see a screen like Figure 4.1.

Account			
Groups	Users		
Add Group			
Group ID	Group Name	Built-in Group	Permit Sudo
0	wheel	true	false
1	daemon	true	false
2	kmem	true	false
3	sys	true	false
4	tty	true	false
5	operator	true	false
6	mail	true	false
7	bin	true	false
8	news	true	false
9	man	true	false
13	games	true	false
14	ftp	true	false
20	staff	true	false
22	sshd	true	false
25	smmsp	true	false
26	mailnull	true	false
31	guest	true	false
53	bind	true	false

Fig. 4.1: Group Management

The *Groups* page lists all groups, including those built-in and used by the operating system. The table displays group names, group IDs (GID), built-in groups, and if sudo is permitted. Clicking a group entry causes a *Members* button to appear. Click the button to view and modify the group membership

The *Add Group* button opens the screen shown in Figure 4.2. Table 4.1 summarizes the available options when creating a group.

A	dd Group		ж
	Group ID:	1001	
	Group Name:		
	Permit Sudo:		
	Allow repeated GIDs:		
	OK Cancel		

Fig. 4.2: Creating a New Group

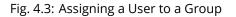
Setting	Value	Description
Group ID	string	The next available group ID is suggested. UNIX groups containing user ac- counts typically have an ID greater than 1000 and groups required by a service have an ID equal to the default port number used by the service. Example: the sshd group has an ID of 22.
Group Name	string	Required. Enter a descriptive name for the new group.
Permit Sudo	checkbox	Set to allow group members to use sudo (https://www.sudo.ws/). When using sudo, a user is prompted for their own password.
Allow repeated GIDs	checkbox	Set to allow multiple groups to share the same group id (GID). This is use- ful when a GID is already associated with the UNIX permissions for exist- ing data, but is generally not recommended.

Table 4.1: Group Creation Options

After a group and users are created, users can be added to a group. Highlight the group where users will be assigned, then click the *Members* button. Highlight the user in the *Member users* list. This shows all user accounts on the system. Click >> to move that user to the right frame. The user accounts which appear in the right frame are added as members of the group.

Figure 4.3, shows user1 added as a member of group data1.

Account				
Groups	Users			
Add Group				
Group ID	Group Name	Built-in Group	Permit	Sudo
1001	data1	false	false	Members 82
1002	user1	false	false	
0	wheel	true	false	
1	daemon	true	false	Available Selected
2	kmem	true	false	root
3	sys	true	false	daemon >>
4	tty	true	false	operator
5	operator	true	false	bin 🗸
6	mail	true	false	OK Cancel
7	bin	true	false	
8	news	true	false	
9	man	true	false	
13	games	true	false	
14	ftp	true	false	
20	staff	true	false	
22	sshd	true	false	
25	smmsp	true	false	
26	mailnull	true	false	



The *Delete Group* button deletes a group. The pop-up message asks whether all members of that group should also be deleted. Note that the built-in groups do not provide a *Delete Group* button.

4.2 Users

FreeNAS[®] supports users, groups, and permissions, allowing flexibility in configuring which users have access to the data stored on FreeNAS[®]. To assign permissions to shares, **one of these options** must be done:

- Create a guest account for all users, or create a user account for every user in the network where the name of each account is the same as a login name used on a computer. For example, if a Windows system has a login name of *bobsmith*, create a user account with the name *bobsmith* on FreeNAS[®]. A common strategy is to create groups with different sets of permissions on shares, then assign users to those groups.
- 2. If the network uses a directory service, import the existing account information using the instructions in *Directory Services* (page 169).

Account \rightarrow Users \rightarrow View Users lists all system accounts installed with the FreeNAS[®] operating system, as shown in Figure 4.4.

Groups	Users										
citabo											
Add User											
User ID	Username	Primary Group ID	Home Directory	Shell	Full Name	Built-in User	E-mail	Disable password login	Lock user	Permit Sudo	Microsoft Account
0	root	0	/root	/bin/csh	root	true		false	false	false	false
1	daemon	1	/root	/usr/sbin /nologin	Owner of many system processes	true		false	false	false	false
2	operator	5	1	/usr/sbin /nologin	System &	true		false	false	false	false
3	bin	7	/	/usr/sbin /nologin	Binaries Commands and Source	true		false	false	false	false
4	tty	65533	1	/usr/sbin /nologin	Tty Sandbox	true		false	false	false	false
5	kmem	2	1	/usr/sbin /nologin	KMem Sandbox	true		false	false	false	false
7	games	13	1	/usr/sbin /nologin	Games pseudo-user	true		false	false	false	false
8	news	8	1	/usr/sbin /nologin	News Subsystem	true		false	false	false	false
9	man	9	/usr/share/man	/usr/sbin /nologin	Mister Man Pages	true		false	false	false	false
14	ftp	14	/nonexistent	/bin/csh		true		false	false	false	false
22	sshd	22	/var/empty	/usr/sbin /nologin	Secure Shell Daemon	true		false	false	false	false
25	smmsp	25	/var/spool /clientmqueue	/usr/sbin /nologin	Sendmail Submission User	true		false	false	false	false
26	mailnull	26	/var/spool /mqueue	/usr/sbin /nologin	Sendmail Default User	true		false	false	false	false
53	bind	53	1	/usr/sbin /nologin	Bind Sandbox	true		false	false	false	false
62	proxy	62	/nonexistent	/usr/sbin /nologin	Packet Filter pseudo-user	true		false	false	false	false
64	_pflogd	64	/var/empty	/usr/sbin /nologin	pflogd privsep user	true		false	false	false	false
65	_dhcp	65	/var/empty	/usr/sbin /nologin	dhcp programs	true		false	false	false	false
66	uucp	66	/var/spool /uucppublic	/usr/local /libexec /uucp/uucico	UUCP pseudo- user	true		false	false	false	false
68	рор	6	/nonexistent	/usr/sbin /nologin	Post Office Owner	true		false	false	false	false
78	auditdistd	77	/var/empty	/usr/sbin /nologin	Auditdistd unprivileged user	true		false	false	false	false
79	ladvd	78	/var/empty	/usr/sbin /nologin	ladvd user	true		false	false	false	false
80	www	80	/nonexistent	/usr/sbin	World Wide	true		false	false	false	false

Modify User

Fig. 4.4: Managing User Accounts

Each account entry indicates the user ID, username, primary group ID, home directory, default shell, full name, whether it is a built-in user that came with the FreeNAS[®] installation, the email address, if logins are disabled, if the user account is locked, whether the user is allowed to use sudo, and if the user connects from a Windows 8 or newer system. To reorder the list, click the desired column name. An arrow indicates which column controls the view sort order. Click the arrow to reverse the sort order.

Click a user account to cause these buttons to appear:

- Modify User: used to modify the account's settings, as listed in Table 4.2.
- **Change E-mail:** used to change the email address associated with the account.

Note: Setting the the email address for the built-in *root* user account is recommended as important system messages are sent to the *root* user. For security reasons, password logins are disabled for the *root* account and changing this setting is discouraged.

Except for the *root* user, the accounts that come with FreeNAS[®] are system accounts. Each system account is used by a service and should not be used as a login account. For this reason, the default shell on system accounts is nologin(8) (https://www.freebsd.org/cgi/man.cgi?query=nologin). For security reasons and to prevent breakage of system services, do not modify the system accounts.

The *Add User* button opens the screen shown in Figure 4.5. Some settings are only available in *Advanced Mode*. To see these settings, either click *Advanced Mode* or configure the system to always display these settings by setting *Show advanced fields* by *default* in *System* \rightarrow *Advanced*. Table 4.2 summarizes the options which are available when user accounts are created or modified.

Add User		_	88
User ID:	1001		
Username:			
Create a new primary group for the user:			
Primary Group:	v		
Create Home Directory In:	/nonexistent	Browse	
Shell:	csh		
Full Name:			
E-mail:			
Password:			
Password confirmation:		(i)	
Disable password login:			

Fig. 4.5: Adding or Editing a User Account

Setting	Value	Advanced Mode	Description
User ID	integer		Grayed out if the user already exists. When creating an ac- count, the next numeric ID is suggested. User accounts typi- cally have an ID greater than 1000 and system accounts have an ID equal to the default port number used by the service.
Username	string		Enter an alphanumeric username of eight to sixteen characters. Keeping usernames to eight characters or less is recommended for compatibility with legacy clients. Note that \$ can only be used as the last character. Usernames cannot begin with a hyphen – or contain a space, tab, or these characters: , : + & #
Create a new primary	checkbox		A primary group with the same name as the user is created au-
group			tomatically. Unset to select a different primary group name.
Primary Group	drop-down menu		Unset <i>Create a new primary group</i> to access this menu. For se- curity reasons, FreeBSD does not give a user su permissions if <i>wheel</i> is their primary group. To give a user su access, add them to the <i>wheel</i> group in <i>Auxiliary groups</i> .
Create Home Direc-	browse button		Browse to the name of an existing volume or dataset that the
tory ln			user will be assigned permission to access.
Home Directory	checkboxes	\checkmark	Sets default Unix permissions of the user's home directory.
Mode			This is read-only for built-in users.
Shell	drop-down menu		Select the shell to use for local and SSH logins. See Table 4.3 for an overview of available shells.
Full Name	string		Required. This field may contain spaces.
E-mail	string		The email address associated with the account.
Password	string		Required unless <i>Disable password login</i> is set. Cannot contain a ?.
Password confirma- tion	string		This must match the value of <i>Password</i> .
Disable password login	checkbox		Set to disable password logins and authentication to SMB shares. To undo this setting, create a password for the user by clicking <i>Modify User</i> for the user in the <i>View Users</i> screen. Setting this grays out <i>Lock user</i> and <i>Permit Sudo</i> .
Lock user	checkbox		Set to prevent the user from logging in until this box is unset. Setting this grays out <i>Disable password login</i> .
Permit Sudo	checkbox		Set to give group members permission to use sudo (https://www.sudo.ws/). When using sudo, a user is prompted for their own password.
Microsoft Account	checkbox		Set this when the user is connecting from a Windows 8 or newer system.
SSH Public Key	string		Enter or paste the user's public SSH key to be used for key- based authentication. Do not paste the private key!
Auxiliary groups	mouse selection		Highlight groups to add the user. Click the >> to add the user to the highlighted groups.

Table 4.2: User Account Configuration

Note: Some fields cannot be changed for built-in users and will be grayed out.

Shell	Description
netcli.sh	User is shown the Console Setup menu (Figure 3.1) on connection, even if it is
	disabled in System \rightarrow Advanced \rightarrow Enable Console Menu. The user must be root
	or have root permissions (effective user ID 0, like toor).
csh	C shell (https://en.wikipedia.org/wiki/C_shell)
sh	Bourne shell (https://en.wikipedia.org/wiki/Bourne_shell)
tcsh	Enhanced C shell (https://en.wikipedia.org/wiki/Tcsh)
nologin	Use when creating a system account or to create a user account that can au-
	thenticate with shares but which cannot login to the FreeNAS system using
	ssh.
bash	Bourne Again shell (https://en.wikipedia.org/wiki/Bash_%28Unix_shell%29)
ksh93	Korn shell (http://www.kornshell.com/)
mksh	mirBSD Korn shell (https://www.mirbsd.org/mksh.htm)
rbash	Restricted bash (http://www.gnu.org/software/bash/manual/html_node/The-
	Restricted-Shell.html)
rzsh	Restricted zsh (http://www.csse.uwa.edu.au/programming/linux/zsh-
	doc/zsh_14.html)
scponly	Select scponly (https://github.com/scponly/scponly/wiki) to restrict the user's
	SSH usage to only the scp and sftp commands.
zsh	Z shell (http://www.zsh.org/)
git-shell	restricted git shell (https://git-scm.com/docs/git-shell)

Tabla	1 7.	Available	Challe
Table	4 ⊰`	Available	Snells

Built-in user accounts needed by the system cannot be removed. A *Remove User* button appears for custom users that were added by the system administrator. If the user to be removed is the last user in a custom group, an option is offered to keep the user primary group after deleting the user.

SYSTEM

The System section of the administrative GUI contains these entries:

- *Information* (page 65) provides general FreeNAS[®] system information such as hostname, operating system version, platform, and uptime
- General (page 66) configures general settings such as HTTPS access, the language, and the timezone
- Boot (page 69) creates, renames, and deletes boot environments. It also shows the condition of the Boot Volume.
- Advanced (page 72) configures advanced settings such as the serial console, swap space, and console messages
- Email (page 75) configures the email address to receive notifications
- System Dataset (page 76) configures the location where logs and reporting graphs are stored
- Tunables (page 77) provides a front-end for tuning in real-time and to load additional kernel modules at boot time
- Update (page 79) performs upgrades and checks for system updates
- · Cloud Credentials (page 82) is used to enter connection credentials for remote cloud service providers
- *Alerts* (page 85) lists the available *Alert* (page 296) conditions and provides configuration of the notification frequency for each alert.
- Alert Services (page 85) configures services used to notify the administrator about system events.
- CAs (page 87): import or create internal or intermediate CAs (Certificate Authorities)
- Certificates (page 89): import existing certificates or create self-signed certificates
- Support (page 92): report a bug or request a new feature.

Each of these is described in more detail in this section.

5.1 Information

System \rightarrow *Information* displays general information about the FreeNAS[®] system. An example is seen in Figure 5.1.

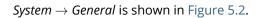
The information includes hostname, build version, type of CPU (platform), amount of memory, current system time, system uptime, number of users connected at the console or by serial, telnet, or SSH connections, and current load average. On systems supplied or certified by iXsystems, an additional *Serial Number* field showing the hardware serial number is displayed.

To change the system hostname, click the *Edit* button, type in the new hostname, and click *OK*. The hostname must include the domain name. If the network does not use a domain name, add *.local* after the hostname.

System													
Information	General	Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts	Alert Services	CAs	Certificates	Support
System I	nformat	ion											
Hostname	freen	nas.local	Edit										
Build	Free	NAS-11.2-	BETA1										
Platform	Intel(R) Atom(T	FM) CPU C2750	@ 2.40G	Hz								
Memory	3270	2MB											
System Tin	ne Mon,	25 Jun 20	018 06:13:20 -0	700									
Uptime	6:13/	AM up 2 d	ays, 23:31, 0 ι	isers									
Load Avera	ige 0.14,	.14, 0.15, 0.14											
System Se	rial												
System Pro	duct FREE	REENAS-MINI-2.0											
System Pro	duct FREE	NAS-MINI-	2.0										

Fig. 5.1: System Information Tab

5.2 General



System													
Information	General	Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts	Alert Services	CAs	Certificates	Support
Protoco	:			HTTP 🔻									
Certifica	te:			💌									
WebGUI	IPv4 Addres	s:		0.0.0.0 🔻									
WebGUI	IPv6 Addres	s:		:: 💌									
WebGUI	HTTP Port:			80									
WebGUI	HTTPS Port:			443									
WebGUI	HTTP -> HTT	PS Redire	ct:	🔽 (Ì)									
Languag	je (Require U	I reload):		English	•								
Console	Keyboard Ma	ap:			•								
Timezon	e:			America/Lo	s_Angeles 🔻								
Syslog l	evel:			Info 🔻 (i)								
Syslog s	erver:					i							
Save	Reset Configu	ration to De	faults	ve Config	Upload Config	NTP Server	5						

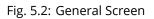


Table 5.1 summarizes the configurable settings in the General tab:

Setting	Value	Description
Protocol	drop-down	Set the web protocol to use when connecting to the administrative GUI
	menu	from a browser. To change the default <i>HTTP</i> to <i>HTTPS</i> or to <i>HTTP</i> + <i>HTTPS</i> ,
		select a certificate to use in <i>Certificate</i> . If there are no certificates, first cre-
		ate a <i>CA</i> (page 87) then a <i>certificate</i> (page 89).
Certificate	drop-down	Required for HTTPS. Browse to the location of the certificate to use for en-
	menu	crypted connections.
WebGUI IPv4 Address	drop-down	Choose a recent IP address to limit the usage when accessing the admin-
	menu	istrative GUI. The built-in HTTP server binds to the wildcard address of
		0.0.0.0 (any address) and issues an alert if the specified address becomes
		unavailable.
WebGUI IPv6 Address	drop-down	Choose a recent IPv6 address to limit the usage when accessing the ad-
	menu	ministrative GUI. The built-in HTTP server binds to any address issues an
		alert if the specified address becomes unavailable.
WebGUI HTTP Port	integer	Allow configuring a non-standard port for accessing the
		administrative GUI over HTTP. Changing this setting can
		also require changing a Firefox configuration setting
		(https://www.redbrick.dcu.ie/~d_fens/articles/Firefox:_This_Address_is_Restricter
WebGUI HTTPS Port	integer	Allow configuring a non-standard port for accessing the administrative
		GUI over HTTPS.
WebGUI HTTP -> HTTPS	checkbox	Set to redirect HTTP connections to HTTPS. HTTPS must be selected in Pro-
Redirect		tocol.
Language	drop-down	Select a localization. View the status of the localization at we-
	menu	blate.trueos.org (https://weblate.trueos.org/projects/freenas/).
Console Keyboard Map	drop-down	Select a keyboard layout.
	menu	
Timezone	drop-down	Select a timezone.
	menu	
Syslog level	drop-down	When <i>Syslog server</i> is defined, only logs matching this level are sent.
	menu	
Syslog server	string	Select an IP address_or_hostname:optional_port_number to send logs to. Set
		to write log entries to both the console and the remote server.

Table 5.1: General Configuration Settings

After making any changes, click the *Save* button.

This screen also contains these buttons:

Reset Configuration to Defaults: reset the configuration database to the default base version. This does not delete user SSH keys or any other data stored in a user home directory. Since configuration changes stored in the configuration database are erased, this option is useful when a mistake has been made or to return a test system to the original configuration.

Save Config: save a backup copy of the current configuration database in the format *hostname-version-architecture* to the computer accessing the administrative interface. Saving the configuration after making any configuration changes is highly recommended. FreeNAS[®] automatically backs up the configuration database to the system dataset every morning at 3:45. However, this backup does not occur if the system is shut down at that time. If the system dataset is stored on the boot pool and the boot pool becomes unavailable, the backup will also not be available. The location of the system dataset is viewed or set using *System* \rightarrow *System Dataset*.

Note: *SSH* (page 245) keys are not stored in the configuration database and must be backed up separately.

There are two types of passwords. User account passwords for the base operating system are stored as hashed values, do not need to be encrypted to be secure, and are saved in the system configuration backup. Other passwords, like iSCSI CHAP passwords, Active Directory bind credentials, and cloud credentials are stored in an encrypted form to prevent them from being visible as plain text in the saved system configuration. The key or *seed* for this encryption is normally stored

only on the boot device. When *Save Config* is chosen, a dialog gives the option to *Export Password Secret Seed* with the saved configuration, allowing the configuration file to be restored to a different boot device where the decryption seed is not already present. Configuration backups containing the seed must be physically secured to prevent decryption of passwords and unauthorized access.

Warning: The *Export Password Secret Seed* option is off by default and should only be used when making a configuration backup that will be stored securely. After moving a configuration to new hardware, media containing a configuration backup with a decryption seed should be securely erased before reuse.

Upload Config: allows browsing to the location of a previously saved configuration file to restore that configuration. The screen turns red as an indication that the system will need to reboot to load the restored configuration.

NTP Servers: The network time protocol (NTP) is used to synchronize the time on the computers in a network. Accurate time is necessary for the successful operation of time sensitive applications such as Active Directory or other directory services. By default, FreeNAS[®] is pre-configured to use three public NTP servers. If the network is using a directory service, ensure that the FreeNAS[®] system and the server running the directory service have been configured to use the same NTP servers.

Available NTP servers can be found at https://support.ntp.org/bin/view/Servers/NTPPoolServers. For time accuracy, choose NTP servers that are geographically close to the physical location of the FreeNAS[®] system.

Click *NTP Servers* \rightarrow *Add NTP Server* to add an NTP server. Figure 5.3 shows the screen that appears. Table 5.2 summarizes the options available when adding an NTP server. ntp.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=ntp.conf) explains these options in more detail.

NTP Servers								
Add NTP Server								
Address	Burst	IBurst	Prefer	Min. Poll	Max. Poll	Add NTP Server	86	
0.freebsd.pool.ntp.org	false	true	false	6	10	Add NTP Server	~	
1.freebsd.pool.ntp.org	false	true	false	6	10			
2.freebsd.pool.ntp.org	false	true	false	6	10	Address:		
						Burst:		
						IBurst: 🔽 🚺		
						Prefer:		
						Min. Poll: 6	i	
						Max. Poll: 10	i	
						Force:		
						ОК Сапсе!		

Fig. 5.3: Add an NTP Server

Table 5.2: NTP Servers Configuration Options

AddressstringEnter the hostname or IP address of the NTP server.BurstcheckboxRecommended when Max. Poll is greater than 10. Only use on private servers. Do not use with a public NTP server.	Setting	Value	Description
	Address	string	Enter the hostname or IP address of the NTP server.
	Burst	checkbox	

Continued on next page

Setting	Value	Description
lBurst	checkbox	Speed up the initial synchronization, taking seconds rather than minutes.
Prefer	checkbox	This option is only recommended for highly accurate NTP servers, such as
		those with time monitoring hardware.
Min. Poll	integer	Minimum polling time in seconds. Must be a power of 2, and cannot be
		lower than 4 or higher than Max. Poll.
Max. Poll	integer	Maximum polling time in seconds. Must be a power of 2, and cannot be
		higher than 17 or lower than <i>Min. Poll</i> .
Force	checkbox	Force the addition of the NTP server, even if it is currently unreachable.

Table 5.2 – continued f	from	previous	page
-------------------------	------	----------	------

5.3 Boot

FreeNAS[®] supports a ZFS feature known as multiple boot environments. With multiple boot environments, the process of updating the operating system becomes a low-risk operation. The updater automatically creates a snapshot of the current boot environment and adds it to the boot menu before applying the update.

If an update fails, reboot the system and select the previous boot environment, using the instructions in *If Something Goes Wrong* (page 33), to instruct the system to go back to that system state.

Note: Boot environments are separate from the configuration database. Boot environments are a snapshot of the *operating system* at a specified time. When a FreeNAS[®] system boots, it loads the specified boot environment, or operating system, then reads the configuration database to load the current configuration values. If the intent is to make configuration changes rather than operating system changes, make a backup of the configuration database first using *System* \rightarrow *General* \rightarrow *Save Config.*

As seen in Figure 5.4, FreeNAS[®] displays the condition and statistics of the *Boot Volume*. It also shows the two boot environments that are created when FreeNAS[®] is installed. The system will boot into the *default* boot environment and users can make their changes and update from this version. The *Initial-Install* boot environment can be booted into if the system needs to be returned to a non-configured version of the installation.

If the *Wizard* (page 281) was used, a third boot environment called <code>Wizard-date</code> is also created, indicating the date and time the *Wizard* (page 281) was run.

System							
Information General Boot Advanced I	nail System Dataset Tunables Cloud Credentials	Update Alerts Alert Services CAs Cer	tificates Support				
Scrub Boot Status Boot Volume Condition: HEALTHY Size: 29.5 GiB Used: 873.2 MiB (2%)							
Name	Active	Created	Кеер				
default	On Reboot, Now	2018-06-22 06:32:00	No				
Initial-Install		2018-06-22 06:44:00	No				

Fig. 5.4: Viewing Boot Environments

Each boot environment entry contains this information:

- **Name:** the name of the boot entry as it will appear in the boot menu.
- Active: indicates which entry will boot by default if the user does not select another entry in the boot menu.
- **Created:** indicates the date and time the boot entry was created.
- **Keep:** indicates whether or not this boot environment can be pruned if an update does not have enough space to proceed. Click *Keep* for an entry if that boot environment should not be automatically pruned.

Highlight an entry to view the configuration buttons for it. These configuration buttons are shown:

- **Rename:** used to change the name of the boot environment.
- **Keep/Unkeep:** used to toggle whether or not the updater can prune (automatically delete) this boot environment if there is not enough space to proceed with the update.
- Clone: makes a new boot environment from the selected boot environment.
- **Delete:** used to delete the highlighted entry, which also removes that entry from the boot menu. Since an activated entry cannot be deleted, this button does not appear for the active boot environment. To delete an entry that is currently activated, first activate another entry, which will clear the *On reboot* field of the currently activated entry. Note that this button does not appear for the *default* boot environment as this entry is needed to return the system to the original installation state.
- Activate: only appears on entries which are not currently set to *Active*. Changes the selected entry to the default boot entry on next boot. The status changes to *On Reboot* and the current *Active* entry changes from *On Reboot*, *Now* to *Now*, indicating that it was used on the last boot but will not be used on the next boot.

The buttons above the boot entries can be used to:

- **Create:** makes a new boot environment from the active environment. The active boot environment contains the text On Reboot, Now in the Active column. Only alphanumeric characters, underscores, and dashes are allowed in the name.
- **Scrub Boot:** can be used to perform a manual scrub of the boot devices. By default, the boot device is scrubbed every 7 days. To change the default interval, change the number in the *Automatic scrub interval (in days)* field. The date and results of the last scrub are also listed in this screen. The condition of the boot device should be listed as *HEALTHY*.
- Status: click this button to see the status of the boot devices. Figure 5.5, shows only one boot device, which is ONLINE.

Note: Using *Clone* to clone the active boot environment functions the same as using *Create*.

Boot Status				
Name	Read	Write	Checksum	Status
⊿ freenas-boot	0	0	0	ONLINE
▲ stripe	0	0	0	ONLINE
da0p2	0	0	0	ONLINE

Replace

Fig. 5.5: Viewing the Status of the Boot Device

If the system has a mirrored boot pool, there will be a *Detach* button in addition to the *Replace* button. To remove a device from the boot pool, highlight the device and click its *Detach* button. Alternately, if one of the boot devices has an *OFFLINE Status*, click the device to replace, then click *Replace* to rebuild the boot mirror.

Note that **the boot device cannot be replaced if it is the only boot device** because it contains the operating system itself.

5.3.1 Mirroring the Boot Device

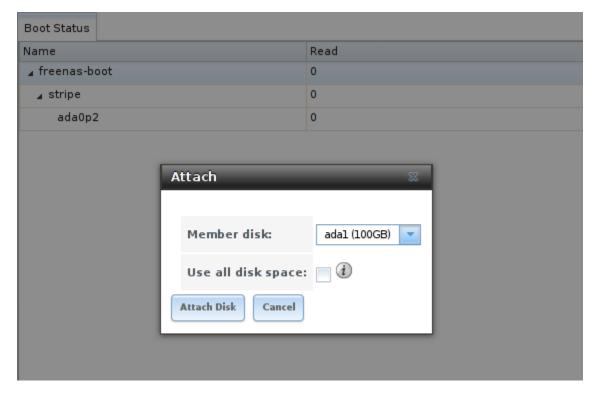
If the system is currently booting from a device, another device can be added to create a mirrored boot device. If one device in a mirror fails, the remaining device can still be used to boot the system.

Note: When adding another boot device for a mirror, the new device must have at least the same capacity as the existing boot device. Larger capacity devices can be added, but the mirror will only have the capacity of the smallest device. Different models of devices which advertise the same nominal size are not necessarily the same actual size. For this reason, adding another of the same model of boot device is recommended.

In the example shown in Figure 5.6, the user has clicked *System* \rightarrow *Boot* \rightarrow *Status* to display the current status of the boot device. The example indicates that there is currently one device, ada0p2, its status is *ONLINE*, and it is currently the only boot device as indicated by the word *stripe*. To create a mirrored boot device, click either the entry called *freenas-boot* or *stripe*, then click the *Attach* button. If another device is available, it appears in the *Member disk* drop-down menu. Select the desired device.

The *Use all disk space* option gives control of how much of the new device is made available to ZFS. The new device is partitioned to the same size as the existing device by default. Select *Use all disk space* to use all available space on the new device. If either device in the mirror fails, it can be replaced with another of the same size as the original boot device.

When *Use all disk space* is enabled, the entire capacity of the new device is used. If the original boot device fails and is removed, the boot mirror will consist of just the newer drive, and will grow to whatever capacity it provides. However, new devices added to this mirror must now be as large as the new capacity.



Click *Attach Disk* to attach the new disk to the mirror.

Fig. 5.6: Mirroring a Boot Device

After the mirror is created, the *Status* screen indicates that it is now a *mirror*. The number of devices in the mirror are shown as in Figure 5.7.

Boot Status				
Name	Read	Write	Checksum	Status
▲ freenas-boot	0	0	0	ONLINE
⊿ mirror-0	0	0	0	ONLINE
ada1p2	0	0	0	ONLINE
ada0p2	0	0	0	ONLINE

Fig. 5.7: Viewing the Status of a Mirrored Boot Device

5.4 Advanced

System \rightarrow *Advanced* is shown in Figure 5.8. The configurable settings are summarized in Table 5.3.

System	
Information General Boot Advanced Email System Dataset Tunables Cloud Credentials Update Alerts	Alert Services CAs Certificates Support
Show Text Console without Password Prompt:	
Use Serial Console:	
Serial Port Address:	oas -
Serial Port Speed:	
Enable powerd (Power Saving Daemon):	
Swap size on each drive in GiB, affects new disks only. Setting this to 0 disables swap creation completely (STRONGLY DISCOURAGED)	2
Show console messages in the footer:	
Show tracebacks in case of fatal errors:	
Show advanced fields by default:	
Enable autotune:	
Enable debug kernel:	
MOTD banner:	Welcome to FreeMAS
Periodic Notification User:	root 🖉
Report CPU usage in percentage:	
Remote Graphite Server Hostname:	
Use FQDN for logging:	
ATA Security User:	User 👻 🕄
SED Password:	
Reset SED Password:	
Save Debug	

Fig. 5.8: Advanced Screen

Setting	Value	Description
Show Text Console without	checkbox	Set for the system to immediately display the text console after boot-
Password Prompt		ing. Unset to require logging into the system before the console menu is
		shown.
Use Serial Console	checkbox	Do not enable this option if the serial port is disabled.
Serial Port Address	string	Select the serial port address in hex.
Serial Port Speed	drop-down	Select the speed used by the serial port.
	menu	
Enable powerd (Power Sav-	checkbox	powerd(8) (https://www.freebsd.org/cgi/man.cgi?query=powerd) monitors
ing Daemon)		the system state and sets the CPU frequency accordingly.
Swap size	non-zero	By default, all data disks are created with this amount of swap. Log or
	integer rep-	cache devices do not create with swap and are unaffected. Setting to 0
	resenting	disables swap creation completely. This is <i>strongly</i> discouraged.
	GiB	
Show console messages in	checkbox	Set to display console messages in real time at the bottom of the browser.
the footer		Click the console to bring up a scrollable screen. Set <i>Stop refresh</i> in the
		scrollable screen to pause updating, and deselect the option to continue
		to watch the messages as they occur.
Show tracebacks in case of	checkbox	Open a pop-up of diagnostic information when a fatal error occurs.
fatal errors		
Show advanced fields by	checkbox	Show Advanced Mode fields by default.
default		
Enable autotune	checkbox	Enable an <i>Autotune</i> (page 73) script which attempts to optimize the sys-
		tem based on the installed hardware. <i>Warning</i> : Autotuning is only used
		as a temporary measure and is not a permanent fix for system hardware
<u> </u>		issues.
Enable debug kernel	checkbox	Use a debug version of the kernel on the next boot.
MOTD banner	string	This message is shown when a user logs in with SSH.
Periodic Notification User	drop-down	Choose a user to receive security output emails. This output runs nightly
	menu	but only sends email when the system reboots or encounters an error.
Report CPU usage in per-	checkbox	Display CPU usage as percentages in <i>Reporting</i> (page 279).
centage		
Remote Graphite Server	string	IP address or hostname of a remote server running Graphite
hostname		(http://graphiteapp.org/).
Use FQDN for logging	checkbox	Include the Fully-Qualified Domain Name in logs to precisely identify sys-
		tems with similar hostnames.
ATA Security User	drop-down	User passed to camcontrol security -u for unlocking Self-Encrypting
	menu	Drives (page 74). Values are User or Master.
SED Password	string	Global password used to unlock <i>Self-Encrypting Drives</i> (page 74).
Reset SED Password	checkbox	Select to clear the Password for SED column of Storage $ ightarrow$ View Disks.

Table 5.3:	Advanced	Configurati	on Settings

Click the *Save* button after making any changes.

This tab also contains this button:

Save Debug: used to generate a text file of diagnostic information. After the debug data is collected, the system prompts for a location to save the compressed .tgz text file.

5.4.1 Autotune

FreeNAS[®] provides an autotune script which optimizes the system depending on the installed hardware. For example, if a ZFS volume exists on a system with limited RAM, the autotune script automatically adjusts some ZFS sysctl values in an attempt to minimize ZFS memory starvation issues. It should only be used as a temporary measure on a system that hangs until the underlying hardware issue is addressed by adding more RAM. Autotune will always slow such a system, as it caps the ARC.

The *Enable autotune* option in *System* \rightarrow *Advanced* is off by default. Enable this option to run the autotuner at boot time. To run the script immediately, reboot the system.

If the autotune script adjusts any settings, the changed values appear in *System* \rightarrow *Tunables*. These values can be modified and overridden. Note that deleting tunables that were created by autotune only affects the current session, as autotune-set tunables are recreated at boot.

When attempting to increase the performance of the FreeNAS[®] system, and particularly when the current hardware may be limiting performance, try enabling autotune.

For those who wish to see which checks are performed, the autotune script is located in /usr/local/bin/autotune.

5.4.2 Self-Encrypting Drives

FreeNAS[®] version 11.1-U5 introduced Self-Encrypting Drive (SED) support.

Three types of SED devices are supported:

- Legacy interface for older ATA devices. Not recommended for security-critical environments
- TCG OPAL 2 standard for newer consumer-grade devices (HDD or SSD over PCIe or SATA)
- TCG Enterprise standard for newer enterprise-grade SAS devices

The FreeNAS[®] middleware implements the security capabilities of camcontrol (https://www.freebsd.org/cgi/man.cgi?query=camcontrol (for legacy devices) and sedutil-cli (https://www.mankier.com/8/sedutil-cli) (for TCG devices). When managing SED devices from the command line, it is important to use sedutil-cli rather than camcontrol to access the full capabilities of the device. FreeNAS[®] provides the sedhelper wrapper script to ease SED device administration from the command line.

By default, SED devices are not locked until the administrator explicitly configures a global or per-device password and initializes the devices.

Once configured, the system automatically unlocks all SEDs during the boot process, without requiring manual intervention. This allows a pool to contain a mix of SED and non-SED devices.

A password-protected SED device protects the data stored on the device when the device is physically removed from the FreeNAS[®] system. This allows secure disposal of the device without having to first wipe its contents. If the device is instead removed to be repurposed on another system, it can only be unlocked if the password is known.

Warning: It is important to remember the password! Without it, the device is unlockable and its data remains unavailable. While it is possible to specify the PSID number on the label of the device with the sedutil-cli command, doing so will erase the contents of the device rather than unlock it. Always record SED passwords whenever they are configured or modified and store them in a safe place!

When SED devices are detected during system boot, the middleware checks for global and device-specific passwords. Devices with their own password are unlocked with their password and any remaining devices, without a device-specific password, are unlocked using the global password.

To configure a global password, go to System \rightarrow Advanced \rightarrow SED Password and enter the password. Recording the password and storing it in a safe place is recommended.

To determine which devices support SED and their device names:

sedutil-cli --scan

In the results:

- no indicates a non-SED device
- 1 indicates a legacy TCG OPAL 1 device
- 2 indicates a modern TCG OPAL 2 device
- E indicates a TCG Enterprise device

To specify a password for a device, go to *Storage* \rightarrow *View Disks*. Highlight the device name for the confirmed SED device and click *Edit*. Enter and confirm the password in the *Password for SED* and *Confirm SED Password* fields. Disks that have a configured password will show bullets in their row of the *Password for SED* column of *Storage* \rightarrow *View Disks*. Conversely, the rows in that column will be empty for disks that do not support SED or which are unlocked using the global password.

Next, remember to initialize the devices:

sedhelper setup password

This command ensures that all detected SED disks are properly setup using the specified password.

Note: Rerun sedhelper setup password every time a new SED disk is placed in the system.

This command is used to unlock all available SED disks:

sedhelper unlock

5.5 Email

An automatic script sends a nightly email to the *root* user account containing important information such as the health of the disks. *Alert* (page 296) events are also emailed to the *root* user account. Problems with *Scrubs* (page 162) are reported separately in an email sent at 03:00AM.

Note: *S.M.A.R.T.* (page 237) reports are mailed separately to the address configured in that service.

The administrator typically does not read email directly on the FreeNAS[®] system. Instead, these emails are usually sent to an external email address where they can be read more conveniently. It is important to configure the system so it can send these emails to the administrator's remote email account so they are aware of problems or status changes.

The first step is to set the remote address where email will be sent. Select *Account* \rightarrow *Users*, click on *root* to highlight that user, then click *Modify User*. In the *E-mail* field, enter the email address on the remote system where email is to be sent, like *admin@example.com*. Click *OK* to save the settings.

System General Email System Dataset Tunables Cloud Credentials Information Boot Advanced Update Alerts Alert Services CAs Certificates Support i From email: root@freenas.local a Outgoing mail server: a Port to connect to: 25 Plain i TLS/SSL: Use SMTP Authentication: ì Username: Password: ì Password confirmation: HINT: Test e-mails are sent to root user. To configure it use Account -> Users -> View Users -> root -> Modify User Send Test Mail Save

Additional configuration is performed with *System* \rightarrow *Email*, shown in Figure 5.9.

Fig. 5.9: Email Screen

Setting	Value	Description
From email	string	Setting a known From address is helpful in filtering mail on the receiv-
		ing system.
Outgoing mail server	string or IP address	Hostname or IP address of SMTP server used for sending this email.
Port to connect to	integer	SMTP port number. Typically 25, 465 (secure SMTP), or 587 (submis-
		sion).
TLS/SSL	drop-down menu	Choose an encryption type. Choices are <i>Plain</i> , SSL, or <i>TLS</i>
Use SMTP Authenti-	checkbox	Enable or disable SMTP AUTH
cation		(https://en.wikipedia.org/wiki/SMTP_Authentication) using PLAIN
		SASL. If enabled, enter the required <i>Username</i> and <i>Password</i> .
Username	string	Enter the SMTP username if the SMTP server requires authentication.
Password	string	Enter the SMTP password if the SMTP server requires authentication.
Password Confirma-	string	Confirm the SMTP password.
tion		

Table 5.4: Email Configuration Settings

Click the Send Test Mail button to verify that the configured email settings are working. If the test email fails, double-check that the *E-mail* field of the *root* user is correctly configured by clicking the *Modify User* button for the *root* account in Account \rightarrow Users \rightarrow View Users.

Configuring email for TLS/SSL email providers is described in Are you having trouble getting FreeNAS to email you in Gmail? (https://forums.freenas.org/index.php?threads/are-you-having-trouble-getting-freenas-to-email-you-in-gmail.22517/).

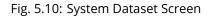
Note: The FreeNAS[®] user who receives periodic email is set in the *Periodic Notification User* field in *System* \rightarrow *Advanced*.

5.6 System Dataset

System \rightarrow System Dataset, shown in Figure 5.10, is used to select the pool which contains the persistent system dataset. The system dataset stores debugging core files and Samba4 metadata such as the user or group cache and share level permissions. If the FreeNAS[®] system is configured to be a Domain Controller, all of the domain controller state is stored there as well, including domain controller users and groups.

Note: When the system dataset is moved, a new dataset is created and set active. The old dataset is intentionally not deleted by the system because the move might be transient or the information in the old dataset might be useful for later recovery.

System													
Information	General	Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts	Alert Services	CAs	Certificates	Support
System d	ataset pool:	volume1											
Syslog:													
Reporting	g Database:	Ì											
Save													



Note: Encrypted, locked volumes are not displayed in the *System dataset pool* drop-down menu.

The system dataset can optionally be configured to also store the system log and *Reporting* (page 279) information. If there are lots of log entries or reporting information, moving these to the system dataset will prevent /var/ on the device holding the operating system from filling up as /var/ has limited space.

Use the drop-down menu to select the ZFS volume (pool) to contain the system dataset. Whenever the location of the system dataset is changed, a pop-up warning indicates that the SMB service must be restarted, causing a temporary outage of any active SMB connections.

To store the system log on the system dataset, enable the *Syslog* option.

To store the reporting information on the system dataset, enable the *Reporting Database* option. When this option is not enabled, a RAM disk is created to prevent reporting information from filling up /var.

Click the *Save* button to save changes.

If the pool storing the system dataset is changed at a later time, FreeNAS[®] migrates the existing data in the system dataset to the new location.

Note: Depending on configuration, the system dataset can occupy a large amount of space and receive frequent writes. Do not put the system dataset on a flash drive or other media with limited space or write life.

5.7 Tunables

System \rightarrow *Tunables* can be used to manage:

- 1. **FreeBSD sysctls:** a sysctl(8) (https://www.freebsd.org/cgi/man.cgi?query=sysctl) makes changes to the FreeBSD kernel running on a FreeNAS[®] system and can be used to tune the system.
- 2. **FreeBSD loaders:** a loader is only loaded when a FreeBSD-based system boots and can be used to pass a parameter to the kernel or to load an additional kernel module such as a FreeBSD hardware driver.
- 3. FreeBSD rc.conf options: rc.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rc.conf&manpath=FreeBSD+11.0-RELEASE) is used to pass system configuration options to the system startup scripts as the system boots. Since FreeNAS[®] has been optimized for storage, not all of the services mentioned in rc.conf(5) are available for configuration. Note that in FreeNAS[®], customized rc.conf options are stored in /tmp/rc.conf.freenas.

Warning: Adding a sysctl, loader, or rc.conf option is an advanced feature. A sysctl immediately affects the kernel running the FreeNAS[®] system and a loader could adversely affect the ability of the FreeNAS[®] system to successfully boot. **Do not create a tunable on a production system unless it is understood and ramifications have been tested for that change.**

Since sysctl, loader, and rc.conf values are specific to the kernel parameter to be tuned, the driver to be loaded, or the service to configure, descriptions and suggested values can be found in the man page for the specific driver and in many sections of the FreeBSD Handbook (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/).

To add a loader, sysctl, or rc.conf option, go to System \rightarrow Tunables \rightarrow Add Tunable, to access the screen shown in Figure 5.11.

Add Tunable	ж
Variable:	
Value:	
Type: Loader 💌	
Comment:	
Enabled: 💟	
OK Cancel	

Fig. 5.11: Adding a Tunable

Table 5.5 summarizes the options when adding a tunable.

Table 5.5: Adding a Tunable

Setting	Value	Description
Variable	string	The name of the sysctl or driver to load.
Value	integer or string	Set a value for the <i>Variable</i> . Refer to the man page for the specific driver or the FreeBSD Handbook (https://www.freebsd.org/doc/en_US.ISO8859- 1/books/handbook/) for suggested values.
Туре	drop-down menu	Choices are Loader, rc.conf, or Sysctl.
Comment	string	Enter a userful description of this tunable.
Enabled	checkbox	Unset this option to disable the tunable without deleting it.

Note: As soon as a *Sysctl* is added or edited, the running kernel changes that variable to the value specified. However, when a *Loader* or *rc.conf* value is changed, it does not take effect until the system is rebooted. Regardless of the type of tunable, changes persist at each boot and across upgrades unless the tunable is deleted or the *Enabled* option is deselected.

Any added tunables are listed in *System* \rightarrow *Tunables*. To change the value of an existing tunable, click its *Edit* button. To remove a tunable, click its *Delete* button.

Restarting the FreeNAS[®] system after making sysctl changes is recommended. Some sysctls only take effect at system startup, and restarting the system guarantees that the setting values correspond with what is being used by the running system.

The GUI does not display the sysctls that are pre-set when FreeNAS[®] is installed. FreeNAS[®] 11.2 ships with these sysctls set:

```
kern.metadelay=3
kern.dirdelay=4
kern.filedelay=5
kern.coredump=1
vfs.timestamp_precision=3
net.link.lagg.lacp.default_strict_mode=0
vfs.zfs.min_auto_ashift=12
```

Do not add or edit these default sysctls as doing so may render the system unusable.

The GUI does not display the loaders that are pre-set when FreeNAS[®] is installed. FreeNAS[®] 11.2 ships with these loaders set:

```
autoboot_delay="2"
loader_logo="freenas"
loader_menu_title="Welcome to FreeNAS"
loader_brand="freenas-brand"
loader_version=" "
kern.cam.boot_delay="30000"
debug.debugger_on_panic=1
debug.ddb.textdump.pending=1
hw.hptrr.attach_generic=0
vfs.mountroot.timeout="30"
ispfw_load="YES"
freenas_sysctl_load="YES"
hint.isp.0.role=2
hint.isp.1.role=2
hint.isp.2.role=2
hint.isp.3.role=2
hint.isp.0.topology="nport-only"
hint.isp.1.topology="nport-only"
hint.isp.2.topology="nport-only"
hint.isp.3.topology="nport-only"
module_path="/boot/kernel;/boot/modules;/usr/local/modules"
net.inet6.ip6.auto_linklocal="0"
vfs.zfs.vol.mode=2
kern.geom.label.disk_ident.enable="0"
hint.ahciem.0.disabled="1"
hint.ahciem.1.disabled="1"
kern.msgbufsize="524288"
hw.mfi.mrsas_enable="1"
hw.usb.no_shutdown_wait=1
hw.cxgbe.toecaps_allowed=0
hw.cxgbe.rdmacaps_allowed=0
hw.cxgbe.iscsicaps_allowed=0
vfs.nfsd.fha.write=0
vfs.nfsd.fha.max_nfsds_per_fh=32
```

Do not add or edit the default tunables. Changing the default tunables can make the system unusable.

The ZFS version used in 11.2 deprecates these tunables:

```
vfs.zfs.write_limit_override
vfs.zfs.write_limit_inflated
vfs.zfs.write_limit_max
vfs.zfs.write_limit_min
vfs.zfs.write_limit_shift
vfs.zfs.no_write_throttle
```

After upgrading from an earlier version of FreeNAS[®], these tunables are automatically deleted. Please do not manually add them back.

5.8 Update

FreeNAS[®] has an integrated update system to make it easy to keep up to date.

5.8.1 Preparing for Updates

It is best to perform updates at times the FreeNAS[®] system is idle, with no clients connected and no scrubs or other disk activity going on. Most updates require a system reboot. Plan updates around scheduled maintenance times to avoid disrupting user activities.

The update process will not proceed unless there is enough free space in the boot pool for the new update files. If a space warning is shown, use *Boot* (page 69) to remove unneeded boot environments.

5.8.2 Updates and Trains

FreeNAS[®] uses signed update files. This provides flexibility in deciding when to upgrade the system with patches, new drivers, or new features. It also allows "test driving" an upcoming release. Combined with boot environments, new features or system patches can be tested while maintaining the ability to revert to a previous version of the operating system, using the instructions in *If Something Goes Wrong* (page 33). Digitally signed update files eliminate the need to manually download both an upgrade file and the associated checksum to verify file integrity.

Figure 5.12 shows an example of the *System* \rightarrow *Update* screen.

System									
Information	General	Boot /	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts
🔽 Check	for Updates Da	aily and Dov	vnload if Ava	ailable					
Current Tra	ain: FreeNAS-:	11.2-STABLE	Ξ()				Manual Update		
Update Se	ver: http://up	date.ixsyste	ems.com/Fre	eeNAS					
Apply Per	iding Updates	Check No	Verify	Install		Fre	eNAS-11.2-STABLE	r	
Pending Up	dates								
Name									
	base-os-11.2-RC2-83a6344522ee9463bb43955cca021e24 -> base-os-11.2-RELEASE- d2f5bbb81785fcff78abaf7e4a32bb6e								
docs-11.2	docs-11.2-RELEASE-d2f5bbb81785fcff78abaf7e4a32bb6e								
	freebsd-pkgdb-11.2-RC2-83a6344522ee9463bb43955cca021e24 -> freebsd-pkgdb-11.2-RELEASE- d2f5bbb81785fcff78abaf7e4a32bb6e								
	kg-tools-11.2- L785fcff78aba			3bb43955	cca021e24 -> free	nas-pkg-too	ols-11.2-RELEASE-		
Train Des	criptions								

Fig. 5.12: Update Options

The system checks daily for updates and downloads an update if one is available. An alert is issued when a new update becomes available. The automatic check and download of updates can be disabled by unsetting *Check for Updates Daily and Download if Available*.

This screen lists the URL of the official update server in case that information is needed in a network with outbound firewall restrictions. It also shows which software branch, or *train*, is being tracked for updates.

Several trains are available for updates. Update trains are labeled with a numeric version and a short description.

The current version of FreeNAS[®] receives regular bug fixes and new features. Supported older versions of FreeNAS[®] only receive maintenance updates. Several specific words are used to describe the type of train:

- **STABLE:** Bug fixes and new features are available from this train. Upgrades available from a *STABLE* train are tested and ready to apply to a production environment.
- Nightlies: Experimental train used for testing future versions of FreeNAS[®].
- SDK: Software Developer Kit train. This has additional development tools for testing and debugging FreeNAS[®].

Warning: Only STABLE trains are recommended for regular usage. Other trains are made available for preproduction testing and updates to legacy versions. Pre-production testing trains are provided only to permit testing of new versions before switching to a new branch. Before using a non-production train, be prepared to experience bugs or problems. Testers are encouraged to submit bug reports at https://redmine.ixsystems.com/projects/freenas/issues.

The train selector does not allow downgrades. For example, a FreeNAS[®] system using a *Nightlies* upgrade train is not allowed to switch to a *STABLE* train. A version 9.10 train cannot be selected while booted in a version 11 boot environment. To go back to an earlier version after testing or running a more recent version of FreeNAS[®], reboot and select a *boot environment* (page 69) for that earlier version. *System* \rightarrow *Update* can then be used to check for updates from the related train.

The *Verify Install* button verifies that the operating system files in the current installation do not have any inconsistencies. If any problems are found, a pop-up menu lists the files with checksum mismatches or permission errors.

5.8.3 Checking for Updates

Check for updates by making sure the desired train is selected and clicking the *Check Now* button. Any available updates are listed. In the example shown in Figure 5.13, the numbers which begin with a # represent the issue number from the issue tracker (https://redmine.ixsystems.com/projects/freenas/issues). Numbers which do not begin with a # represent a git commit. Click the *ChangeLog* link to open the log of changes in a web browser. Click the *ReleaseNotes* link to open the Release Notes in the browser.

Check Now 🕺
ChangeLog
ReleaseNotes
The following packages will be downloaded:
Upgrade: base-os-9.10.1-U1-435aefc42da5265860019b42c921a40b -> base-os-9.10.1-ca82ba222c0be179a6983636c50732c3 Upgrade: docs-9.10.1-U1-435aefc42da5265860019b42c921a40b -> docs-9.10.1-ca82ba222c0be179a6983636c50732c3 Upgrade: freebsd-pkgdb-9.10.1-U1-435aefc42da5265860019b42c921a40b -> freebsd-pkgdb-9.10.1-ca82ba222c0be179a6983636c50732c3 Upgrade: freenas-pkg-tools-9.10.1-U1-435aefc42da5265860019b42c921a40b -> freenas-pkg-tools-9.10.1-ca82ba222c0be179a6983636c50732c3 Upgrade: FreeNASUI-9.10.1-U1-435aefc42da5265860019b42c921a40b -> FreeNASUI-9.10.1-ca82ba222c0be179a6983636c50732c3
🔀 Apply updates after downloading (The system will reboot)
Do you want to continue?
OK

Fig. 5.13: Reviewing Updates

5.8.4 Applying Updates

Make sure the system is in a low-usage state as described above in *Preparing for Updates* (page 80).

Click the *OK* button to immediately download and install an update. Be aware that some updates automatically reboot the system after they are applied.

Warning: Each update creates a boot environment. If the update process needs more space, it attempts to remove old boot environments. Boot environments marked with the *Keep* attribute as shown in *Boot* (page 69) will not be removed. If space for a new boot environment is not available, the upgrade fails. Space on the boot device can be manually freed using *System* \rightarrow *Boot*. Review the boot environments and remove the *Keep* attribute or delete any boot environments that are no longer needed.

During the update process a progress dialog appears. **Do not** interrupt the update until it completes.

Updates can also be downloaded and applied later. To do so, unset the *Apply updates after downloading* option before pressing *OK*. In this case, this screen closes after updates are downloaded. Downloaded updates are listed in the *Pending Updates* section of the screen shown in Figure 5.12. When ready to apply the previously downloaded updates, click the *Apply Pending Updates* button. Remember that the system reboots after the updates are applied.

Warning: After updates have completed, reboot the system. Configuration changes made after an update but before that final reboot will not be saved.

5.8.5 Manual Updates

Updates can be manually downloaded as a file ending with <code>-manual-update-unsigned.tar</code>. These updates are then applied with the *Manual Update* button. After obtaining the update file, click *Manual Update* and choose a location to temporarily store the file on the FreeNAS[®] system. Use the file browser to locate the update file, then click *Apply Update* to apply it.

There is also an option to back up the system configuration before updating. Click *Click here* and select any options to export in the configuration file. Click *OK* to open a popup window to save the system configuration. A progress dialog is displayed during the update. **Do not** interrupt the update.

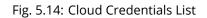
Tip: Manual updates cannot be used to upgrade from older major versions.

5.9 Cloud Credentials

FreeNAS[®] can use cloud services for features like *Cloud Sync* (page 95). The credentials to provide secure connections with cloud services are entered here. Amazon Cloud Drive, Amazon S3, Backblaze B2, Box, Dropbox, FTP, Google Cloud Storage, Google Drive, HTTP, Hubic, Mega, Microsoft Azure Blob Storage, Microsoft OneDrive, pCloud, SFTP, WebDAV, and Yandex are supported.

Select System \rightarrow Cloud Credentials to see the screen shown in Figure 5.14.

FreeNAS								()	🗸 syst	:ems`
Account System Tasks Network	Storage Directory	Sharing Services	Plugins Jails	VMs Repo) Support	Guide	О К
expand all collapse all Account	System Information General	Boot Advanced	Email System Datase	et Tunables Cloud	Credentials Update	Alerts Alert Services	CAs Certificates Support			
 ★ System ★ Jasks ★ Setwork 	Add Cloud Credential									
 Storage 	Account Name				Provider					
📧 🚺 Directory Service	Private Cloud Google Storage				HTTP GOOGLE_	DRIVE				_
 Sharing Image: Services 	Amazon S3 Storage				S3					
 Services Perjoins Jails VMs Guide Wizard Display System Processes Shell Log Out Reboot Shutdown 										



The list shows the *Account Name* and *Provider* for each credential. There are options to *Edit* and *Delete* a credential after selecting it. Click *Add Cloud Credential* to display the dialog shown in Figure 5.15.

Add Cloud Credential	8
Account Name:	
Provider:	Amazon Cloud Drive
Amazon Application Client ID	
Application Key	
OK Cancel	

Fig. 5.15: Adding Cloud Credentials

Amazon Cloud Drive options are shown by default. Enter a descriptive and unique name for the cloud credential in the *Account Name* field, then select a *Provider*. The remaining options vary by provider, and are shown in Table 5.6.

Provider	Setting	Description
Amazon Cloud Drive	Application Client ID, Application Key	Enter the Amazon application client ID and application key.
Amazon S3	Access Key, Secret Key	Enter the Amazon account access key and secret key.
Amazon S3	Endpoint URL	Enter the Endpoint URL for the web service.
Backblaze B2	Account ID or Appli- cation Key ID, Appli- cation Key	Enter the Account ID and Master Application Key (https://help.backblaze.com/hc/en-us/articles/224991568-Where-can- I-find-my-Account-ID-and-Application-Key-) for the Backblaze B2 account. These are visible after logging into the account, clicking <i>Buckets</i> , and clicking <i>Show Account ID and Application Key</i> . An <i>Application Key</i> with limited permis- sions can be used in place of the <i>Account ID</i> . Create a new Application Key, enter the key string in the <i>Application Key</i> field, and replace the <i>Account ID</i> with the <i>keyID</i> .
Box	Access Token	Enter the Box access token.
Dropbox	Access Token	Enter the Dropbox access token. The token is located on the App Console (https://www.dropbox.com/developers/apps). After creating an app, go to <i>Settings</i> and click <i>Generate</i> under the Generated access token field.
FTP	Host, Port	Enter the FTP host and port.
FTP	Username, Password	Enter the FTP username and password.
Google Cloud Storage	JSON Service Account Key	<i>Browse</i> to the location of the saved Google Cloud Storage key and select it.
Google Drive	Access Token, Team Drive ID	Enter the Google Drive Access Token. <i>Team Drive</i> <i>ID</i> is only used when connecting to a Team Drive (https://developers.google.com/drive/api/v3/reference/teamdrives). The ID is also the ID of the top level folder of the Team Drive.
HTTP	URL	Enter the URL.
Hubic	Access Token	Enter the access token.
Mega	Username, Password	Enter the Mega (https://mega.nz) username and password.
Microsoft Azure Blob Storage	Account Name, Ac- count Key	Enter the Azure Blob Storage account name and key.
Microsoft OneDrive	Access Token, Drive Account Type, Drive ID	Enter the access token. Choose the account type: <i>PERSONAL</i> , <i>BUSINESS</i> , or SharePoint (https://products.office.com/en-us/sharepoint/collaboration) <i>DOCUMENT_LIBRARY</i> . Enter the unique drive identifier. Open the <i>Shell</i> (page 289), enter rclone config, and follow the prompts to find these values. The rclone OneDrive documentation (https://rclone.org/onedrive/) guides through the configuration process.
pCloud	Access Token	Enter the access token.
SFTP	Host, Port	Enter the SFTP host and port.
SFTP	Username, Password, key file path	Enter the SFTP username, password, and PEM-encoded private key file path.
WebDAV	URL, WebDAV Service	Enter URL and use the dropdown to select the WebDAV service.
WebDAV	Username, Password	Enter the username and password.
Yandex	Access Token	Enter the access token.

Table 5.6:	Cloud	Credential	Options
10010 0.01	cioaa	creachtai	options

Additional fields are displayed after *Provider* is selected. For Amazon S3, *Access Key* and *Secret Key* are shown. These values are found on the Amazon AWS website by clicking on the account name, then *My Security Credentials* and *Access Keys (Access Key ID and Secret Access Key)*. Copy the Access Key value to the FreeNAS[®] Cloud Credential *Access Key* field, then enter the *Secret Key* value saved when the key pair was created. If the Secret Key value is unknown, a new key pair can be created on the same Amazon screen. The Google Cloud Storage *JSON Service Account Key* is found on the Google Cloud Platform Console (https://console.cloud.google.com/apis/credentials).

More details about individual Provider settings are available in the rclone documentation (https://rclone.org/about/).

5.10 Alerts

System \rightarrow Alerts displays the default notification frequency for each type of Alert (page 296). An example is seen in Figure 5.16.

nformation	General	Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts	Alert Service	CAs	Certificates	Support	
										-				
Settings:	ActiveDir	ectory did	not bind to the	e domain									Im	mediately
	collectd e	error											Im	mediately
	Encrypte	d volume f	ailed to rekey	some dis	s. Please make su	re you have	e working recovery k	eys, check	logs files ar	id correct the e	rror as it m	ay result to dat	ta loss. Im	mediately
	FreeNAS	HTTP serve	er SSL misconf	figuration									Im	mediately
	FreeNAS	Mini Critica	l IPMI Firmwar	e Update .	Available								Im	mediately
	IPs bound	d to iSCSI P	ortal were not	t found in	the system								Im	mediately
	LAGG int	erface erro	or										Im	mediately
	LDAP did	not bind to	o the domain										Im	mediately
	Multipath	is not opti	mal										Im	mediately
	NFS serv	ices could i	not bind specif	fic IPs, usi	ng wildcard								Im	mediately
	Replicatio	on failed											Im	mediately
	Samba e	rror											Im	mediately
	Scrub is	paused											Im	mediately
	Self-test	error											Im	mediately
	Service is	s not runnii	ng										Im	mediately
	SMART e	rror											Im	mediately
	smartd n	ot running											Im	mediately
	The boot	volume sta	ate is not HEAI	LTHY									Im	mediately
	The capa	city for the	e volume is ab	ove recon	nmended value								Im	mediately
	The volu	me status i	s not HEALTHY	,									Im	mediately
	The Web	GUI could n	not bind to spe	cified add	ress								Im	mediately
	There is	a new upda	ate available										Im	mediately
	Update fa	ailed. Chec	k /data/update	e.failed for	further details								Im	mediately
	Update n	ot applied											Im	mediately
	VMWare	failed to lo	g in to snapsh	ot									Im	mediately
	VMWare	snapshot d	elete failed										Im	mediately
	VMWare	snapshot fa	ailed										Im	mediately
	ZFS vers	ion is out o	of date										Im	mediately

Fig. 5.16: Configure Alert Notification Frequency

To change the notification frequency of an alert, click its drop-down menu and select IMMEDIATELY, HOURLY, DAILY, or NEVER.

Note: To configure where to send alerts, use Alert Services (page 85).

5.11 Alert Services

FreeNAS[®] can use a number of methods to notify the administrator of system events that require attention. These events are system *Alerts* (page 296) marked *WARN* or *CRITICAL*.

Currently available alert services:

- AWS-SNS (https://aws.amazon.com/sns/)
- Hipchat (https://www.atlassian.com/software/hipchat)
- InfluxDB (https://www.influxdata.com/)
- Slack (https://slack.com/)
- Mattermost (https://about.mattermost.com/)
- OpsGenie (https://www.opsgenie.com/)
- PagerDuty (https://www.pagerduty.com/)
- VictorOps (https://victorops.com/)

Warning: These alert services might use a third party commercial vendor not directly affiliated with iXsystems. Please investigate and fully understand that vendor's pricing policies and services before using their alert service. iXsystems is not responsible for any charges incurred from the use of third party vendors with the Alert Services feature.

Select System \rightarrow Alert Services to show the Alert Services screen. Click Add Service to display the dialog shown in Figure 5.17.

Add Alert Servic	e	ж
Service Name:	AWS-SNS	
Region:		ì
ARN:		ì
Key Id:		ì
Secret Key:		ì
Enabled:		
OK Cancel		

Fig. 5.17: Add Alert Service

The *Service Name* drop-down menu is used to pick a specific alert service. The fields shown in the rest of the dialog change to those required by that service. Enter the required information, set the *Enabled* option, then click *OK* to save the settings.

System alerts marked WARN or CRITICAL are sent to each alert service that has been configured and enabled.

Alert services are deleted from this list by clicking them and then clicking the *Delete* button at the bottom of the window. To disable an alert service temporarily, click *Edit* and remove the checkmark from the *Enabled* option.

Note: To send a test alert, highlight an alert entry, click Edit, and click the Send Test Alert button.

5.11.1 How it Works

A *nas-health* service is registered with Consul. This service runs /usr/local/etc/consul-checks/freenas_health.sh periodically, currently every two minutes. If an alert marked *WARNING* or *CRITICAL* is found, the *nas-health* service is marked as "unhealthy", triggering consul-alerts to notify configured alert services.

5.12 CAs

FreeNAS[®] can act as a Certificate Authority (CA). When encrypting SSL or TLS connections to the FreeNAS[®] system, either import an existing certificate, or create a CA on the FreeNAS[®] system, then create a certificate. This certificate will appear in the drop-down menus for services that support SSL or TLS.

For secure LDAP, the public key of an existing CA is imported with *Import CA*, or a new CA created on the FreeNAS[®] system and used on the LDAP server also.

Figure 5.18 shows the screen after clicking *System* \rightarrow *CAs*.

System													
Information	General	Boot Advanced	Email S	System Dataset	Tunables	Cloud Credentials	Update Al	erts A	lert Services	CAs	Certificates	Support	
Import CA	Create Interna	I CA Create Interme	ediate CA										
Name		Internal		Issuer		Certificates		Distingu	uished Name		From		Until
No entry has	been found												

Fig. 5.18: Initial CA Screen

If the organization already has a CA, the CA certificate and key can be imported. Click the *Import CA* button to open the configuration screen shown in Figure 5.19. The configurable options are summarized in Table 5.7.

Import CA		88
Identifier:	Internal identifier of the calphanumeric, "_" and "-"	ertificate. Only " are allowed.
Certificate:		٢
Private Key:		٢
Passphrase:	i	
Confirm Passphrase:		
Serial:	(1)	
OK Cancel		

Fig. 5.19: Importing a CA

Setting	Value	Description
Identifier	string	Enter a descriptive name for the CA using only alphanumeric, underscore
		(_), and dash (–) characters.
Certificate	string	Paste in the certificate for the CA.
Private Key	string	If there is a private key associated with the <i>Certificate</i> , paste it here.
Passphrase	string	If the <i>Private Key</i> is protected by a passphrase, enter it here and repeat it in
		the "Confirm Passphrase" field.
Serial	string	Enter the serial number for the certificate.

Table 5.7:	Importing	a CA	Options
------------	-----------	------	---------

To create a new CA, first decide if it will be the only CA which will sign certificates for internal use or if the CA will be part of a certificate chain (https://en.wikipedia.org/wiki/Root_certificate).

To create a CA for internal use only, click the Create Internal CA button which will open the screen shown in Figure 5.20.

Create Internal	CA		88
Identifier:		Internal identifier of the alphanumeric, "_" and "	certificate. Only -" are allowed.
Key length:	2048 💌		
Digest Algorith	m: SHA256 💌		
Lifetime:		3,650	
Country:	United States 💌 🕡		
State:		۱	
Locality:		(i)	
Organization:		(i)	
Email Address:		(i)	
Common Name	:	(i)	
Subject Alterna	ite Names:		i
OK Cancel			

Fig. 5.20: Creating an Internal CA

The configurable options are described in Table 5.8. When completing the fields for the certificate authority, supply the information for the organization.

Setting	Value	Description
Identifier	string	Enter a descriptive name for the CA using only alphanumeric, underscore
		(_), and dash (–) characters.
Key Length	drop-down menu	For security reasons, a minimum of 2048 is recommended.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different algo-
		rithm.
Lifetime	integer	The lifetime of the CA is specified in days.
Country	drop-down menu	Select the country for the organization.
State	string	Enter the state or province of the organization.
Locality	string	Enter the location of the organization.
Organization	string	Enter the name of the company or organization.
Email Address	string	Enter the email address for the person responsible for the CA.
Common Name	string	Enter the fully-qualified hostname (FQDN) of the system. The Common
		<i>Name</i> must be unique within a certificate chain.
Subject Alter-	string	Multi-domain support. Enter additional domain names and separate them
nate Names		with a space.

Table 5.8: Internal CA Options

To create an intermediate CA which is part of a certificate chain, click *Create Intermediate CA*. This screen adds one more option to the screen shown in Figure 5.20:

• **Signing Certificate Authority:** this drop-down menu is used to specify the root CA in the certificate chain. This CA must first be imported or created.

Imported or created CAs are added as entries in *System* \rightarrow *CAs*. The columns in this screen indicate the name of the CA, whether it is an internal CA, whether the issuer is self-signed, the number of certificates that have been issued by the CA, the distinguished name of the CA, the date and time the CA was created, and the date and time the CA expires.

Clicking the entry for a CA causes these buttons to become available:

- **Sign CSR:** used to sign internal Certificate Signing Requests created using *System* → *Certificates* → *Create Certificate Signing Request.*
- **Export Certificate:** prompts to browse to the location to save a copy of the CA X.509 certificate on the computer being used to access the FreeNAS[®] system.
- **Export Private Key:** prompts to browse to the location to save a copy of the CA private key on the computer being used to access the FreeNAS[®] system. This option only appears if the CA has a private key.
- Delete: prompts for confirmation before deleting the CA.

5.13 Certificates

FreeNAS[®] can import existing certificates, create new certificates, and issue certificate signing requests so that created certificates can be signed by the CA which was previously imported or created in *CAs* (page 87).

Figure 5.21 shows the initial screen after clicking System \rightarrow Certificates.

System												
Information Ge	neral Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts	Alert Services	CAs	Certificates	Support
Import Certificate	Create Internal	Certificate	Create Cer	tificate Signing Reqເ	iest							
Name		Issue	r			Distinguished Name			From			Until
No entry has been	found											



To import an existing certificate, click *Import Certificate* to open the configuration screen shown in Figure 5.22. When importing a certificate chain, paste the primary certificate, followed by any intermediate certificates, followed by the root CA certificate.

The configurable options are summarized in Table 5.9.

Import Certificate	×
Identifier:	certificate. Only " are allowed.
Certificate:	i
Private Key:	۲
Passphrase:	
Confirm Passphrase:	
OK Cancel	

Fig. 5.22: Importing a Certificate

Setting	Value	Description
Identifier	string	Enter a descriptive name for the certificate using only alphanumeric, under-
		score (_), and dash (–) characters.
Certificate	string	Paste the contents of the certificate.
Private Key	string	Paste the private key associated with the certificate.
Passphrase	string	If the private key is protected by a passphrase, enter it here and repeat it in
		the Confirm Passphrase field.

Table 5.9: Certificate Import Options

To create a new self-signed certificate, click the *Create Internal Certificate* button to see the screen shown in Figure 5.23. The configurable options are summarized in Table 5.10. When completing the fields for the certificate authority, use the information for the organization. Since this is a self-signed certificate, use the CA that was imported or created with *CAs* (page 87) as the signing authority.

Create Internal Certificate		_	Ж
Signing Certificate Authority:			
Identifier:		(i)	
Key length:	2048		
Digest Algorithm:	SHA256 -		
Lifetime:	3,650		
Country:	United States 👻 (1)		
State:		١	
Locality:		١	
Organization:		١	
Email Address:		١	
Common Name:		١	
Subject Alternate Names:			ì
OK Cancel			

Fig. 5.23: Creating a New Certificate

Setting	Value	Description
Signing Certificate	drop-down menu	Select the CA which was previously imported or created using CAs
Authority		(page 87).
Identifier	string	Enter a descriptive name for the certificate using only alphanumeric,
		underscore (_), and dash (–) characters.
Key Length	drop-down menu	For security reasons, a minimum of 2048 is recommended.
Digest Algorithm	drop-down menu	The default is acceptable unless the organization requires a different
		algorithm.
Lifetime	integer	The lifetime of the certificate is specified in days.
Country	drop-down menu	Select the country for the organization.
State	string	State or province for the organization.
Locality	string	Location of the organization.
Organization	string	Name of the company or organization.
Email Address	string	Email address for the person responsible for the CA.

Continued on next page

Setting	Value	Description
Common Name	string	Enter the fully-qualified hostname (FQDN) of the system. The Common
		<i>Name</i> must be unique within a certificate chain.
Subject Alternate	string	Multi-domain support. Enter additional domain names and separate
Names		them with a space.

Table 5.10 – continued from previous page

If the certificate is signed by an external CA, such as Verisign, instead create a certificate signing request. To do so, click *Create Certificate Signing Request*. A screen like the one in Figure 5.23 opens, but without the *Signing Certificate Authority* field.

Certificates that are imported, self-signed, or for which a certificate signing request is created are added as entries to System \rightarrow Certificates. In the example shown in Figure 5.24, a self-signed certificate and a certificate signing request have been created for the fictional organization *My Company*. The self-signed certificate was issued by the internal CA named *My Company* and the administrator has not yet sent the certificate signing request to Verisign so that it can be signed. Once that certificate is signed and returned by the external CA, it should be imported using *Import Certificate* so it is available as a configurable option for encrypting connections.

System						
Information General Boot Advar	nced Email System Dataset Tunables	Cloud Credentials Update Alerts	Alert Services CAs Certificates St	ipport		
Import Certificate Create Internal Certifica	Create Certificate Signing Request					
Name	Issuer	Distinguished Name	From	Until		
Name Distinguished Name From Until reeNAS_Internal_Certificate FreeNAS_Internal_CA /C=US/ST=CA/L=Silicon Valley/O=iXsystems /CN=realmini.tn.ixsystems /emailAddress=ix-docs@ixsystems.com Mon Jun 25 13:44:21 2018 Sat Jun 30 13:44:21 2018						

Fig. 5.24: Managing Certificates

Clicking an entry activates these configuration buttons:

- View: use this option to view the contents of an existing certificate or to edit the *Identifier*.
- **Export Certificate** saves a copy of the certificate or certificate signing request to the system being used to access the FreeNAS[®] system. For a certificate signing request, send the exported certificate to the external signing authority so that it can be signed.
- **Export Private Key** saves a copy of the private key associated with the certificate or certificate signing request to the system being used to access the FreeNAS[®] system.
- **Delete** is used to delete a certificate or certificate signing request.

5.14 Support

The FreeNAS[®] *Support* tab, shown in Figure 5.25, provides a built-in ticketing system for generating bug reports and feature requests.

System														
Information	General	Boot	Advanced	Email	System Dataset	Tunables	Cloud Credentials	Update	Alerts	Alert Services	CAs	Certificates	Support	
Before filin For enterp	ig a bug report rise-grade stora	or featu age solut	re request, sea tions and suppo	arch http://b ort, please v	ugs.freenas.org visit http://www.i:	to ensure t systems.co	he issue has not alre om/storage/.	ady been re	ported. If i	it has, add a comn	nent to th	e existing issue	e instead of	creating a new one.
lf you do n	ot have an acco	ount, ple	ase register.											
Username	Ð	[
Password	I	[
Туре		[Bug 🔻											
Category		[•											
Attach De	ebug Info													
Subject		[
Descripti	on													
Attachme	ents													
× -	Browse	No file	selected.											
Submit														

Fig. 5.25: Support Tab

This screen provides a built-in interface to the FreeNAS[®] issue tracker located at https://redmine.ixsystems.com/projects/ freenas/issues. When using the FreeNAS[®] bug tracker for the first time, go to the website, click the *Register* link, fill out the form, and reply to the registration email. This will create a username and password which can be used to create bug reports and receive notifications as the reports are actioned.

Before creating a bug report or feature request, ensure that an existing report does not already exist at https://redmine. ixsystems.com/projects/freenas/issues. If a similar issue is already present and has not been marked as *Closed* or *Resolved*, comment on that issue, adding new information to help solve it. If similar issues have already been *Closed* or *Resolved*, create a new issue and refer to the previous issue.

Note: Update the system to the latest version of STABLE and retest before reporting an issue. Newer versions of the software might have already fixed the problem.

To generate a report using the built-in *Support* screen, complete these fields:

- Username: enter the login name created when registering at https://redmine.ixsystems.com/projects/freenas/issues.
- · Password: enter the password associated with the registered login name.
- **Type:** select *Bug* when reporting an issue or *Feature* when requesting a new feature.
- **Category:** this drop-down menu is empty until a registered *Username* and *Password* are entered. An error message is displayed if either value is incorrect. After the *Username* and *Password* are validated, possible categories are populated to the drop-down menu. Select the one that best describes the bug or feature being reported.
- Attach Debug Info: enabling this option is recommended so an overview of the system hardware, build string, and configuration is automatically generated and included with the ticket. Generating and attaching a debug to the ticket can take some time. An error will occur if the debug is more than the file size limit of 20 MiB.

- Subject: enter a descriptive title for the ticket. A good Subject makes it easy to find similar reports.
- **Description:** enter a one- to three-paragraph summary of the issue that describes the problem, and if applicable, what steps can be taken to reproduce it.
- Attachments: this is the only optional field. It is useful for including configuration files or screenshots of any errors or tracebacks.

After completing the fields, click the *Submit* button to automatically generate and upload the report to https://redmine. ixsystems.com/projects/freenas/issues. A pop-up menu provides a clickable URL so to view status or add additional information to the report.

TASKS

The Tasks section of the administrative GUI is used to configure repetitive tasks:

- Cloud Sync (page 95) schedules data synchronization to cloud providers
- Cron Jobs (page 100) schedules a command or script to automatically execute at a specified time
- Init/Shutdown Scripts (page 102) configures a command or script to automatically execute during system startup or shutdown
- Rsync Tasks (page 103) schedules data synchronization to another system
- S.M.A.R.T. Tests (page 110) schedules disk tests

Each of these tasks is described in more detail in this section.

Note: By default, *Scrubs* (page 162) are run once a month by an automatically-created task. *S.M.A.R.T. Tests* (page 110) and *Periodic Snapshot Tasks* (page 150) must be set up manually.

6.1 Cloud Sync

Files or directories can be synchronized to remote cloud storage providers with the *Cloud Sync* feature.

Warning: This Cloud Sync task might go to a third party commercial vendor not directly affiliated with iXsystems. Please investigate and fully understand that vendor's pricing policies and services before creating any Cloud Sync task. iXsystems is not responsible for any charges incurred from the use of third party vendors with the Cloud Sync feature.

Cloud Credentials (page 82) must be pre-defined before a cloud sync is created. One set of credentials can be used for more than one cloud sync. For example, a single set of credentials for Amazon S3 can be used for separate cloud syncs that push different sets of files or directories.

A cloud storage area must also exist. With Amazon S3, these are called *buckets*. The bucket must be created before a sync task can be created.

After the credentials and receiving bucket have been configured, $Tasks \rightarrow Cloud Syncs \rightarrow Add Cloud Sync$ is used to define the schedule for running a cloud sync task. An example is shown in Figure 6.1.

Add Cloud Sync

Warning: This Cloud Sync task might go to a third party commercial vendor not directly affiliated with iXsystems. Please investigate and fully understand that vendor's pricing policies and services before creating any Cloud Sync task. iXsystems is not responsible for any charges incurred from the use of third party vendors with the Cloud Sync feature.

Description:	
Direction:	Push 🗸 i
Provider:	Credential 💌 Bucket Folder
Path:	Browse
Transfer Mode:	Sync
Remote encryption:	(i)
Auxiliary arguments:	i
Minute:	Every N minute Each selected minute 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54

Fig. 6.1: Adding a Cloud Sync

Table 6.1 shows the configuration options for Cloud Syncs.

Setting	Value Type	Description
Description	string	Enter a descriptive name for this Cloud Sync.
Direction	string	<i>Push</i> sends data to cloud storage. <i>Pull</i> receives data from cloud storage.
Provider	drop-down menu	Choose the cloud storage provider credentials from the list of entered <i>Cloud Credentials</i> (page 82). The UI tests the credential and displays an error if a connection cannot be made.
Amazon S3	drop-down menu	Only appears when an S3 credential is the <i>Provider</i> . Select the pre-defined
Buckets		S3 bucket to use.
Folder	string	Only appears when an S3 credential is the <i>Provider</i> . Optionally enter the name of the folder within the selected bucket.
Server Side En-	drop-down menu	Only appears when an S3 credential is the <i>Provider</i> . Choices are <i>None</i> (no
cryption		encryption) or AES-256 (encrypted).
Path	browse button	Select the directories or files to be sent to the cloud for <i>Push</i> syncs, or the destination to be written as the destinations for <i>Pull</i> syncs. Be cautious about the destination of <i>Pull</i> jobs to avoid overwriting existing files.
Transfer Mode	drop-down menu	Sync (default) makes files on destination system identical to those on the source. Files removed from the source are also removed from the destination, similar to rsyncdelete. Copy copies files from the source to the destination and skips files that are identical, similar to rsync. Move copies files from the source to the destination and deletes the source files after the copy, similar to mv.
Remote encryp-	checkbox	Set to encrypt files before transfer and store the encrypted files on the
tion		remote system. rclone Crypt (https://rclone.org/crypt/) is used.
Filename en- cryption	checkbox	Only appears when <i>Remote encryption</i> is enabled. Set to encrypt the shared file names.
Encryption pass- word	string	Only appears when <i>Remote encryption</i> is enabled. Enter the password for encrypting and decrypting remote data. <i>Warning</i> : Always save and back up this password. Losing the encryption password can result in data loss.
Encryption salt	string	Only appears when <i>Remote encryption</i> is enabled. En- ter a long string of random characters for use as salt (https://searchsecurity.techtarget.com/definition/salt) for the encryp- tion password. <i>Warning</i> : Save and back up the encryption salt value. Losing the salt value can result in data loss.
Minute	slider or minute se- lections	Select <i>Every N minutes</i> and use the slider to choose a value, or select <i>Each selected minute</i> and choose specific minutes to run the task.
Hour	slider or hour selec- tions	Select <i>Every N hours</i> and use the slider to choose a value, or select <i>Each selected hour</i> and choose specific hours to run the task.
Day of month	slider or day of month selections	Select <i>Every N days of month</i> and use the slider to choose a value, or select <i>Each selected day of month</i> and choose specific days to run the task.
Month	checkboxes	Months when the task runs.
Day of week	checkboxes	Days of the week to run the task.
Enabled	checkbox	Unset to temporarily disable this Cloud Sync.
Lindbied	Checkbox	

Table 6.1: Cloud	Sync	Options
------------------	------	---------

Figure 6.2 shows a cloud sync called *backup-acctg* that "pushes" a file to cloud storage. The last run finished with a status of *SUCCESS*.

d Cloud Sync										
scription	Direction	Path	Status	Minute	Hour	Day of month	Month	Day of week	Credential	Enabled
ackup-acctg	PUSH	/mnt/volume1 /smb-storage /accounting- backup.bin	SUCCESS	0	Every hour	Every day	Every month	Every day of week	S3 Storage	true

Fig. 6.2: Cloud Sync Status

To modify an existing cloud sync, click the entry to access the *Edit*, and *Delete*, and *Run Now* buttons.

Click the *Status* column entry for a cloud sync that is *RUNNING*, *FAILED*, or a *SUCCESS*. This opens the log in a pop-up window to read any error messages or other details.

6.1.1 Cloud Sync Example

This example shows a *Push* cloud sync which writes an accounting department backup file from the FreeNAS[®] system to Amazon S3 storage.

Before the new cloud sync was added, a bucket called *cloudsync-bucket* was created with the Amazon S3 web console for storing data from the FreeNAS[®] system.

System \rightarrow Cloud Credentials \rightarrow Add Cloud Credential is used to enter the credentials for storage on an Amazon AWS account. The credential is given the name S3 Storage, as shown in Figure 6.3:

Add	Cloud Crede	ntial	Ж
Ad	count Name:	S3 Storage	
Pr	ovider:	Amazon S3	
Ac	cess Key:	XYZZYXSQUAWKASQUEEPS	
Se	cret Key:	VutRWwPQEos+TEtQEWE5si	
0	KCancel		

Fig. 6.3: Example: Adding Cloud Credentials

The local data to be sent to the cloud is a single file called accounting-backup.bin on the smb-storage dataset. A cloud sync job is created with $Tasks \rightarrow Cloud Sync \rightarrow Add Cloud Sync$. The *Description* is set to *backup-acctg* to describe the job. This data is being sent to cloud storage, so this is a *Push*. The provider comes from the cloud credentials defined in the previous step, and the destination bucket *cloudsync-bucket* is selected.

The *Path* to the data file is selected.

The remaining fields are for setting a schedule. The default is to send the data to cloud storage once an hour, every day. The options provide great versatility in configuring when a cloud sync runs, anywhere from once a minute to once a year.

The *Enabled* option is set by default, so this cloud sync will run at the next scheduled time.

The completed dialog is shown in Figure 6.4:

Add Cloud Sync	8		
		Hour:	Every N hour Each selected hour
Description:	backup-acctg		(,
Direction:	Push v d		1
Provider:	Credential S3 Storage v Amazon S3 Buckets cloudsync-bucket v Folder		<i>(</i>)
Path:	/mnt/volume1/smb-storage/a	Day of month:	Every N day of month Each selected day of month
Palli	intervention and age in the ag		1
			(i)
Minute:	Every N minute Each selected minute 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59	Month: Day of week:	 January February March April June July July August September October November December December Wonday Tuesday Wednesday Thursday
		Enabled: OK Cancel	 C Thursday Friday Saturday Sunday

Fig. 6.4: Example: Adding a Cloud Sync

6.2 Cron Jobs

cron(8) (https://www.freebsd.org/cgi/man.cgi?query=cron) is a daemon that runs a command or script on a regular schedule as a specified user.

Figure 6.5 shows the screen that opens after clicking $Tasks \rightarrow Cron Jobs \rightarrow Add Cron Job$.

Add Cron Job	88
User:	The user to run the comman
Command:	
Short description:	
Minute: Hour:	Every N minute Each selected minute 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 Image: The selected hour Image: The selec
Day of month:	Every N day of month Each selected day of month T 1
Month:	• 🔽 January

Fig. 6.5: Creating a Cron Job

Table 6.2 lists the configurable options for a cron job.

Value	Description	
drop-down menu	Choose a user account to run the command or script. The user must have	
	permissions to run the command.	
string	Enter the full path to the command or script to be run. Test a script at the	
	command line first to make sure it works as expected.	
string	Optional. Describe the new cron job.	
slider or minute se-	With the slider, the cron job occurs every N minutes. With minute selec-	
lections	tions, the cron job occurs at the highlighted minutes	
slider or hour selec-	With the slider, the cron job occurs every N hours. With hour selections,	
tions	the cron job occurs at the highlighted hours.	
slider or month selec-	With the slider, the cron job occurs every N days. With day selections, the	
tions	cron job occurs on the highlighted days each month.	
checkboxes	Cron job occurs on the selected months.	
checkboxes	Cron job occurs on the selected days.	
checkbox	Disables emailing standard output to the <i>root</i> user account.	
checkbox	Disables emailing errors to the <i>root</i> user account.	
checkbox	Deselect disable the cron job without deleting it.	
	drop-down menu string string slider or minute se- lections slider or hour selec- tions slider or month selec- tions checkboxes checkboxes checkbox checkbox	

Table 6.2: Cron Job Options

Cron jobs are shown in *View Cron Jobs*. Highlight a cron job entry to display buttons to *Edit*, *Delete*, or *Run Now*.

Note: % symbols are automatically escaped and should not be prefixed with backslashes. For example, use date '+%Y-%m-%d' in a cron job to generate a filename based on the date.

6.3 Init/Shutdown Scripts

FreeNAS[®] provides the ability to schedule commands or scripts to run at system startup or shutdown.

Figure 6.6 shows the screen that opens after clicking *Tasks* \rightarrow *Init/Shutdown Scripts* \rightarrow *Add Init/Shutdown Script*. Table 6.3 summarizes the options.

Scheduled commands must be in the default path. The full path to the command can also be included in the entry. The path can be tested by typing which commandname. If the command is not found, it is not in the path.

When scheduling a script, make sure that the script is executable and has been fully tested to ensure it achieves the desired results.

Add Init/Shutdown Script 🛛 🕅	
Туре:	Command
Command:	
When:	💌
Enabled:	
OK Cancel	

Fig. 6.6: Add an Init/Shutdown Script

Table 6.3: Options When Adding an Init/Shi	utdown Script
--	---------------

Setting	Value	Description
Туре	drop-down menu	Select <i>Command</i> for an executable or <i>Script</i> for an executable script.
Command	string	If <i>Command</i> is selected, enter the command plus any desired options. If <i>Script</i> is selected, <i>Browse</i> to the location of the script.
When	drop-down menu	Select when the command or script runs. <i>Pre Init</i> is very early in boot pro- cess before mounting filesystems, <i>Post Init</i> is towards end of boot process before FreeNAS services start, or at <i>Shutdown</i> .
Enabled	checkbox	Unset to disable the task.

6.4 Rsync Tasks

Rsync (https://www.samba.org/ftp/rsync/rsync.html) is a utility that copies specified data from one system to another over a network. Once the initial data is copied, rsync reduces the amount of data sent over the network by sending only the differences between the source and destination files. Rsync is used for backups, mirroring data on multiple systems, or for copying files between systems.

Rsync is most effective when only a relatively small amount of the data has changed. There are also some limitations when using Rsync with Windows files (https://forums.freenas.org/index.php?threads/impaired-rsync-permissions-support-for-windows-datasets.43973/). For large amounts of data, data that has many changes from the previous copy, or Windows files, *Replication Tasks* (page 152) are often the faster and better solution.

Rsync is single-threaded and gains little from multiple processor cores. To see whether rsync is currently running, use pgrep rsync from the *Shell* (page 289).

Both ends of an rsync connection must be configured:

- the rsync server: this system pulls (receives) the data. This system is referred to as PULL in the configuration examples.
- the rsync client: this system pushes (sends) the data. This system is referred to as PUSH in the configuration examples.

FreeNAS[®] can be configured as either an *rsync client* or an *rsync server*. The opposite end of the connection can be another FreeNAS[®] system or any other system running rsync. In FreeNAS[®] terminology, an *rsync task* defines which data is synchronized between the two systems. To synchronize data between two FreeNAS[®] systems, create the *rsync task* on the *rsync client*.

FreeNAS[®] supports two modes of rsync operation:

- rsync module mode: exports a directory tree, and the configured settings of the tree as a symbolic name over an unencrypted connection. This mode requires that at least one module be defined on the rsync server. It can be defined in the FreeNAS[®] GUI under *Services* → *Rsync* → *Rsync* Modules. In other operating systems, the module is defined in rsyncd.conf(5) (https://www.samba.org/ftp/rsync/rsyncd.conf.html).
- **rsync over SSH:** synchronizes over an encrypted connection. Requires the configuration of SSH user and host public keys.

This section summarizes the options when creating an rsync task. It then provides a configuration example between two FreeNAS[®] systems for each mode of rsync operation.

Note: If there is a firewall between the two systems or if the other system has a built-in firewall, make sure that TCP port 873 is allowed.

Figure 6.7 shows the screen that appears after selecting *Tasks* \rightarrow *Rsync Tasks* \rightarrow *Add Rsync Task*. Table 6.4 summarizes the options that can be configured when creating an rsync task.

Add Rsync Task	×
Path:	Browse
User:	- (i)
Remote Host:	(i)
Rsync mode:	Rsync module
Remote Module Name:	(d)
Direction:	Push 💌 🚺
Short description:	
Minute:	Every N minute Each selected minute
	00 01 02 03 04 05 06 07 08 09
	10 11 12 13 14 15 16 17 18 19
	20 21 22 23 24 25 26 27 28 29
	30 31 32 33 34 35 36 37 38 39
	40 41 42 43 44 45 46 47 48 49
	50 51 52 53 54 55 56 57 58 59
Hour:	Every N hour Each selected hour
Day of month:	Every N day of month Each selected day of month
	I
	1

Setting	Value	Description
Path	browse button	<i>Browse</i> to the path to be copied. Path lengths cannot be greater than 255
		characters.
User	drop-down menu	The chosen user must have write permissions for the specified remote
		directory. The user name cannot contain spaces or exceed 17 characters.
Remote Host	string	Enter the IP address or hostname of the remote system that will store the
		copy. Use the format <i>username@remote_host</i> if the username differs on
		the remote host.
Remote SSH	integer	Only available in <i>Rsync over SSH</i> mode. Allows specifying an SSH port other
Port		than the default of 22.
Rsync mode	drop-down menu	Choices are <i>Rsync module</i> or <i>Rsync over SSH</i> .
Remote Module	string	At least one module must be defined in rsyncd.conf(5)
Name	Stille	(https://www.samba.org/ftp/rsync/rsyncd.conf.html) of the rsync
Name		server or in the <i>Rsync Modules</i> of another system.
Remote Path	string	
Remote Path	string	Only appears when using <i>Rsync over SSH</i> mode. Enter the existing path on
		the remote host to sync with. Example: <i>/mnt/volume</i> . Note that maximum
		path length is 255 characters.
Validate Remote	checkbox	Verifies the existence of the <i>Remote Path</i> .
Path		
Direction	drop-down menu	Direct the flow of the data to the remote host. Choices are <i>Push</i> or <i>Pull</i> .
		Default is to <i>Push</i> to a remote host.
Short Descrip-	string	Enter an optional description of the new rsync task.
tion		
Minute	slider or minute se-	When the slider is used the sync occurs every N minutes. Use Each selected
	lections	<i>minute</i> for the sync to occur at the highlighted minutes.
Hour	slider or hour selec-	When the slider is used the sync occurs every N hours. Use <i>Each selected</i>
	tions	<i>hour</i> for the sync to occur at the highlighted hours.
Day of month	slider or day selec-	When the slider is used the sync occurs every N days. Use <i>Each selected</i>
-)	tions	<i>day of the month</i> for the sync to occur on the highlighted days.
Month	checkboxes	Define which months to run the task.
Day of week	checkboxes	Define which days of the week to run the task.
Recursive	checkbox	Set to include all subdirectories of the specified volume during the rsync
Recursive		task.
Times	checkbox	Set to preserve the modification times of the files.
	checkbox	Set to preserve the modification times of the mes. Set to reduce the size of data to transmit. Recommended for slower con-
Compress	CHECKDOX	
A		nections.
Archive	checkbox	Equivalent to -rlptgoD. This will run the task as recursive, copy symlinks
		as symlinks, preserve permissions, preserve modification times, preserve
		group, preserve owner (root only), and preserve device and special files.
Delete	checkbox	Set to delete files in the destination directory that do not exist in the send-
		ing directory.
Quiet	checkbox	Set to suppresses informational messages from the remote server.
Preserve per- missions	checkbox	Set to preserve original file permissions. Useful if User is set to <i>root</i> .
Preserve ex-	checkbox	Both systems must support extended attributes.
tended at-		(https://en.wikipedia.org/wiki/Xattr).
tributes		
Delay Updates	checkbox	Set to save the temporary file from each updated file to a holding direc-
Delay opuales		tory. At the end of the transfer, all transferred files are renamed into place
Futra cationa	string	and temporary files deleted.
Extra options	string	Add any other rsync(1) (http://rsync.samba.org/ftp/rsync/rsync.html) op-
		tions. The $*$ character must be escaped with a backslash ($\ \pm \pm$) or used
		inside single quotes ('*.txt').

Table 6.4: Rsync	Configuration	Options
------------------	---------------	---------

Continued on next page

Table 6.4 – continued from previous page			
Setting	Value	Description	
Enabled	checkbox	Unset to disable the rsync task without deleting it.	

If the rysnc server requires password authentication, enter --password-file=/PATHTO/FILENAME in the *Extra options* option, replacing /PATHTO/FILENAME with the appropriate path to the file containing the password.

Created rsync tasks will be listed in *View Rsync Tasks*. Highlight the entry for an rsync task to display buttons for *Edit*, *Delete*, or *Run Now*.

6.4.1 Rsync Module Mode

This configuration example configures rsync module mode between these two FreeNAS[®] systems:

- 192.168.2.2 has existing data in /mnt/local/images. It will be the rsync client, meaning that an rsync task needs to be defined. It will be referred to as PUSH.
- 192.168.2.6 has an existing volume named /mnt/remote. It will be the rsync server, meaning that it will receive the contents of /mnt/local/images. An rsync module needs to be defined on this system and the rsyncd service needs to be started. It will be referred to as *PULL*.

On *PUSH*, an rsync task is defined in *Tasks* \rightarrow *Rsync Tasks* \rightarrow *Add Rsync Task*. In this example:

- the Path points to /usr/local/images, the directory to be copied
- the Remote Host points to 192.168.2.6, the IP address of the rsync server
- the *Rsync Mode* is *Rsync module*
- the Remote Module Name is backups; this will need to be defined on the rsync server
- the Direction is Push
- the rsync is scheduled to occur every 15 minutes
- the User is set to root so it has permission to write anywhere
- the Preserve Permissions option is enabled so that the original permissions are not overwritten by the root user

On *PULL*, an rsync module is defined in *Services* \rightarrow *Rsync Modules* \rightarrow *Add Rsync Module*. In this example:

- the Module Name is backups; this needs to match the setting on the rsync client
- the Path is /mnt/remote; a directory called images will be created to hold the contents of /usr/local/images
- the User is set to root so it has permission to write anywhere
- Hosts allow is set to 192.168.2.2, the IP address of the rsync client

Descriptions of the configurable options can be found in Rsync Modules.

To finish the configuration, start the rsync service on *PULL* in *Services* \rightarrow *Control Services*. If the rsync is successful, the contents of /mnt/local/images/ will be mirrored to /mnt/remote/images/.

6.4.2 Rsync over SSH Mode

SSH replication mode does not require the creation of an rsync module or for the rsync service to be running on the rsync server. It does require SSH to be configured before creating the rsync task:

- a public/private key pair for the rsync user account (typically *root*) must be generated on *PUSH* and the public key copied to the same user account on *PULL*
- to mitigate the risk of man-in-the-middle attacks, the public host key of PULL must be copied to PUSH
- the SSH service must be running on PULL

To create the public/private key pair for the rsync user account, open *Shell* (page 289) on *PUSH* and run ssh-keygen. This example generates an RSA type public/private key pair for the *root* user. When creating the key pair, do not enter the passphrase as the key is meant to be used for an automated task.

ssh-keygen -t rsa Generating public/private rsa key pair. Enter file in which to save the key (/root/.ssh/id_rsa): Created directory '/root/.ssh'. Enter passphrase (empty for no passphrase): Enter same passphrase again: Your identification has been saved in /root/.ssh/id_rsa. Your public key has been saved in /root/.ssh/id_rsa.pub. The key fingerprint is: f5:b0:06:d1:33:e4:95:cf:04:aa:bb:6e:a4:b7:2b:df root@freenas.local The key's randomart image is: +--[RSA 2048]----+ .0. 00 0+0. . 1 L . =0 + + + 0 1 so. .0 ο. 0 00 **0E ____|

FreeNAS[®] supports RSA keys for SSH. When creating the key, use -t rsa to specify this type of key. Refer to Key-based Authentication (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/openssh.html#security-ssh-keygen) for more information.

Note: If a different user account is used for the rsync task, use the su - command after mounting the filesystem but before generating the key. For example, if the rsync task is configured to use the *user1* user account, use this command to become that user:

su – user1

Next, view and copy the contents of the generated public key:

more .ssh/id_rsa.pub ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQC11BEXRgw1W8y8k+1XP1VR3xsmVSjtsoyIzV/PlQPo SrWotUQzqILq0SmUpViAAv4Ik3T8NtxXyohKmFNbBczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4h dcD7Y5mvU3MAEeDC1t02/xoi5xS/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmasti00qmDDcp/k0 xT+S6DFNDBy6IYQN4heqmhTPRXqPhXqcD1G+rWr/nZK4H8Ckzy+19RaEXMRuTyQgqJB/rsRcmJX5fApd DmNfwrRSxLjDvUzfywnjFH1Kk/+TQIT1gg1QQaj21PJD9pnDVF0AiJrWyWnR root@freenas.local

Go to *PULL* and paste (or append) the copied key into the *SSH Public Key* field of *Account* \rightarrow *Users* \rightarrow *View Users* \rightarrow *root* \rightarrow *Modify User*, or the username of the specified rsync user account. The paste for the above example is shown in Figure 6.8. When pasting the key, ensure that it is pasted as one long line and, if necessary, remove any extra spaces representing line breaks.

Account			
Groups	ers	Microsoft Account:	
Add User		SSH Public Key:	ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQC1lBEXRgw1 W8y8k+lXPlVR3xsmVSjtsoyIzV/PlQPo SrWotUQzqILq0SmUpViAAv4Ik3T8NtxXyohKmFNb
User ID	Userr		BczU6tEsVGHo/2BLjvKiSHRPHc/1DX9hofcFti4h dcD7Y5mvU3MAEeDClt02/xoi5xS
0	root		/RLxgP0R5dNrakw958Yn001sJS9VMf528fknUmas ti00gmDDcp/k0
1	daem		xT+S6DFNDBy6IYQN4heqmhTPRXqPhXqcD1G+ <u>rWr</u> /nZK4H8Ckzy+l9RaEXMRuTyQgqJB /rSRcmJX5fApd
2	opera		DmNfwrRSxLiDvUzfywnjFHlKk
3	bin		/+TQIT1gg1QQaj21PJD9pnDVF0AiJrWyWnR root@freenas.local
		line a Birraham	Owner Group Other
4	tty	Home Directory	Read 🗹 🗹
5	kmer	Mode:	Write
7	game		Execute 🖌 🖌
8	news	Auxiliary groups:	Available Selected
9	man		_uncp
14	ftp		_pflogd >>
22	sshd		authpf
25	smm		avahi bin -
26	mailn	OK Cancel	
Modify User Ch	ange E-	Cancer	

Fig. 6.8: Pasting the User SSH Public Key

While on *PULL*, verify that the SSH service is running in *Services* \rightarrow *Control Services* and start it if it is not.

Next, copy the host key of *PULL* using Shell on *PUSH*. The command below copies the RSA host key of the *PULL* server used in our previous example. Be sure to include the double bracket >> to prevent overwriting any existing entries in the known_hosts file:

ssh-keyscan -t rsa 192.168.2.6 >> /root/.ssh/known_hosts

Note: If *PUSH* is a Linux system, use this command to copy the RSA key to the Linux system:

cat ~/.ssh/id_rsa.pub | ssh user@192.168.2.6 'cat >> .ssh/authorized_keys'

The rsync task can now be created on *PUSH*. To configure rsync SSH mode using the systems in the previous example, use this configuration:

- the Path points to /mnt/local/images, the directory to be copied
- the Remote Host points to 192.168.2.6, the IP address of the rsync server
- the Rsync Mode is Rsync over SSH
- the rsync is scheduled to occur every 15 minutes
- the User is set to root so it has permission to write anywhere; the public key for this user must be generated on PUSH and copied to PULL
- the Preserve Permissions option is enabled so that the original permissions are not overwritten by the root user

Save the rsync task and the rsync will automatically occur according to the schedule. In this example, the contents of /mnt/local/images/ will automatically appear in /mnt/remote/images/ after 15 minutes. If the content does not appear, use

Shell on *PULL* to read /var/log/messages. If the message indicates a *n* (newline character) in the key, remove the space in the pasted key-it will be after the character that appears just before the *n* in the error message.

6.5 S.M.A.R.T. Tests

S.M.A.R.T. (https://en.wikipedia.org/wiki/S.M.A.R.T.) (Self-Monitoring, Analysis and Reporting Technology) is a monitoring system for computer hard disk drives to detect and report on various indicators of reliability. Replace the drive when a failure is anticipated by S.M.A.R.T. Most modern ATA, IDE, and SCSI-3 hard drives support S.M.A.R.T. – refer to the drive documentation for confirmation.

Figure 6.9 shows the configuration screen that appears after selecting *Tasks* \rightarrow *S.M.A.R.T. Tests* \rightarrow *Add S.M.A.R.T. Test.* Tests are listed under *View S.M.A.R.T. Tests.* After creating tests, check the configuration in *Services* \rightarrow *S.M.A.R.T.*, then click the slider to *ON* for the S.M.A.R.T. service in *Services* \rightarrow *Control Services.* The S.M.A.R.T. service will not start if there are no volumes.

Note: To prevent problems, do not enable the S.M.A.R.T. service if the disks are controlled by a RAID controller. It is the job of the controller to monitor S.M.A.R.T. and mark drives as Predictive Failure when they trip.

Add S.M.A.R.T. Tes	it 🛛 🕺
Disks:	ada0 ada1 ada2 ada3
Туре:	
Short description:	
Hour:	Every N hour Each selected hour
	1
	<i>i</i>
Day of month:	Every N day of month Each selected day of month
	1
	(i)
Month:	• 💟 January • 💟 February
	• 🗹 March • 💟 April
	• 🔽 May • 🔽 June
	• 🔽 July
	• 📝 August • 📝 September
	• 🔽 October • 🔽 November
	• 🔽 December

Fig. 6.9: Adding a S.M.A.R.T. Test

Table 6.5 summarizes the configurable options when creating a S.M.A.R.T. test.

Setting	Value	Description
Disks	list	Select the disks to monitor.
Туре	drop-down menu	Choose the test type. See smartctl(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.ir for descriptions of each type of test. Some test types will degrade performance or take disks offline. Avoid scheduling S.M.A.R.T. tests simultaneously with scrub or resilver operations.
Short descrip- tion	string	Optional. Enter a short description of this test.
Hour	slider or hour selec- tions	When the slider is used the sync occurs every N hours. Use <i>Each selected hour</i> for the test to occur at the highlighted hours.
Day of month	slider or day selec- tions	When the slider is used the sync occurs every N days. Use <i>Each selected day of the month</i> for the sync to occur on the highlighted days.
Month	checkboxes	Select which months to run the test.
Day of week	checkboxes	Select which days of the week to run the test.

Table 6.5:	S.M.A.R.T.	Test Options
------------	------------	--------------

Note: Scrub tasks are run if and only if the threshold is met or exceeded *and* the task is scheduled to run on the date marked.

An example configuration is to schedule a *Short Self-Test* once a week and a *Long Self-Test* once a month. These tests do not have a performance impact, as the disks prioritize normal I/O over the tests. If a disk fails a test, even if the overall status is *Passed*, start to think about replacing that disk.

Warning: Some S.M.A.R.T. tests cause heavy disk activity and can drastically reduce disk performance. Do not schedule S.M.A.R.T. tests to run at the same time as scrub or resilver operations or during other periods of intense disk activity.

Which tests will run and when can be verified by typing smartd -q showtests within Shell (page 289).

The results of a test can be checked from *Shell* (page 289) by specifying the name of the drive. For example, to see the results for disk *ada0*, type:

smartctl -l selftest /dev/ada0

If an email address is entered in the *Email to report* field of *Services* \rightarrow *S.M.A.R.T.*, the system will send an email to that address when a test fails. Logging information for S.M.A.R.T. tests can be found in /var/log/daemon.log.

NETWORK

The Network section of the administrative GUI contains these components for viewing and configuring network settings on the FreeNAS[®] system:

- Global Configuration (page 113): general network settings.
- Interfaces (page 115): settings for each network interface.
- *IPMI* (page 117): settings controlling connection to the appliance through the hardware side-band management interface if the graphical user interface becomes unavailable.
- Link Aggregations (page 119): settings for network link aggregation and link failover.
- Network Summary (page 123): display an overview of the current network settings.
- Static Routes (page 123): add static routes.
- VLANs (page 123): configure IEEE 802.1q tagging for virtual LANs.

Each of these is described in more detail in this section.

Warning: Making changes to the network interface the web interface uses can result in losing connection to the FreeNAS[®] system! Misconfiguring network settings might require command line knowledge or physical access to the FreeNAS[®] system to fix. Be very careful when configuring *Interfaces* (page 115) and *Link Aggregations* (page 119).

7.1 Global Configuration

Network \rightarrow *Global Configuration*, shown in Figure 7.1, is for general network settings that are not unique to any particular network interface.

lobal Configuration Interfa	ces Link Aggregations Network Summary Static Routes	VLAN
Hostname:	freenas	
Domain:	local	
Additional domains:		i
IPv4 Default Gateway:		
IPv6 Default Gateway:		
Nameserver 1:		
Nameserver 2:		
Nameserver 3:		
HTTP Proxy:		
Enable netwait feature:		
Netwait IP list:	(i)	
Host name data base:		i
Save		

Fig. 7.1: Global Network Configuration

Table 7.1 summarizes the settings on the Global Configuration tab. *Hostname* and *Domain* fields are pre-filled as shown in Figure 7.1, but can be changed to meet requirements of the local network.

SettingValueDescriptionHostnamestringSystem host name. Cannot contain the underscore character.DomainstringSystem domain name.Additional do- mainsstringCan enter up to 6 space delimited search domains. Adding multiple do- mains may result in slower DNS lookups.IPv4 DefaultIP addressTypically not set. See <i>this note about Gateways</i> (page 115). If set, used in- stead of default gateway provided by DHCP.			6 6
DomainstringSystem domain name.Additional do- mainsstringCan enter up to 6 space delimited search domains. Adding multiple do- mains may result in slower DNS lookups.IPv4 DefaultIP addressTypically not set. See this note about Gateways (page 115). If set, used in-	Setting	Value	Description
Additional do- mainsstringCan enter up to 6 space delimited search domains. Adding multiple do- mains may result in slower DNS lookups.IPv4 DefaultIP addressTypically not set. See this note about Gateways (page 115). If set, used in-	Hostname	string	System host name. Cannot contain the underscore character.
mainsmains may result in slower DNS lookups.IPv4 DefaultIP addressTypically not set. See this note about Gateways (page 115). If set, used in-	Domain	string	System domain name.
IPv4 DefaultIP addressTypically not set. See this note about Gateways (page 115). If set, used in-	Additional do-	string	Can enter up to 6 space delimited search domains. Adding multiple do-
	mains		mains may result in slower DNS lookups.
Gateway stead of default gateway provided by DHCP.	IPv4 Default	IP address	Typically not set. See this note about Gateways (page 115). If set, used in-
	Gateway		stead of default gateway provided by DHCP.

Continued on next page

Setting	Value	Description
IPv6 Default	IP address	Typically not set. See <i>this note about Gateways</i> (page 115).
Gateway		
Nameserver 1	IP address	Primary DNS server.
Nameserver 2	IP address	Secondary DNS server.
Nameserver 3	IP address	Tertiary DNS server.
HTTP Proxy	string	Enter the proxy information for the network in the format
		http://my.proxy.server:3128 or http://user:password@my.proxy.server:3128.
Enable netwait	checkbox	If enabled, network services do not start at boot until the interface is able
feature		to ping the addresses listed in the <i>Netwait IP list</i> .
Netwait IP list	string	If Enable netwait feature is unset, list of IP addresses to ping. Otherwise,
		ping the default gateway.
Host name	string	Used to add one entry per line which will be appended to /etc/hosts.
database		Use the format <i>IP_address space hostname</i> where multiple hostnames can
		be used if separated by a space.

Table 7.1 - continued from previous page

When using Active Directory, set the IP address of the realm's DNS server in the *Nameserver 1* field.

If the network does not have a DNS server, or NFS, SSH, or FTP users are receiving "reverse DNS" or timeout errors, add an entry for the IP address of the FreeNAS[®] system in the *Host name database* field.

Note: In many cases, a FreeNAS[®] configuration does not include default gateway information as a way to make it more difficult for a remote attacker to communicate with the server. While this is a reasonable precaution, such a configuration does **not** restrict inbound traffic from sources within the local network. However, omitting a default gateway will prevent the FreeNAS[®] system from communicating with DNS servers, time servers, and mail servers that are located outside of the local network. In this case, it is recommended to add *Static Routes* (page 123) to be able to reach external DNS, NTP, and mail servers which are configured with static IP addresses. When a gateway to the Internet is added, make sure the FreeNAS[®] system is protected by a properly configured firewall.

7.2 Interfaces

 $Network \rightarrow Interfaces$ shows which interfaces have been manually configured and allows adding or editing a manually configured interface.

Note: Typically, the interface used to access the FreeNAS[®] administrative GUI is configured by DHCP. This interface does not appear in this screen, even though it is already dynamically configured and in use.

Creating a Link Aggregation (page 120) that does **not** include the NIC used to access the FreeNAS[®] administrative GUI may require adding an *Interfaces* entry for this interface with DHCP enabled. See this *warning* (page 113) about changing the interface that the web interface uses.

Figure 7.2 shows the screen that opens on clicking *Interfaces* \rightarrow *Add Interface*. Table 7.2 summarizes the configuration options shown when adding an interface or editing an already configured interface. Note that if any changes to this screen require a network restart, the screen will turn red when the *OK* button is clicked and a pop-up message will point out that network connectivity to the FreeNAS[®] system will be interrupted while the changes are applied.

Add Interface 🕺 🚖		
NIC:	em0 🔽	
Interface Name:	a	
DHCP:	(
IPv4 Address:		
IPv4 Netmask:		
Auto configure IPv6:	i	
IPv6 Address:		
IPv6 Prefix Length:		
Options:		
Alias	-	
IPv4 Address:		



Setting	Value	Description
NIC	drop-down menu	The FreeBSD device name of the interface. This is a read-only field when
		editing an interface.
Interface Name	string	Description of interface.
DHCP	checkbox	Requires static IPv4 or IPv6 configuration if unselected. Only one interface
		can be configured for DHCP.
IPv4 Address	IP address	Enter a static IP address if <i>DHCP</i> is unset.
IPv4 Netmask	drop-down menu	Enter a netmask if DHCP is unset.
Auto configure	checkbox	Only one interface can be configured for this option. If unset, manual con-
IPv6		figuration is required to use IPv6.
IPv6 Address	IPv6 address	Must be unique on the network.
IPv6 Prefix	drop-down menu	Match the prefix used on the network.
Length		
Options	string	Additional parameters from ifconfig(8)
		(https://www.freebsd.org/cgi/man.cgi?query=ifconfig). Separate mul-
		tiple parameters with a space. For example: <i>mtu 9000</i> increases the MTU
		for interfaces which support jumbo frames (but see <i>this note</i> (page 122)
		about MTU and lagg interfaces).

Table 7.2	Interface	Configuration	h Settings
10010 7.2.	michace	Configuration	Julias

This screen also provides for the configuration of IP aliases, making it possible for a single interface to have multiple IP addresses. To set multiple aliases, click the *Add extra alias* link for each alias. Aliases are deleted by clicking the interface in the tree, clicking the *Edit* button, checking the *Delete* checkbox below the alias, then clicking the *OK* button.

Warning: Aliases are deleted by checking the Delete checkbox in the alias area, then clicking OK for the interface. Do

not click the Delete button at the bottom of this screen, which deletes the entire interface.

Multiple interfaces **cannot** be members of the same subnet. See Multiple network interfaces on a single subnet (https://forums.freenas.org/index.php?threads/multiple-network-interfaces-on-a-single-subnet.20204/) for more information. Check the subnet mask if an error is shown when setting the IP addresses on multiple interfaces.

This screen will not allow an interface's IPv4 and IPv6 addresses to both be set as primary addresses. An error is shown if both the *IPv4 address* and *IPv6 address* fields are filled in. Instead, set only one of these address fields and create an alias for the other address.

7.3 IPMI

Beginning with version 9.2.1, FreeNAS[®] provides a graphical screen for configuring an IPMI interface. This screen will only appear if the system hardware includes a Baseboard Management Controller (BMC).

IPMI provides side-band management if the graphical administrative interface becomes unresponsive. This allows for a few vital functions, such as checking the log, accessing the BIOS setup, and powering on the system without requiring physical access to the system. IPMI is also used to give another person remote access to the system to assist with a configuration or troubleshooting issue. Before configuring IPMI, ensure that the management interface is physically connected to the network. The IPMI device may share the primary Ethernet interface, or it may be a dedicated separate IPMI interface.

Warning: It is recommended to first ensure that the IPMI has been patched against the Remote Management Vulnerability before enabling IPMI. This article (https://www.ixsystems.com/blog/how-to-fix-the-ipmi-remote-managementvulnerability/) provides more information about the vulnerability and how to fix it.

Note: Some IPMI implementations require updates to work with newer versions of Java. See PSA: Java 8 Update 131 breaks ASRock's IPMI Virtual console (https://forums.freenas.org/index.php?threads/psa-java-8-update-131-breaks-asrocks-ipmi-virtual-console.53911/) for more information.

IPMI is configured from *Network* \rightarrow *IPMI*. The IPMI configuration screen, shown in Figure 7.3, provides a shortcut to the most basic IPMI configuration. Those already familiar with IPMI management tools can use them instead. Table 7.3 summarizes the options available when configuring IPMI with the FreeNAS[®] GUI.

Network			
Global Configuration Interfac	es IPMI Link Aggregations Network Summary Static Routes VLANs		
Channel:			
Password:			
Password confirmation:	<i>i</i>		
DHCP:			
IPv4 Address:	10.275.1.70		
IPv4 Netmask:	/16 (255.255.0.0)		
IPv4 Default Gateway:	14.4708.0.0		
VLAN ID:			
OK Cancel Identify Light			

Fig. 7.3: IPMI Configuration

Setting	Value	Description
Channel	drop-down menu	Select the channel to use.
Password	string	Enter the password used to connect to the IPMI interface from a web
		browser.
DHCP	checkbox	If left unset, the next three fields must be set.
IPv4 Address	string	IP address used to connect to the IPMI web GUI.
IPv4 Netmask	drop-down menu	Subnet mask associated with the IP address.
IPv4 Default	string	Default gateway associated with the IP address.
Gateway		
VLAN ID	string	Enter the VLAN identifier if the IPMI out-of-band management interface is
		not on the same VLAN as management networking.

Table 7.3: IPMI Options

The *Identify Light* button can be used to identify a system in a multi-system rack by flashing its IPMI LED light. Clicking this button will present a pop-up with a menu of times, ranging from 15 seconds to 4 minutes, to flash the LED light.

After configuration, the IPMI interface is accessed using a web browser and the IP address specified in the configuration. The management interface prompts for a username and the configured password. Refer to the IPMI device's documentation to determine the default administrative username.

After logging in to the management interface, the default administrative username can be changed, and additional users created. The appearance of the IPMI utility and the functions that are available vary depending on the hardware.

A command-line utility called ipmitool is available to control many features of the IPMI interface. See How To: Change IPMI Sensor Thresholds using ipmitool (https://forums.freenas.org/index.php?resources/how-to-change-ipmi-sensor-thresholds-using-ipmitool.35/) for some examples.

7.4 Link Aggregations

FreeNAS[®] uses the FreeBSD lagg(4) (https://www.freebsd.org/cgi/man.cgi?query=lagg) interface to provide link aggregation and link failover support. A lagg interface allows combining multiple network interfaces into a single virtual interface. This provides fault-tolerance and high-speed multi-link throughput. The aggregation protocols supported by lagg both determines the ports to use for outgoing traffic and if a specific port accepts incoming traffic. The link state of the lagg interface is used to validate whether the port is active.

Aggregation works best on switches supporting LACP, which distributes traffic bi-directionally while responding to failure of individual links. FreeNAS[®] also supports active/passive failover between pairs of links. The LACP and load-balance modes select the output interface using a hash that includes the Ethernet source and destination address, VLAN tag (if available), IP source and destination address, and flow label (IPv6 only). The benefit can only be observed when multiple clients are transferring files *from* the NAS. The flow entering *into* the NAS depends on the Ethernet switch load-balance algorithm.

The lagg driver currently supports several aggregation protocols, although only *Failover* is recommended on network switches that do not support *LACP*:

Failover: the default protocol. Sends traffic only through the active port. If the master port becomes unavailable, the next active port is used. The first interface added is the master port. Any interfaces added later are used as failover devices. By default, received traffic is only accepted when received through the active port. This constraint can be relaxed, which is useful for certain bridged network setups, by creating a tunable with a *Variable* of *net.link.lagg.failover_rx_all*, a *Value* of a non-zero integer, and a *Type* of *Sysctl* in *System* \rightarrow *Tunables* \rightarrow *Add Tunable*.

LACP: supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP) and the Marker Protocol. LACP negotiates a set of aggregable links with the peer into one or more link aggregated groups (LAGs). Each LAG is composed of ports of the same speed, set to full-duplex operation. Traffic is balanced across the ports in the LAG with the greatest total speed; in most cases there will only be one LAG which contains all ports. In the event of changes in physical connectivity, link aggregation will quickly converge to a new configuration. LACP must be configured on the switch, and LACP does not support mixing interfaces of different speeds. Only interfaces that use the same driver, like two *igb* ports, are recommended for LACP. Using LACP for iSCSI is not recommended, as iSCSI has built-in multipath features which are more efficient.

Note: When using LACP, verify the switch is configured for active LACP. Passive LACP is not supported.

Load Balance: balances outgoing traffic across the active ports based on hashed protocol header information and accepts incoming traffic from any active port. This is a static setup and does not negotiate aggregation with the peer or exchange frames to monitor the link. The hash includes the Ethernet source and destination address, VLAN tag (if available), and IP source and destination address. Requires a switch which supports IEEE 802.3ad static link aggregation.

Round Robin: distributes outgoing traffic using a round-robin scheduler through all active ports and accepts incoming traffic from any active port. This mode can cause unordered packet arrival at the client. This has a side effect of limiting throughput as reordering packets can be CPU intensive on the client. Requires a switch which supports IEEE 802.3ad static link aggregation.

None: this protocol disables any traffic without disabling the lagg interface itself.

7.4.1 LACP, MPIO, NFS, and ESXi

LACP bonds Ethernet connections to improve bandwidth. For example, four physical interfaces can be used to create one mega interface. However, it cannot increase the bandwidth for a single conversation. It is designed to increase bandwidth when multiple clients are simultaneously accessing the same system. It also assumes that quality Ethernet hardware is used and it will not make much difference when using inferior Ethernet chipsets such as a Realtek.

LACP reads the sender and receiver IP addresses and, if they are deemed to belong to the same TCP connection, always sends the packet over the same interface to ensure that TCP does not need to reorder packets. This makes LACP ideal for load balancing many simultaneous TCP connections, but does nothing for increasing the speed over one TCP connection.

MPIO operates at the iSCSI protocol level. For example, if four IP addresses are created and there are four simultaneous TCP connections, MPIO will send the data over all available links. When configuring MPIO, make sure that the IP addresses

on the interfaces are configured to be on separate subnets with non-overlapping netmasks, or configure static routes to do point-to-point communication. Otherwise, all packets will pass through one interface.

LACP and other forms of link aggregation generally do not work well with virtualization solutions. In a virtualized environment, consider the use of iSCSI MPIO through the creation of an iSCSI Portal with at least two network cards on different networks. This allows an iSCSI initiator to recognize multiple links to a target, using them for increased bandwidth or redundancy. This how-to (https://fojta.wordpress.com/2010/04/13/iscsi-and-esxi-multipathing-and-jumbo-frames/) contains instructions for configuring MPIO on ESXi.

NFS does not understand MPIO. Therefore, one fast interface is needed, since creating an iSCSI portal will not improve bandwidth when using NFS. LACP does not work well to increase the bandwidth for point-to-point NFS (one server and one client). LACP is a good solution for link redundancy or for one server and many clients.

7.4.2 Creating a Link Aggregation

Before creating a link aggregation, make sure that all interfaces to use in the lagg are not manually configured in *Network* \rightarrow *Interfaces* \rightarrow *View Interfaces*.

Lagg creation fails if any of the included interfaces are manually configured. See this *warning* (page 113) about changing the interface that the web interface uses.

Figure 7.4 shows the configuration options when adding a lagg interface using Network \rightarrow Link Aggregations \rightarrow Create Link Aggregation.

Add Link Aggregation	X
Protocol Type:	 Failover LACP Load Balance Round Robin None
Physical NICs in the LAGG:	em0
OK Cancel	

Fig. 7.4: Creating a lagg Interface

Note: If interfaces are installed but do not appear in the *Physical NICs* list, check that a FreeBSD driver for the interface exists here (https://www.freebsd.org/releases/11.1R/hardware.html#ethernet).

To create a link aggregation, select the desired *Protocol Type*. *LACP* is preferred. If the network switch does not support LACP, choose *Failover*. Highlight the interfaces to associate with the lagg device, and click the *OK* button.

Once the lagg device has been created, click its entry to enable its *Edit*, *Delete*, and *Edit Members* buttons.

Clicking the *Edit* button for a lagg opens the configuration screen shown in Figure 7.5. Table 7.4 describes the options in this screen.

Network								
Global Configuration	n Interfaces	Link Ag	gregation	Network Su	mmary	Static Routes	VLAN	
Add Link Aggregation		Ec	lit	_		_		88
Interface	Protocol Type	e	NIC:		lagg0			
lagg0 (none: em0)	none		Interface	e Name:	lagg0			
			DHCP:		i			
			IPv4 Add	dress:				
			IPv4 Net	mask:		-	•	
		L	Auto con IPv6:	nfigure	i			
			IPv6 Add	dress:				
			IPv6 Pre Length:	fix		-	•	
4			Options:					
Edit Delete Ed	dit Members		Alias					
			IPv4	Address:				
								-

Fig. 7.5: Editing a lagg

Table 7 1	Configurabl	a Ontiana	foralage	
Table 7.4:	Configurabl	e Options	for a lagg	

Setting	Value	Description
NIC	string	Read-only. Automatically assigned the next available numeric ID.
Interface Name	string	By default, this is the same as device (NIC) name. This can be changed to a
		more descriptive value.
DHCP	checkbox	Enable if the lagg device will get IP address info from DHCP server. The IP
		address of the new lagg can be set to DHCP only if no other interface uses
		DHCP.
IPv4 Address	string	Enter a static IP address if <i>DHCP</i> is unset.
IPv4 Netmask	drop-down menu	Enter a netmask if DHCP is unset.
Auto configure	checkbox	Set only if DHCP server available to provide IPv6 address info
IPv6		
IPv6 Address	string	This is optional.
IPv6 Prefix	drop-down menu	Required if an IPv6 address is entered.
Length		
Options	string	Additional ifconfig(8) (https://www.freebsd.org/cgi/man.cgi?query=ifconfig)
		options.

This screen also allows the configuration of an alias for the lagg interface. Multiple aliases can be added with the *Add extra Alias* link.

Click the *Edit Members* button, click the entry for a member, then click its *Edit* button to see the configuration screen shown in Figure 7.6. The configurable options are summarized in Table 7.5.

Network								
Global Configuration Ir	nterfaces Link A	ggregation	n Network	Summar	y Static Routes V	LAN	LAGG Members undefine	ed
Add Link Aggregation Merr	ıber			Edit	_		Σ	3
LAGG Interface Group	LAGG Priority Number	Physical NIC	Options	LAG	G Interface Group:	agg0	(none: em0)	
lagg0 (none: em0)	0	em0	up	LAG	G Priority Number:	0		
				LAG	G Physical NIC:	em0	v	
				Opt	ions:	up		
				ок	Cancel Delete)		
				_				
Edit Delete								

Fig. 7.6: Editing a Member Interface

Setting	Value	Description
LAGG Interface	drop-down menu	Select the member interface to configure.
group		
LAGG Priority	integer	Order of selected interface within the lagg. Configure a failover to set the
Number		master interface to 0 and the other interfaces to 1, 2, etc.
LAGG Physical	drop-down menu	Physical interface of the selected member. The drop-down is empty when
NIC		no NICs are available.
Options	string	Additional parameters from ifconfig(8)
		(https://www.freebsd.org/cgi/man.cgi?query=ifconfig).

Table 7.5:	Configuring a	Member	Interface
------------	---------------	--------	-----------

Click Add Link Aggregation Member to see the same options. Click OK to add the new member to the list.

Options can be set at the lagg level using the *Edit* button, or at the individual parent interface level using the *Edit Members* button. Changes are typically made at the lagg level (Figure 7.5) as each interface member will inherit from the lagg. To configure at the interface level (Figure 7.6) instead, repeat the configuration for each interface within the lagg. Some options can only be set on the parent interfaces and are inherited by the lagg interface. For example, to set the MTU on a lagg, use *Edit Members* to set the MTU for each parent interface.

If the MTU settings on the lagg member interfaces are not identical, the smallest value is used for the MTU of the entire lagg.

Note: A reboot is required after changing the MTU to create a jumbo frame lagg.

Link aggregation load balancing can be tested with:

systat -ifstat

More information about this command can be found at systat(1) (https://www.freebsd.org/cgi/man.cgi?query=systat).

7.5 Network Summary

 $Network \rightarrow Network$ Summary shows a quick summary of the addressing information of every configured interface. For each interface name, the configured IPv4 and IPv6 addresses, DNS servers, and default gateway are displayed.

7.6 Static Routes

No static routes are defined on a default FreeNAS[®] system. If a static route is required to reach portions of the network, add the route with *Network* \rightarrow *Static Routes* \rightarrow *Add Static Route*, shown in Figure 7.7.

Add Static Route					
Destination network:					
Gateway:					
Description:					
ОК Сапсеі					

Fig. 7.7: Adding a Static Route

The available options are summarized in Table 7.6.

Table 7.6: Static Route Options

Setting	Value	Description
Destination net- work	integer	Use the format <i>A.B.C.D/E</i> where <i>E</i> is the CIDR mask.
Gateway	integer	Enter the IP address of the gateway.
Description	string	Optional. Add any notes about the route.

Added static routes are shown in View Static Routes. Click a route's entry to access the Edit and Delete buttons.

7.7 VLANs

FreeNAS[®] uses FreeBSD's vlan(4) (https://www.freebsd.org/cgi/man.cgi?query=vlan) interface to demultiplex frames with IEEE 802.1q tags. This allows nodes on different VLANs to communicate through a layer 3 switch or router. A vlan interface must be assigned a parent interface and a numeric VLAN tag. A single parent can be assigned to multiple vlan interfaces provided they have different tags.

Note: VLAN tagging is the only 802.1q feature that is implemented. Additionally, not all Ethernet interfaces support full VLAN processing. See the HARDWARE section of vlan(4) (https://www.freebsd.org/cgi/man.cgi?query=vlan) for details.

Click *Network* \rightarrow *VLANs* \rightarrow *Add VLAN*, to see the screen shown in Figure 7.8.

Add VLAN 🕺					
Virtual Interface:					
Parent Interface:	em0 💌				
VLAN Tag:					
Priority Code Point (CoS):					
Description:					
OK Cancel					

Fig. 7.8: Adding a VLAN

Table 7.7 summarizes the configurable fields.

Table 7.7: Adding a VLAN

Setting	Value	Description
Virtual Interface	string	Use the format <i>vlanX</i> where <i>X</i> is a number representing a vlan interface
		not currently being used as a parent.
Parent Interface	drop-down menu	Usually an Ethernet card connected to a properly configured switch port.
		Newly created <i>Link Aggregations</i> (page 119) do not appear in the drop-
		down until the system is rebooted.
VLAN Tag	integer	Enter a number between 1 and 4095 which matches a numeric tag set up
		in the switched network.
Priority Code	drop-down menu	Available 802.1p Class of Service ranges from Best Effort (default) to Net-
Point		work Control (highest).
Description	string	Optional. Enter any notes about this VLAN.

The parent interface of a VLAN must be up, but it can either have an IP address or be unconfigured, depending upon the requirements of the VLAN configuration. This makes it difficult for the GUI to do the right thing without trampling the configuration. To remedy this, add the VLAN, then select *Network* \rightarrow *Interfaces* \rightarrow *Add Interface*. Choose the parent interface from the *NIC* drop-down menu and in the *Options* field, type up. This will bring up the parent interface. If an IP address is required, it can be configured using the rest of the options in the *Add Interface* screen.

Warning: Creating a VLAN causes an interruption to network connectivity. The GUI provides a warning and an opportunity to cancel the VLAN creation.

STORAGE

The Storage section of the graphical interface allows configuration of these options:

- Volumes (page 125) create and manage storage volumes.
- Periodic Snapshot Tasks (page 150) schedule automatic creation of filesystem snapshots.
- Replication Tasks (page 152) automate the replication of snapshots to a remote system.
- Resilver Priority (page 161) control the priority of resilvers.
- Scrubs (page 162) schedule scrubs as part of ongoing disk maintenance.
- Snapshots (page 165) manage local snapshots.
- VMware-Snapshot (page 167) coordinate OpenZFS snapshots with a VMware datastore.

8.1 Volumes

The *Volumes* section of the FreeNAS[®] graphical interface is used to format volumes, attach a disk to copy data onto an existing volume, or import a ZFS volume. It is also used to create ZFS datasets and zvols and to manage their permissions.

Note: In ZFS terminology, groups of storage devices managed by ZFS are referred to as a *pool*. The FreeNAS[®] graphical interface uses the term *volume* to refer to a ZFS pool.

Proper storage design is important for any NAS. Please read through this entire chapter before configuring storage disks. Features are described to help make it clear which are beneficial for particular uses, and caveats or hardware restrictions which limit usefulness.

8.1.1 Volume Manager

Before creating a volume, determine the level of required redundancy, how many disks will be added, and if any data exists on those disks. Creating a volume overwrites disk data, so save any required data to different media before adding disks to a pool. Refer to the *ZFS Primer* (page 320) for information on ZFS redundancy with multiple disks before using *Volume Manager*. It is important to realize that different layouts of virtual devices (*vdevs*) affect which operations can be performed on that volume later. For example, drives can be added to a mirror to increase redundancy, but that is not possible with RAIDZ arrays.

To create a volume, click *Storage* \rightarrow *Volume Manager*. This opens a screen like the example shown in Figure 8.1.

Volume Manager 8	Ж
Volume Name Volume to extend The formula of the f	
Image: Capacity: 0 B Add Extra Device	
Add Volume Existing data will be cleared	•

Fig. 8.1: Creating a ZFS Pool Using Volume Manager

Table 8.1 summarizes the configuration options of this screen.

Setting	Value	Description	
Volume name string		ZFS volumes must conform to these naming conventions	
		(https://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html)	
		Choose a memorable name that sticks out in the logs and avoid generic	
		names.	
Volume to extend	drop-down	Extend an existing ZFS pool. See <i>Extending a ZFS Volume</i> (page 129) for	
	menu	more details.	
Encryption	checkbox	See the warnings in <i>Encryption</i> (page 127) before enabling encryption.	
Available disks	display	Display the number and size of available disks. Hover over show to list the	
		available device names, and click the + to add all of the disks to the pool.	
Volume layout	drag and	Click and drag the icon to select the desired number of disks for a vdev.	
	drop	When at least one disk is selected, the layouts supported by the selected	
		number of disks are added to the drop-down menu.	
Add Extra Device	button	Configure multiple vdevs or add log or cache devices during pool creation.	
Manual setup	button	Create a pool manually, which is not recommended. See <i>Manual Setup</i>	
		(page 128) for more details.	

Table 8.1: ZFS Volume Creation Options

Click the *Volume name* field and enter a name for the pool. Ensure that the chosen name conforms to these naming conventions (http://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html).

If the underlying disks need to be encrypted as a protection against physical theft, enable the *Encryption* option.

Warning: Refer to the warnings in *Encryption* (page 127) before enabling encryption! Be aware that this form of encryption will be replaced by OpenZFS native encryption in a future version. Pools created with the current encryption mechanism will have to be backed up and destroyed to be recreated with native encryption when it becomes available.

Drag the slider to select the desired number of disks. *Volume Manager* displays the resulting storage capacity, taking reserved swap space into account. To change the layout or the number of disks, drag the slider to the desired volume layout. The *Volume layout* drop-down menu can also be clicked if a different level of redundancy is required.

Note: For performance and capacity reasons, this screen does not allow creating a volume from disks of differing sizes. While it is not recommended, it is possible to create a volume of differently-sized disks with the *Manual setup* button. Follow the instructions in *Manual Setup* (page 128).

Volume Manager only allows choosing a configuration if enough disks have been selected to create that configuration. These layouts are supported:

- **Stripe:** requires at least one disk
- Mirror: requires at least two disks
- RAIDZ1: requires at least three disks
- RAIDZ2: requires at least four disks
- RAIDZ3: requires at least five disks
- log device: requires at least one dedicated device, a fast, low-latency, power-protected SSD is recommended
- cache device: requires at least one dedicated device, SSD is recommended

When more than five disks are used, consideration must be given to the optimal layout for the best performance and scalability. An overview of the recommended disk group sizes as well as more information about log and cache devices can be found in the *ZFS Primer* (page 320).

The *Add Volume* button warns that **existing data will be cleared**. In other words, creating a new volume **reformats the selected disks**. To preserve existing data, click the *Cancel* button and refer to *Import Disk* (page 137) and *Import Volume* (page 138) to see if the existing format is supported. If so, perform that action instead. If the current storage format is not supported, it is necessary to back up the data to external media, format the disks, then restore the data to the new volume.

Depending on the size and number of disks, the type of controller, and whether encryption is selected, creating the volume may take some time. After the volume is created, the screen refreshes and the new volume is listed in the tree under *Storage* \rightarrow *Volumes*. Click the + next to the volume name to access *Change Permissions* (page 130), *Create Dataset* (page 132), and *Create zvol* (page 135) options for that volume.

8.1.1.1 Encryption

Note: The encryption facility used by FreeNAS[®] is designed to protect against physical theft of the disks. It is not designed to protect against unauthorized software access. Ensure that only authorized users have access to the administrative GUI and that proper permissions are set on shares if sensitive data is stored on the system.

FreeNAS[®] supports GELI (https://www.freebsd.org/cgi/man.cgi?query=geli) full disk encryption for ZFS volumes. It is important to understand the details when considering whether encryption is right for the intended use:

- FreeNAS[®] encryption is different from the encryption used in Oracle's proprietary, non-open source version of ZFS.
- In FreeNAS[®], entire disks are encrypted, not individual filesystems. Encrypted devices are created from the underlying drives, then the volume (pool) is created on top of the encrypted devices. Data is encrypted as it is written and decrypted as it is read.
- This type of encryption is primarily useful for users storing sensitive data but wanting the ability to remove disks from the pool without having to first wipe the disk contents.

- The FreeNAS[®] encryption design is only suitable for safe disposal of disks independent of the encryption key. As long as the key and the disks are intact, the system is vulnerable to being decrypted. Protect the key with a strong passphrase and secure all backups of it.
- If the encryption key is lost, the data on the disks is inaccessible. Always back up the key!
- Encryption keys are per ZFS volume (pool). Each pool has a separate encryption key. Technical details about how encryption keys are used, stored, and managed within FreeNAS[®] are described in this forum post (https://forums.freenas.org/index.php?threads/recover-encryption-key.16593/#post-85497).
- Data in memory, including ARC, is not encrypted. ZFS data on disk, including ZIL and SLOG, are encrypted if the underlying disks are encrypted. Swap data on disk is always encrypted.
- All drives in an encrypted volume are encrypted, including L2ARC (read cache) and SLOG (write cache). Drives added to an existing encrypted volume are encrypted with the same method specified when the volume was created.
- At present, there is no one-step way to encrypt an existing, unencrypted volume. Instead, the data must be backed up, the existing pool destroyed, a new encrypted volume created, and the backup restored to the new volume.
- Hybrid pools are not supported. Added vdevs must match the existing encryption scheme. *Volume Manager* (page 125) automatically encrypts a new vdev being added to an existing encrypted pool.

To create an encrypted volume, enable the *Encryption* option shown in Figure 8.1. A pop-up message shows a reminder that **it is extremely important to make a backup of the key**. Without the key, the data on the disks is inaccessible. See *Managing Encrypted Volumes* (page 143) for instructions.

8.1.1.2 Encryption Performance

Encryption performance depends upon the number of disks encrypted. The more drives in an encrypted volume, the more encryption and decryption overhead, and the greater the impact on performance. **Encrypted volumes composed of more than eight drives can suffer severe performance penalties**. If encryption is desired, please benchmark such volumes before using them in production.

Note: Processors with support for the AES-NI (https://en.wikipedia.org/wiki/AES_instruction_set#Supporting_x86_CPUs) instruction set are strongly recommended. These processors can handle encryption of a small number of disks with negligible performance impact. They also retain performance better as the number of disks increases. Older processors without the AES-NI instructions see significant performance impact with even a single encrypted disk. This forum post (https://forums.freenas.org/index.php?threads/encryption-performance-benchmarks.12157/) compares the performance of various processors.

8.1.1.3 Manual Setup

The *Manual Setup* button shown in Figure 8.1 can be used to create a ZFS volume manually. While this is **not** recommended, it can, for example, be used to create a non-optimal volume containing disks of different sizes.

Note: The usable space of each disk in a volume is limited to the size of the smallest disk in the volume. Because of this, creating volumes with disks of the same size through the *Volume Manager* is recommended.

Figure 8.2 shows the *Manual Setup* screen. Table 8.2 shows the available options.

lanual Setup	
Volume name	
Encryption	
Member disks (0)	ada1 (21.5 GB) ada2 (21.5 GB) ada3 (21.5 GB) ada4 (21.5 GB) ada5 (21.5 GB)
Deduplication	off 💌
ZFS Extra	Disk None Log Cache Spare ada1
Add Volume Existing data will be cleared	ada5 Cancel

Fig. 8.2: Manually Creating a ZFS Volume

Note: Because of the disadvantages of creating volumes with disks of different sizes, the displayed list of disks is sorted by size.

Setting	Value	Description
Volume name	string	ZFS volumes must conform to these naming conventions
		(https://docs.oracle.com/cd/E53394_01/index.html). Choosing a memo-
		rable name is recommended.
Encryption	checkbox	See the warnings in <i>Encryption</i> (page 127) before using encryption.
Member disks	list	Highlight desired number of disks from list of available disks.
Deduplication	drop-down	Choices are Off, Verify, and On. Carefully consider the section on Deduplica-
	menu	<i>tion</i> (page 134) before changing this setting.
ZFS Extra	bullet selec-	Specify disk usage: storage (<i>None</i>), a log device, a cache device, or a spare.
	tion	

Table 8.2: Ma	anual Setup	Options
---------------	-------------	---------

8.1.1.4 Extending a ZFS Volume

The *Volume to extend* drop-down menu in *Storage* \rightarrow *Volumes* \rightarrow *Volume Manager*, shown in Figure 8.1, is used to add disks to an existing ZFS volume to increase capacity. This menu is empty if there are no ZFS volumes yet.

If more than one disk is added, the arrangement of the new disks into stripes, mirrors, or RAIDZ vdevs can be specified. Mirrors and RAIDZ arrays provide redundancy for data protection if an individual drive fails.

Note: If the existing volume is encrypted, a warning message shows a reminder that **extending a volume resets the passphrase and recovery key**. After extending the volume, immediately recreate both using the instructions in *Managing Encrypted Volumes* (page 143).

After an existing volume has been selected from the drop-down menu, drag and drop the desired disks and select the desired volume layout. For example, disks can be added to increase the capacity of the volume.

When adding disks to increase the capacity of a volume, ZFS supports the addition of virtual devices, or *vdevs*, to an existing ZFS pool. A vdev can be a single disk, a stripe, a mirror, a RAIDZ1, RAIDZ2, or a RAIDZ3. **After a vdev is created, more drives cannot be added to that vdev**. However, a new vdev can be striped with another of the **same type of existing vdev** to increase the overall size of the volume. Extending a volume often involves striping similar vdevs. Here are some examples:

- to extend a ZFS stripe, add one or more disks. Since there is no redundancy, disks do not have to be added in the same quantity as the existing stripe.
- to extend a ZFS mirror, add the same number of drives. The resulting striped mirror is a RAID 10. For example, if ten new drives are available, a mirror of two drives could be created initially, then extended by creating another mirror of two drives, and repeating three more times until all ten drives have been added.
- to extend a three drive RAIDZ1, add three additional drives. The result is a RAIDZ+0, similar to RAID 50 on a hardware controller.
- to extend a RAIDZ2 requires a minimum of four additional drives. The result is a RAIDZ2+0, similar to RAID 60 on a hardware controller.

If an attempt is made to add a non-matching number of disks to the existing vdev, an error message appears, indicating the number of disks that are required. Select the correct number of disks to continue.

Adding L2ARC or SLOG Devices

Storage \rightarrow Volumes \rightarrow Volume Manager (see Figure 8.1) is also used to add L2ARC or SLOG SSDs to improve volume performance for specific use cases. Refer to the ZFS Primer (page 320) to determine if the system will benefit or suffer from the addition of the device.

Once the SSD has been physically installed, click the *Volume Manager* button and choose the volume from the *Volume to extend* drop-down menu. Click the + next to the SSD in the *Available disks* list. In the *Volume layout* drop-down menu, select *Cache (L2ARC)* to add a cache device, or *Log (ZIL)* to add a log device. Finally, click *Extend Volume* to add the SSD.

Removing L2ARC or SLOG Devices

Cache or log devices can be removed by going to *Storage* \rightarrow *Volumes*. Choose the desired pool and click *Volume Status*. Choose the log or cache device to remove, then click *Remove*.

8.1.2 Change Permissions

Setting permissions is an important aspect of managing data access. The graphical administrative interface is meant to set the **initial** permissions for a volume or dataset to make it available as a share. After a share has been created, the client operating system is used to fine-tune the permissions of the files and directories that are created by the client.

Sharing (page 180) contains configuration examples for several types of permission scenarios. This section provides an overview of the options available for configuring the initial set of permissions.

Note: For users and groups to be available, they must either be first created using the instructions in *Account* (page 58) or imported from a directory service using the instructions in *Directory Services* (page 169). If more than 50 users or groups are available, the drop-down menus described in this section will automatically truncate their display to 50 for performance reasons. In this case, start to type in the desired user or group name so that the display narrows its search to matching results.

After a volume or dataset is created, it is listed by its mount point name in *Storage* \rightarrow *Volumes*. Clicking the *Change Permissions* icon for a specific volume or dataset displays the screen shown in Figure 8.3. Table 8.3 summarizes the options in this screen.

С	hange Permissions	88	-
(Change permission		
Cł	nange permission on /mnt/vo	lume1 to:	
	Apply Owner (user):		
	Owner (user):	root	
	Apply Owner (group):		
	Owner (group):	wheel	
	Apply Mode:		
	Mode:	Owner Group Other Read 🔽 🔽 💟 Write 🔽 📄 📄 Execute 💟 💟 💟	
	Permission Type:	• (©) Unix • (©) Mac • (©) Windows	
	Set permission		•

Fig. 8.3: Changing Permissions on a Volume or Dataset

Table 8.3: Options When	Changing Permissions
-------------------------	----------------------

Setting	Value	Description
Apply Owner (user)	checkbox	Deselect to prevent new permission change from being applied to Owner
		<i>(user)</i> , see Note below.
Owner (user)	drop-down	Select the user to control the volume or dataset. Users manually created
	menu	or imported from a directory service will appear in the drop-down menu.
Apply Owner (group)	checkbox	Deselect to prevent new permission change from being applied to Owner
		(group), see Note below for more information.
Owner (group)	drop-down	Select the group to control the volume or dataset. Groups manually cre-
	menu	ated or imported from a directory service will appear in the drop-down
		menu.
Apply Mode	checkbox	Deselect to prevent new permission change from being applied to <i>Mode</i> ,
		see Note below.
Mode	checkboxes	Only applies to the <i>Unix</i> or <i>Mac</i> "Permission Type". Will be grayed out if
		Windows is selected.
Permission Type	bullet selec-	Select the type which matches the type of client accessing the volume or
	tion	dataset. Choices are Unix, Mac, or Windows.
Set permission recursively	checkbox	If enabled, permissions will also apply to subdirectories of the volume or
		dataset. If data already exists on the volume or dataset, change the per-
		missions on the client side to prevent a performance lag.

Note: The *Apply Owner (user)*, *Apply Owner (group)*, and *Apply Mode* options allow fine-tuning of the change permissions behavior. By default, all options are enabled and FreeNAS[®] resets the owner, group, and mode when the *Change* button is clicked. These optionss allow choosing which settings to change. For example, to change just the *Owner (group)* setting, deselect the *Apply Owner (user)* and *Apply Mode* options.

The *Windows Permission Type* is used for *Windows (SMB) Shares* (page 195) or when the FreeNAS[®] system is a member of an Active Directory domain. This type adds ACLs to traditional *Unix* permissions. When the *Windows Permission Type* is set, ACLs are set to the Windows defaults for new files and directories. A Windows client can be used to further fine-tune permissions as needed.

Warning: Changing a volume or dataset with *Windows* permissions back to *Unix* permissions will overwrite and destroy some of the extended permissions provided by *Windows* ACLs.

The Unix Permission Type is usually used with Unix (NFS) Shares (page 188). Unix permissions are compatible with most network clients and generally work well with a mix of operating systems or clients. However, Unix permissions do not support Windows ACLs. Do not use them with Windows (SMB) Shares (page 195).

The Mac Permission Type can be used with Apple (AFP) Shares (page 181).

8.1.3 Create Dataset

An existing ZFS volume can be divided into datasets. Permissions, compression, deduplication, and quotas can be set on a per-dataset basis, allowing more granular control over access to storage data. Like a folder or directory, permissions can be set on dataset. Datasets are also similar to filesystems in that properties such as quotas and compression can be set, and snapshots created.

Note: ZFS provides thick provisioning using quotas and thin provisioning using reserved space.

Selecting an existing ZFS volume in the tree and clicking *Create Dataset* shows the screen in Figure 8.4.

Create Dataset	2
Create ZFS dataset in volume1	
Dataset Name:	
Comments:	
Sync:	Inherit (standard) 💌
Compression level:	Inherit (Iz4)
Share type:	
Enable atime:	 Inherit (on) On Off
ZFS Deduplication:	Enabling dedup can drastically reduce performance and affect the ability to access data. Compression usually offers similar space savings with much lower performance impact and overhead.
	Inherit (off)
Case Sensitivity:	Sensitive 🔽
Add Dataset Cancel Ad	dvanced Mode

Fig. 8.4: Creating a ZFS Dataset

Table 8.4 shows the options available when creating a dataset. Some settings are only available in *Advanced Mode*. To see these settings, either click the *Advanced Mode* button, or configure the system to always display advanced settings by enabling the *Show advanced fields by default* option in *System* \rightarrow *Advanced*. Most attributes, except for the *Dataset Name, Case Sensitivity*, and *Record Size*, can be changed after dataset creation by highlighting the dataset name and clicking the *Edit Options* button in *Storage* \rightarrow *Volumes*.

Setting	Value	Description	
Dataset Name	string Enter a mandatory unique name for the dataset.		
Comments	string	Enter optional comments or notes about this dataset.	
Sync	drop-down menu	Sets the data write synchronization. <i>Inherit</i> inherits the sync settings from the parent dataset. <i>Always</i> always waits. <i>Standard</i> uses the sync settings that are requested by the client software for data writes to complete. <i>Disabled</i> never waits for writes to complete.	

Table 8.4: ZFS Dataset Options

Continued on next page

Setting	Value	Description
Compression Level	drop-down	Refer to the section on <i>Compression</i> (page 135) for a description of the
	menu	available algorithms.
Share type	drop-down	Select the type of share that will be used on the dataset. Choices are UNIX
	menu	for an NFS share, <i>Windows</i> for a SMB share, or <i>Mac</i> for an AFP share.
Enable atime	Inherit, On,	Choose On to update the access time for files when they are read. Choose
	or Off	Off to prevent producing log traffic when reading files. This can result in
		significant performance gains.
Quota for this dataset	integer	Only available in Advanced Mode. Default of 0 disables quotas. Specify-
		ing a value uses no more than the specified size and is suitable for user
		datasets to prevent users from taking all available space.
Quota for this dataset and	integer	Only available in <i>Advanced Mode</i> . A specified value applies to both this
all children		dataset and any child datasets.
Reserved space for this	integer	Only available in <i>Advanced Mode</i> . Default of <i>0</i> is unlimited. Specifying a
dataset		value keeps at least this much space free and is suitable for datasets with
		logs that could take all free space.
Reserved space for this	integer	Only available in Advanced Mode. A specified value applies to both this
dataset and all children		dataset and any child datasets.
ZFS Deduplication	drop-down	Read the section on <i>Deduplication</i> (page 134) before making a change to
	menu	this setting.
Read-Only	drop-down	Only available in Advanced Mode. Choices are Inherit (off), On, or Off.
	menu	
Exec	drop-down	Only available in Advanced Mode. Choices are Inherit (on), On, or Off. Set-
	menu	ting to <i>Off</i> prevents the installation of <i>Plugins</i> (page 254) or <i>Jails</i> (page 256).
Record Size	drop-down	Only available in Advanced Mode. While ZFS automatically adapts the
	menu	record size dynamically to adapt to data, if the data has a fixed size,
		matching that size can result in better performance.
Case Sensitivity	drop-down	Sensitive is the default and assumes filenames are case sensitive. Insen-
	menu	sitive assumes filenames are not case sensitive. Mixed understands both
		types of filenames.

Table	84 –	continued	from	previous	nage
Iable	0.4 -	continueu	nom	previous	page

Create a nested dataset by clicking on an existing dataset and selecting *Create Dataset*. A zvol can also be created within a dataset.

8.1.3.1 Deduplication

Deduplication is the process of ZFS transparently reusing a single copy of duplicated data to save space. Depending on the amount of duplicate data, deduplicaton can improve storage capacity, as less data is written and stored. However, deduplication is RAM intensive. A general rule of thumb is 5 GiB of RAM per terabyte of deduplicated storage. **In most cases, compression provides storage gains comparable to deduplication with less impact on performance.**

In FreeNAS[®], deduplication can be enabled during dataset creation. Be forewarned that **there is no way to undedup the data within a dataset once deduplication is enabled**, as disabling deduplication has **NO EFFECT** on existing data. The more data written to a deduplicated dataset, the more RAM it requires. When the system starts storing the DDTs (dedup tables) on disk because they no longer fit into RAM, performance craters. Further, importing an unclean pool can require between 3-5 GiB of RAM per terabyte of deduped data, and if the system does not have the needed RAM, it will panic. The only solution is to add more RAM or recreate the pool. **Think carefully before enabling dedup!** This article (https://constantin.glez.de/2011/07/27/zfs-to-dedupe-or-not-dedupe/) provides a good description of the value versus cost considerations for deduplication.

Unless a lot of RAM and a lot of duplicate data is available, do not change the default deduplication setting of "Off". For performance reasons, consider using compression rather than turning this option on.

If deduplication is changed to *On*, duplicate data blocks are removed synchronously. The result is that only unique data is stored and common components are shared among files. If deduplication is changed to *Verify*, ZFS will do a byte-to-byte comparison when two blocks have the same signature to make sure that the block contents are identical. Since hash collisions are extremely rare, *Verify* is usually not worth the performance hit.

Note: After deduplication is enabled, the only way to disable it is to use the zfs set dedup=off dataset_name command from *Shell* (page 289). However, any data that has already been deduplicated will not be un-deduplicated. Only newly stored data after the property change will not be deduplicated. The only way to remove existing deduplicated data is to copy all of the data off of the dataset, set the property to off, then copy the data back in again. Alternately, create a new dataset with *ZFS Deduplication* left disabled, copy the data to the new dataset, and destroy the original dataset.

Tip: Deduplication is often considered when using a group of very similar virtual machine images. However, other features of ZFS can provide dedup-like functionality more efficiently. For example, create a dataset for a standard VM, then clone a snapshot of that dataset for other VMs. Only the difference between each created VM and the main dataset are saved, giving the effect of deduplication without the overhead.

8.1.3.2 Compression

When selecting a compression type, try to balance performance with the amount of disk space saved by compression. Compression is transparent to the client and applications as ZFS automatically compresses data as it is written to a compressed dataset or zvol and automatically decompresses that data as it is read. These compression algorithms are supported:

- **Iz4:** default and recommended compression method as it allows compressed datasets to operate at near real-time speed. This algorithm only compresses the files that will benefit from compression.
- **gzip:** varies from levels 1 to 9 where *gzip fastest* (level 1) gives the least compression and *gzip maximum* (level 9) provides the best compression but is discouraged due to its performance impact.
- zle: fast but simple algorithm which eliminates runs of zeroes.
- Izjb: provides decent data compression, but is considered deprecated as Iz4 provides much better performance.

If selecting *Off* as the *Compression level* when creating a dataset or zvol, compression will not be used on that dataset/zvol. This is not recommended as using *lz4* has a negligible performance impact and allows for more storage capacity.

8.1.4 Create zvol

A zvol is a feature of ZFS that creates a raw block device over ZFS. The zvol can be used as an *iSCSI* (page 230) device extent.

To create a zvol, select an existing ZFS volume or dataset from the tree then click *Create zvol* to open the screen shown in Figure 8.5.

Create zvol	8
Create zvol on volume1	
zvol name:	
Comments:	
Size for this zvol:	(i)
Force size:	<u>(</u>
Sync:	Inherit (standard) 💌
Compression level:	Inherit (Iz4)
ZFS Deduplication:	Enabling dedup can drastically reduce performance and affect the ability to access data. Compression usually offers similar space savings with much lower performance impact and overhead.
Sparse volume:	
Add zvol Cancel	Advanced Mode

Fig. 8.5: Creating a Zvol

The configuration options are described in Table 8.5. Some settings are only available in *Advanced Mode*. To see these settings, either click the *Advanced Mode* button or configure the system to always display these settings by enabling *Show* advanced fields by default in System \rightarrow Advanced.

Setting	Value	Description
zvol Name	string	Enter a short name for the zvol. Using a zvol name longer than 63- characters can prevent accessing zvols as devices. For example, a zvol with a 70-character filename or path cannot be used as an iSCSI extent. This setting is mandatory.
Comments	string	Enter any notes about this zvol.
Size for this zvol	integer	Specify size and value such as <i>10Gib</i> . If the size is more than 80% of the available capacity, the creation will fail with an "out of space" error unless <i>Force size</i> is also enabled.
Force size	checkbox	By default, the system does not create a zvol when it brings the pool above 80% capacity. While NOT recommended , enabling this option will force the creation of the zvol.

Table 8.5: zvol Configuration Option	IS
--------------------------------------	----

Continued on next page

Setting	Value	Description
Compression level	drop-down	Refer to the section on <i>Compression</i> (page 135) for a description of the
	menu	available algorithms.
Sparse volume	checkbox	Used to provide thin provisioning. Caution: when this option is set, writes
		will fail when the pool is low on space.
Block size	drop-down	Only available in <i>Advanced Mode</i> . The default is based on the number of
	menu	disks in the pool. Can be set to match the block size of the filesystem to be
		formatted onto the iSCSI target.

Table 8.5 – continued from previous page

8.1.5 Import Disk

The *Volume* \rightarrow *Import Disk* screen, shown in Figure 8.6, is used to import a **single** disk that has been formatted with the UFS, NTFS, MSDOS, or EXT2 filesystem. The import is meant to be a temporary measure to copy the data from a disk to an existing ZFS dataset. Only one disk can be imported at a time.

Note: Imports of EXT3 or EXT4 filesystems are possible in some cases, although neither is fully supported. EXT3 journaling is not supported, so those filesystems must have an external *fsck* utility, like the one provided by E2fsprogs utilities (http://e2fsprogs.sourceforge.net/), run on them before import. EXT4 filesystems with extended attributes or inodes greater than 128 bytes are not supported. EXT4 filesystems with EXT3 journaling must have an *fsck* run on them before import, as described above.

Import Disk	_	×
Member disk:	ada0p2 (3.0 TB) 💌 🤃	
File System type:	 UFS NTFS MSDOSFS EXT2FS 	
MSDOSFS locale:	Default	
Destination:	() ()	Browse
Import Disk Cance	el	

Fig. 8.6: Importing a Disk

Use the drop-down menu to select the disk to import, select the type of filesystem on the disk, and browse to the ZFS dataset that will hold the copied data. If the *MSDOSFS* filesystem is selected, the *MSDOSFS locale* drop-down menu can be used to select the locale when non-ascii characters are present on the disk.

Once *Import Disk* is clicked, the disk is mounted, its contents are copied to the specified ZFS dataset, and the disk is unmounted after the copy operation completes.

8.1.6 Import Volume

Click *Storage* \rightarrow *Volumes* \rightarrow *Import Volume*, to configure FreeNAS[®] to use an **existing** ZFS pool. This action is typically performed when an existing FreeNAS[®] system is re-installed. Since the operating system is separate from the storage disks, a new installation does not affect the data on the disks. However, the new operating system needs to be configured to use the existing volume.

Figure 8.7 shows the initial pop-up window that appears when a volume is imported.

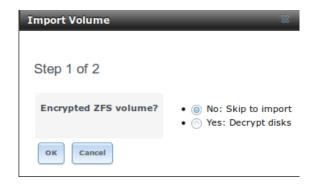


Fig. 8.7: Initial Import Volume Screen

If importing an unencrypted ZFS pool, select No: Skip to import to open the screen shown in Figure 8.8.

Import Vol	ume	8
Step 2 of	f 2	
		_
Volume:	volume1 [zfs, id=1929756524230885343]	-
ок	ancel	

Fig. 8.8: Importing a Non-Encrypted Volume

Existing volumes are available for selection from the drop-down menu. In the example shown in Figure 8.8, the FreeNAS[®] system has an existing, unencrypted ZFS pool. Once the volume is selected, click the *OK* button to import the volume.

If an existing ZFS pool does not show in the drop-down menu, run zpool import from Shell (page 289) to import the pool.

If physically installing ZFS formatted disks from another system, ensure to export the drives on that system to prevent an "in use by another machine" error during the import.

If the hardware is not being detected, run camcontrol devlist from *Shell* (page 289). If the disk does not appear in the output, check to see if the controller driver is supported or if it needs to be loaded using *Tunables* (page 77).

8.1.6.1 Importing an Encrypted Pool

Disks in existing GELI-encrypted ZFS pools must be decrypted before importing the pool. In the Import Volume dialog shown in Figure 8.7, select *Yes: Decrypt disks*. The screen shown in Figure 8.9 is then displayed.

Import Volume		36
Step 2 of 3		
Disks:	ada2p2 ada1p2	
Encryption Key:	Browse No file selected.	
Passphrase:		
ОК Cancel		

Fig. 8.9: Decrypting Disks Before Importing a ZFS Pool

Select the disks in the encrypted pool, browse to the location of the saved encryption key, enter the passphrase associated with the key, then click *OK* to decrypt the disks.

Note: The encryption key is required to decrypt the pool. If the pool cannot be decrypted, it cannot be re-imported after a failed upgrade or lost configuration. This means that it is **very important** to save a copy of the key and to remember the passphrase that was configured for the key. Refer to *Managing Encrypted Volumes* (page 143) for instructions on how to manage the keys for encrypted volumes.

After the pool is decrypted, it appears in the drop-down menu of Figure 8.8. Click the OK button to finish the volume import.

Note: For security reasons, GELI keys for encrypted volumes are not saved in a configuration backup file. When FreeNAS[®] has been installed to a new device and a saved configuration file restored to it, the GELI keys for encrypted disks will not be present, and the system will not request them. To correct this, export the encrypted volume with Detach Volume, making sure that the options *Mark the disks as new (destroy data)* or *Also delete the share's configuration* are **not** selected. Then import the volume again. During the import, the GELI keys can be entered as described above.

8.1.7 View Disks

Storage \rightarrow Volumes \rightarrow View Disks shows all of the disks recognized by the FreeNAS[®] system. An example is shown in Figure 8.10.

	al	Disk Size	Description	Transfer Mode	HDD Standby	Advanced Power Management	Acoustic Level	Enable S.M.A.R.T.	S.M.A.R.T. extra options	Password for SED
a0 1620:	012B2CB65	128.0 GB		Auto	Always On	Disabled	Disabled	true		
1 16203	012B2DA4C	128.0 GB		Auto	Always On	Disabled	Disabled	true		
la2 WD-V	WCC4M1AHDRR9	2.0 TB		Auto	Always On	Disabled	Disabled	true		
a3 WD-V	WCC4M1AHD6HN	2.0 TB		Auto	Always On	Disabled	Disabled	true		
a4 WD-V	WCC4M3SKL8R4	2.0 TB		Auto	Always On	Disabled	Disabled	true		
da5 WD-V	-WCC4M3DFZZL6	2.0 TB		Auto	Always On	Disabled	Disabled	true		

Fig. 8.10: Viewing Disks

The current configuration of each device is displayed. Click a disk entry and the *Edit* button to change its configuration. The configurable options are described in Table 8.6.

To bulk edit disks, hold Shift and click each disk to edit. *Edit* changes to *Edit In Bulk*. Click it to open the *Edit In Bulk* window. This window displays which disks are being edited and a short list of configurable options. The *Disk Options table* (page 140) indicates the options available when editing multiple disks.

Setting	Value	Bulk Edit	Description
Name	string		This is the FreeBSD device name for the disk.
Serial	string		This is the serial number of the disk.
Description	string		Enter any notes about this disk.
HDD Standby	drop-	\checkmark	Indicates the time of inactivity in minutes before the drive
	down		enters standby mode to conserve energy. This forum post
	menu		(https://forums.freenas.org/index.php?threads/how-to-find-out-if-a-
			drive-is-spinning-down-properly.2068/) demonstrates how to deter- mine if a drive has spun down.
Advanced Power	drop		Select a power management profile from the menu. The default value
	drop- down	\checkmark	is Disabled.
Management			IS DISUDIEU.
Acoustic Level	menu		Default is <i>Disabled</i> . Other values can be
ACOUSTIC LEVEL	drop-	\checkmark	
	down		selected for disks that understand AAM
	menu		(https://en.wikipedia.org/wiki/Automatic_acoustic_management).
Enable S.M.A.R.T.	checkbox	\checkmark	Enabled by default if the disk supports S.M.A.R.T. Unsetting this option will disable any configured <i>S.M.A.R.T. Tests</i> (page 110) for the disk.
S.M.A.R.T. extra op-	string	\checkmark	Enter additional smartctl(8)
tions			(https://www.smartmontools.org/browser/trunk/smartmontools/smartctl.8.in options.
Password for SED	string		Enter and confirm the password which will be used for this device
	U		instead of the global SED password. Refer to <i>Self-Encrypting Drives</i>
			(page 74) for more information.
	1	1	

Table 8.6:	Disk (Options
------------	--------	---------

Continued on next page

Table 8.6 – continued from previous page					
Setting	Value	Bulk Edit	Description		
Reset Password	checkbox		Set to clear the SED password.		

Note: If the serial number of a disk is not displayed in this screen, use the smartctl command from *Shell* (page 289). For example, to determine the serial number of disk *ada0*, type smartctl -a /dev/ada0 | grep Serial.

The *Wipe* function is provided for when an unused disk is to be discarded.

Warning: Make certain that all data has been backed up and that the disk is no longer in use. Triple-check that the correct disk is being selected to be wiped, as recovering data from a wiped disk is usually impossible. If there is any doubt, physically remove the disk, verify that all data is still present on the FreeNAS[®] system, and wipe the disk in a separate computer.

Clicking *Wipe* offers several choices. *Quick* erases only the partitioning information on a disk, making it easy to reuse but without clearing other old data. For more security, *Full with zeros* overwrites the entire disk with zeros, while *Full with random data* overwrites the entire disk with random binary data.

Quick wipes take only a few seconds. A *Full with zeros* wipe of a large disk can take several hours, and a *Full with random data* takes longer. A progress bar is displayed during the wipe to track status.

8.1.8 Volumes

Storage \rightarrow Volumes is used to view and further configure existing ZFS pools, datasets, and zvols. The example shown in Figure 8.11 shows one ZFS pool (volume1) with two datasets (the one automatically created with the pool, volume1, and dataset1) and one zvol (zvol1).

Note that in this example, there are two datasets named *volume1*. The first represents the ZFS pool and its *Used* and *Available* entries reflect the total size of the pool, including disk parity. The second represents the implicit or root dataset and its *Used* and *Available* entries indicate the amount of disk space available for storage.

Buttons are provided for quick access to *Volume Manager, Import Disk, Import Volume*, and *View Disks*. If the system has multipath-capable hardware, an extra button will be added, *View Multipaths*. For each entry, the columns indicate the *Name*, how much disk space is *Used*, how much disk space is *Available*, the type of *Compression*, the *Compression Ratio*, the *Status*, whether it is mounted as read-only, and any *Comments* entered for the volume.

Storage Volumes Periodic Snapshot Tasks Replication Tasks Resilver Priority Scrubs Snapshots VMware-Snapshot									
Volume Manager Import Volume View Disks									
Name	Used	Available	Compression	Compression Ratio	Status	Readonly	Comments		
⊿ volume1	2.7 MiB (0%)	7.9 GiB	-	-	HEALTHY				
volumel	1.1 MiB (0%)	7.7 GiB	lz4	3.08x	-	inherit (off)			

Fig. 8.11: Viewing Volumes

Clicking the entry for a pool causes several buttons to appear at the bottom of the screen.

Detach Volume: allows exporting the pool or deleting the contents of the pool, depending upon the choice made in the screen shown in Figure 8.12. The *Detach Volume* screen displays the current used space and indicates whether there are any shares. It provides options to *Mark the disks as new (destroy data)* and *Also delete the share's configuration*. The browser window turns red to indicate that some choices will make the data inaccessible.**When the option to select the disks as new is left deselected, the volume is exported.** The data is not destroyed and the volume can be re-imported at a later time. When moving a ZFS pool from one system to another, perform this export action first as it flushes any unwritten data to disk, writes data to the disk indicating that the export was done, and removes all knowledge of the pool from the system.

When the option to mark the disks as new is selected, the pool and all the data in its datasets, zvols, and shares is destroyed and the individual disks are returned to their raw state. Desired data must be backed up to another disk or device before using this option.

Detach Volume	ж
You have 2.5 MiB of used space within this volu Mark the disks as new (destroy data):	me
volume1: Are you sure you want to detach? Yes Cancel	

Fig. 8.12: Detach or Delete a Volume

Scrub Volume: scrubs and scheduling them are described in more detail in *Scrubs* (page 162). This button allows manually initiating a scrub. Scrubs are I/O intensive and can negatively impact performance. Avoid initiating a scrub when the system is busy.

A *Cancel* button is provided to cancel a scrub. When a scrub is cancelled, it is abandoned. The next scrub to run starts from the beginning, not where the cancelled scrub left off.

The status of a running scrub or the statistics from the last completed scrub can be seen by clicking the Volume Status button.

Volume Status: as shown in the example in Figure 8.13, this screen shows the device name and status of each disk in the ZFS pool as well as any read, write, or checksum errors. It also indicates the status of the latest ZFS scrub. Clicking the entry for a device causes buttons to appear to edit the device options (shown in Figure 8.14), offline or online the device, or replace the device (as described in *Replacing a Failed Drive* (page 147)).

Upgrade: used to upgrade the pool to the latest ZFS features, as described in *Upgrading a ZFS Pool* (page 35). This button does not appear if the pool is running the latest version of feature flags.

Volume Status							
Scrub Status: Completed Errors: 0 Repaired: 0 Date: Mon Oct 16 13:10:08 2017							
Name	Read	Write	Checksum	Status			
⊿ volume1	0	0	0	ONLINE			
⊿ raidz1-0	0	0	0	ONLINE			
ada3p2	0	0	0	ONLINE			
ada2p2	0	0	0	ONLINE			
adalp2	0	0	0	ONLINE			

Fig. 8.13: Volume Status

Selecting a disk in *Volume Status* and clicking its *Edit Disk* button shows the screen in Figure 8.14. Table 8.6 summarizes the configurable options.

E	Edit				
	Name:	ada0			
	Serial:	JP2940HZ3SNPDC			
	Description:				
	HDD Standby:	Always On 👻			
	Advanced Power Management:	Disabled v			
	Acoustic Level:	Disabled v			
	Enable S.M.A.R.T.				
	S.M.A.R.T. extra options:				
	OK Cancel				



Note: Versions of FreeNAS[®] prior to 8.3.1 required a reboot to apply changes to the *HDD Standby*, *Advanced Power Management*, and *Acoustic Level* settings. As of 8.3.1, changes to these settings are applied immediately.

Clicking a dataset in *Storage* \rightarrow *Volumes* causes buttons to appear at the bottom of the screen, providing these options:

Change Permissions: edit the dataset permissions as described in Change Permissions (page 130).

Create Snapshot: create a one-time snapshot. To schedule the regular creation of snapshots, instead use *Periodic Snapshot Tasks* (page 150).

Promote Dataset: only applies to clones. When a clone is promoted, the origin filesystem becomes a clone of the clone making it possible to destroy the filesystem that the clone was created from. Otherwise, a clone cannot be deleted while the origin filesystem exists.

Destroy Dataset: clicking the *Destroy Dataset* button causes the browser window to turn red to indicate that this is a destructive action. Clicking *Yes* proceeds with the deletion.

Edit Options: edit the volume properties described in Table 8.4. Note that it will not allow changing the dataset name.

Create Dataset: used to create a child dataset within this dataset.

Create zvol: create a child zvol within this dataset.

Clicking a zvol in *Storage* \rightarrow *Volumes* causes icons to appear at the bottom of the screen: *Create Snapshot, Edit zvol*, and *Destroy zvol*. Similar to datasets, a zvol name cannot be changed, and destroying a zvol requires confirmation.

8.1.8.1 Managing Encrypted Volumes

FreeNAS[®] generates and stores a randomized *encryption key* whenever a new encrypted volume is created. This key is required to read and decrypt any data on the volume.

Encryption keys can also be downloaded as a safety measure, to allow decryption on a different system in the event of failure, or to allow the locally stored key to be deleted for extra security. Encryption keys can also be optionally protected with a *passphrase* for additional security. The combination of encryption key location and whether a passphrase is used provide several different security scenarios:

- *Key stored locally, no passphrase*: the encrypted volume is decrypted and accessible when the system running. Protects "data at rest" only.
- *Key stored locally, with passphrase*: the encrypted volume is not accessible until the passphrase is entered by the FreeNAS[®] administrator.
- *Key not stored locally*: the encrypted volume is not accessible until the FreeNAS[®] administrator provides the key. If a passphrase is set on the key, it must also be entered before the encrypted volume can be accessed (two factor authentication (https://en.wikipedia.org/wiki/Multi-factor_authentication)).

Encrypted data cannot be accessed when the disks are removed or the system has been shut down. On a running system, encrypted data cannot be accessed when the volume is locked (see below) and the key is not available. If the key is protected with a passphrase, both the key and passphrase are required for decryption.

Encryption applies to a volume, not individual users. When a volume is unlocked, data is accessible to all users with permissions to access it.

Note: GELI (https://www.freebsd.org/cgi/man.cgi?query=geli) uses *two* randomized encryption keys for each disk. The first has been discussed here. The second, the disk's "master key", is encrypted and stored in the on-disk GELI metadata. Loss of a disk master key due to disk corruption is equivalent to any other disk failure, and in a redundant pool, other disks will contain accessible copies of the uncorrupted data. While it is *possible* to separately back up disk master keys, it is usually not necessary or useful.

8.1.8.2 Additional Controls for Encrypted Volumes

If the *Encryption* option is enabled during the creation of a pool, additional buttons appear in the entry for the volume in *Storage* \rightarrow *Volumes*. An example is shown in Figure 8.15.

Storage Volumes Periodic Snapshot Tasks Replication Tasks Resilver Priority Scrubs Snapshots Volume Hanager Import Disk Import Volume								
Name	Used	Available	Compression	Compression Ratio	Status	Readonly	Comments	
⊿ volume1	2.7 MiB (0%)	7.9 GiB	-	-	HEALTHY			
volumel	1.1 MiB (0%)	7.7 GiB	lz4	1.72x	-	inherit (off)		
				·				

Fig. 8.15: Encryption Icons Associated with an Encrypted Volume

These additional encryption buttons are used to:

Create/Change Passphrase: set and confirm a passphrase associated with the GELI encryption key. The desired passphrase is entered and repeated for verification. A red warning is a reminder to *Remember to add a new recovery key as this action invalidates the previous recovery key*. Unlike a password, a passphrase can contain spaces and is typically a series of words. A good passphrase is easy to remember (like the line to a song or piece of literature) but hard to guess. **Remember this passphrase.** An encrypted volume cannot be reimported without it. In other words, if the passphrase is forgotten, the data on the volume can become inaccessible if it becomes necessary to reimport the pool. Protect this passphrase, as anyone who knows it could reimport the encrypted volume, thwarting the reason for encrypting the disks in the first place.

Create Passphrase 🕺						
Remember to add a new recovery key	as this action invalidates the previous recovery key					
Passphrase:	•••••					
Confirm Passphrase:	•••••					
OK Cancel						

Fig. 8.16: Add or Change a Passphrase to an Encrypted Volume

After the passphrase is set, the name of this button changes to *Change Passphrase*. After setting or changing the passphrase, it is important to *immediately* create a new recovery key by clicking the *Add recovery key* button. This way, if the passphrase is forgotten, the associated recovery key can be used instead.

Encrypted volumes with a passphrase display an additional lock button:

≡°

Fig. 8.17: Lock Button

These encrypted volumes can be *locked*. The data is not accessible until the volume is unlocked by suppying the passphrase or encryption key, and the button changes to an unlock button:

O1 F

Fig. 8.18: Unlock Button

To unlock the volume, click the unlock button to display the Unlock dialog:

Unlock						
Passphrase:						
Recovery Key:	Browse No file selected.					
Restart services:	 AFP CIFS FTP ISCSI NFS WebDAV Jails/Plugins 					
OK Cancel						

Fig. 8.19: Unlock Locked Volume

Unlock the volume by entering a passphrase *or* using the *Browse* button to load the recovery key. Only the passphrase is used when both a passphrase and a recovery key are entered. The services listed in *Restart Services* will restart when the pool is unlocked. This allows them to see the new volume and share or access data on it. Individual services can be prevented from restarting by deselecting them. However, a service that is not restarted might not be able to access the unlocked volume.

Download Key: download a backup copy of the GELI encryption key. The encryption key is saved to the client system, not on the FreeNAS[®] system. The FreeNAS[®] administrative password must be entered, then the directory in which to store the key is chosen. Since the GELI encryption key is separate from the FreeNAS[®] configuration database, **it is highly recommended to make a backup of the key. If the key is ever lost or destroyed and there is no backup key, the data on the disks is inaccessible.**

Encryption Re-key: generate a new GELI encryption key. Typically this is only performed when the administrator suspects that the current key may be compromised. This action also removes the current passphrase.

Add recovery key: generate a new recovery key. This screen prompts for the FreeNAS[®] administrative password and then the directory in which to save the key. Note that the recovery key is saved to the client system, not on the FreeNAS[®] system. This recovery key can be used if the passphrase is forgotten. **Always immediately add a recovery key whenever the passphrase is changed.**

Remove recovery key: Typically this is only performed when the administrator suspects that the current recovery key may be compromised. **Immediately** create a new passphrase and recovery key.

Note: The passphrase, recovery key, and encryption key must be protected. Do not reveal the passphrase to others. On the system containing the downloaded keys, take care that the system and its backups are protected. Anyone who has the keys has the ability to re-import the disks if they are discarded or stolen.

Warning: If a re-key fails on a multi-disk system, an alert is generated. **Do not ignore this alert** as doing so may result in the loss of data.

8.1.9 View Multipaths

FreeNAS[®] uses gmultipath(8) (https://www.freebsd.org/cgi/man.cgi?query=gmultipath) to provide multipath I/O (https://en.wikipedia.org/wiki/Multipath_I/O) support on systems containing hardware that is capable of multipath.

An example would be a dual SAS expander backplane in the chassis or an external JBOD.

Multipath hardware adds fault tolerance to a NAS as the data is still available even if one disk I/O path has a failure.

FreeNAS[®] automatically detects active/active and active/passive multipath-capable hardware. Any multipath-capable devices that are detected will be placed in multipath units with the parent devices hidden. The configuration will be displayed in *Storage* \rightarrow *Volumes* \rightarrow *View Multipaths*. Note that this option is not be displayed in the *Storage* \rightarrow *Volumes* tree on systems that do not contain multipath-capable hardware.

8.1.10 Replacing a Failed Drive

With any form of redundant RAID, failed drives must be replaced as soon as possible to repair the degraded state of the RAID. Depending on the hardware capabilities, it might be necessary to reboot to replace the failed drive. Hardware that supports AHCI does not require a reboot.

Note: Striping (RAID0) does not provide redundancy. If a disk in a stripe fails, the volume will be destroyed and must be recreated and the data restored from backup.

Note: If the volume is encrypted with GELI, refer to Replacing an Encrypted Drive (page 149) before proceeding.

Before physically removing the failed device, go to *Storage* \rightarrow *Volumes*. Select the volume name. At the bottom of the interface are several icons, one of which is *Volume Status*. Click the *Volume Status* icon and locate the failed disk. Then perform these steps:

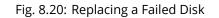
1. Click the disk entry, then its *Offline* button to change the disk status to OFFLINE. This step removes the device from the ZFS pool and prevents swap issues. If the hardware supports hot-pluggable disks, click the disk *Offline* button and pull the disk, then skip to step 3. If there is no *Offline* button but only a *Replace* button, the disk is already offlined and this step can be skipped.

Note: If the process of changing the disk status to OFFLINE fails with a "disk offline failed - no valid replicas" message, the ZFS volume must be scrubbed first with the *Scrub Volume* button in *Storage* \rightarrow *Volumes*. After the scrub completes, try to *Offline* the disk again before proceeding.

- 2. If the hardware is not AHCI capable, shut down the system to physically replace the disk. When finished, return to the GUI and locate the OFFLINE disk.
- 3. After the disk has been replaced and is showing as OFFLINE, click the disk again and then click its *Replace* button. Select the replacement disk from the drop-down menu and click the *Replace Disk* button. After clicking the *Replace Disk* button, the ZFS pool begins resilvering.
- 4. After the drive replacement process is complete, re-add the replaced disk in the S.M.A.R.T. Tests (page 110) screen.

In the example shown in Figure 8.20, a failed disk is being replaced by disk *ada5* in the volume named volume1.

Name	Read	Write	Checksum	Status				
∡ volume1	0	0	0	DEGRADED				
∡ raidz1-0	0	0	0	DEGRADED				
ada4p2	0	0	0	ONLINE				
ada3p2	0	0	0	ONLINE				
ada2p2	0	0	0	ONLINE				
1959638268805654949	0	0	0	OFFLINE				
				k 1959638268805654949 k: ada5 (10.7 GB)				
	Replace Disk Cancel							



After the resilver is complete, *Volume Status* shows a *Completed* resilver status and indicates any errors. Figure 8.21 indicates that the disk replacement was successful in this example.

Note: A disk that is failing but has not completely failed can be replaced in place, without first removing it. Whether this is a good idea depends on the overall condition of the failing disk. A disk with a few newly-bad blocks that is otherwise functional can be left in place during the replacement to provide data redundancy. A drive that is experiencing continuous errors can actually slow down the replacement. In extreme cases, a disk with serious problems might spend so much time retrying failures that it could prevent the replacement resilvering from completing before another drive fails.

Desilion							
Resilver							
Status: Completed							
Errors: 0 Date: Fri Aug 29 11:22:39 2014							
Name Read Write Checksum Status							
⊿ volume1	0	0	0	ONLINE			
⊿ raidz1-0	0	0	0	ONLINE			
ada4p2	0	0	0	ONLINE			
ada3p2	0	0	0	ONLINE			
ada2p2	0	0	0	ONLINE			
ada5p2	0	0	0	ONLINE			

Fig. 8.21: Disk Replacement is Complete

8.1.10.1 Replacing an Encrypted Drive

If the ZFS pool is encrypted, additional steps are needed when replacing a failed drive.

First, make sure that a passphrase has been set using the instructions in *Encryption* (page 127) **before** attempting to replace the failed drive. Then, follow the steps 1 and 2 as described above. During step 3, a prompt will appear to input and confirm the passphrase for the pool. Enter this information then click the *Replace Disk* button. Wait until the resilvering is complete.

Next, restore the encryption keys to the pool. If this additional step is not performed before the next reboot, access to the pool might be permanently lost.

1. Highlight the pool that contains the disk that was just replaced and click the *Add Recovery Key* button to save the new recovery key. The old recovery key will no longer function, so it can be safely discarded.

8.1.10.2 Removing a Log or Cache Device

Added log or cache devices appear in *Storage* \rightarrow *Volumes* \rightarrow *Volume Status*. Clicking the device enables its *Replace* and *Remove* buttons.

Log and cache devices can be safely removed or replaced with these buttons. Both types of devices improve performance, and throughput can be impacted by their removal.

8.1.11 Replacing Drives to Grow a ZFS Pool

The recommended method for expanding the size of a ZFS pool is to pre-plan the number of disks in a vdev and to stripe additional vdevs using *Volume Manager* (page 125) as additional capacity is needed.

However, this is not an option if there are no open drive ports and a SAS/SATA HBA card cannot be added. In this case, one disk at a time can be replaced with a larger disk, waiting for the resilvering process to incorporate the new disk into the pool, then repeating with another disk until all of the original disks have been replaced.

The safest way to perform this is to use a spare drive port or an eSATA port and a hard drive dock. The process follows these steps:

- 1. Shut down the system.
- 2. Install one new disk.
- 3. Start up the system.
- 4. Go to Storage \rightarrow Volumes, select the pool to expand and click the Volume Status button. Select a disk and click the Replace button. Choose the new disk as the replacement.

5. The status of the resilver process can be viewed by running <code>zpool status</code>. When the new disk has resilvered, the old one will be automatically offlined. The system is then shut down to physically remove the replaced disk. One advantage of this approach is that there is no loss of redundancy during the resilver.

If a spare drive port is not available, a drive can be replaced with a larger one using the instructions in *Replacing a Failed Drive* (page 147). This process is slow and places the system in a degraded state. Since a failure at this point could be disastrous, **do not attempt this method unless the system has a reliable backup.** Replace one drive at a time and wait for the resilver process to complete on the replaced drive before replacing the next drive. After all the drives are replaced and the final resilver completes, the added space will appear in the pool.

8.1.12 Hot Spares

ZFS provides the ability to have "hot" *spares*. These are drives that are connected to a volume, but not in use. If the volume experiences the failure of a data drive, the system uses the hot spare as a temporary replacement. If the failed drive is replaced with a new drive, the hot spare drive is no longer needed and reverts to being a hot spare. If the failed drive is instead removed from the volume, the spare is promoted to a full member of the volume.

Hot spares can be added to a volume during or after creation. On FreeNAS[®], hot spare actions are implemented by zfsd(8) (https://www.freebsd.org/cgi/man.cgi?query=zfsd).

8.2 Periodic Snapshot Tasks

A periodic snapshot task allows scheduling the creation of read-only versions of ZFS volumes and datasets at a given point in time. Snapshots can be created quickly and, if little data changes, new snapshots take up very little space. For example, a snapshot where no files have changed takes 0 MiB of storage, but as changes are made to files, the snapshot size changes to reflect the size of the changes.

Snapshots provide a clever way of keeping a history of files, providing a way to recover an older copy or even a deleted file. For this reason, many administrators take snapshots often (perhaps every fifteen minutes), store them for a period of time (possibly a month), and store them on another system (typically using *Replication Tasks* (page 152)). Such a strategy allows the administrator to roll the system back to a specific point in time. If there is a catastrophic loss, an off-site snapshot can be used to restore the system up to the time of the last snapshot.

An existing ZFS volume is required before creating a snapshot. Creating a volume is described in Volume Manager (page 125).

To create a periodic snapshot task, click Storage \rightarrow Periodic Snapshot Tasks \rightarrow Add Periodic Snapshot which opens the screen shown in Figure 8.22. Table 8.7 summarizes the fields in this screen.

Note: If only a one-time snapshot is needed, instead use $Storage \rightarrow Volumes$ and click the *Create Snapshot* button for the volume or dataset to snapshot.

P	eriodic Snapshots	X
	Volume/Dataset:	volumel
	Recursive:	
	Snapshot Lifetime	2 Week(s) Veek(s)
	Begin:	09:00:00 👻 (i)
	End:	18:00:00 🔹 i
	Interval:	1 hour 💌 i
	Weekday:	 Monday Tuesday Wednesday Thursday Friday Saturday Sunday
	Enabled:	
	OK Cancel	

Fig. 8.22: Creating a Periodic Snapshot

Table 8.7: Options	When Creating a	Periodic Snapshot
--------------------	-----------------	-------------------

Setting	Value	Description	
Volume/Dataset	drop-down menu	Select an existing ZFS volume, dataset, or zvol.	
Recursive	checkbox	Set to take separate snapshots of the volume or dataset and each of its child datasets. Unset to take a single snapshot of only the specified volume or dataset.	
Snapshot Life- time	integer and drop- down menu	Define a length of time to retain the snapshot on this system. After the time expires, the snapshot is removed. Snapshots replicated to other systems are not affected.	
Begin	drop-down menu	Choose the hour and minute when the system can begin taking snap- shots.	
End	drop-down menu	Choose the hour and minute when the system will stop taking snapshots.	
Interval	drop-down menu	Define how often the system takes snapshots between the <i>Begin</i> and <i>End</i> times.	
Weekday	checkboxes	Choose the days of the week to take snapshots.	
Enabled	checkbox	Unset to disable this task without deleting it.	

If the *Recursive* option is enabled, child datasets of this dataset are included in the snapshot and there is no need to create snapshots for each child dataset. The downside is that there is no way to exclude particular child datasets from a recursive snapshot.

Click the OK button to save the task. Entries for each task are shown in View Periodic Snapshot Tasks. Click an entry to display *Edit* and *Delete* buttons for it.

8.3 Replication Tasks

Replication is the duplication of snapshots from one FreeNAS[®] system to another computer. When a new snapshot is created on the source computer, it is automatically replicated to the destination computer. Replication is typically used to keep a copy of files on a separate system, with that system sometimes being at a different physical location.

The basic configuration requires a source system with the original data and a destination system where the data will be replicated. The destination system is prepared to receive replicated data, a *periodic snapshot* (page 150) of the data on the source system is created, and then a replication task is created. As snapshots are automatically created on the source computer, they are automatically replicated to the destination computer.

Note: Replicated data is not visible on the receiving system until the replication task completes.

Note: The target dataset on the receiving system is automatically created in read-only mode to protect the data. To mount or browse the data on the receiving system, create a clone of the snapshot and use the clone. Clones are created in read/write mode, making it possible to browse or mount them. See *Snapshots* (page 165) for more information on creating clones.

8.3.1 Examples: Common Configuration

The examples shown here use the same setup of source and destination computers.

8.3.1.1 Alpha (Source)

Alpha is the source computer with the data to be replicated. It is at IP address *10.0.0.102*. A *volume* (page 125) named *alphavol* has already been created, and a *dataset* (page 132) named *alphadata* has been created on that volume. This dataset contains the files which will be snapshotted and replicated onto *Beta*.

This new dataset has been created for this example, but a new dataset is not required. Most users will already have datasets containing the data they wish to replicate.

Create a periodic snapshot of the source dataset by selecting *Storage* \rightarrow *Periodic Snapshot Tasks*. Click the *alphavol/alphadata* dataset to highlight it. Create a *periodic snapshot* (page 150) of it by clicking *Periodic Snapshot Tasks*, then *Add Periodic Snapshot* as shown in Figure 8.23.

This example creates a snapshot of the *alphavol/alphadata* dataset every two hours from Monday through Friday between the hours of 9:00 and 18:00 (6:00 PM). Snapshots are automatically deleted after their chosen lifetime of two weeks expires.

Periodic Snapshots	8
Volume/Dataset:	alphavol/alphadata 💌
Recursive:	
Snapshot Lifetime	2 Week(s) -
Begin:	09:00:00 💌 💰
End:	18:00:00 🔹 🚺
Interval:	1 hour 💌 (i)
Weekday:	 Monday Tuesday Wednesday Thursday Friday Saturday Sunday
Enabled:	
OK Cancel	

Fig. 8.23: Create a Periodic Snapshot for Replication

8.3.1.2 Beta (Destination)

Beta is the destination computer where the replicated data will be copied. It is at IP address *10.0.0.118*. A *volume* (page 125) named *betavol* has already been created.

Snapshots are transferred with *SSH* (page 245). To allow incoming connections, this service is enabled on *Beta*. The service is not required for outgoing connections, and so does not need to be enabled on *Alpha*.

8.3.2 Example: FreeNAS® to FreeNAS® Semi-Automatic Setup

FreeNAS[®] offers a special semi-automatic setup mode that simplifies setting up replication. Create the replication task on *Alpha* by clicking *Replication Tasks* and *Add Replication. alphavol/alphadata* is selected as the dataset to replicate. *betavol* is the destination volume where *alphadata* snapshots are replicated. The *Setup mode* dropdown is set to *Semi-automatic* as shown in Figure 8.24. The IP address of *Beta* is entered in the *Remote hostname* field. A hostname can be entered here if local DNS resolves for that hostname.

Note: If WebGUI HTTP -> HTTPS Redirect has been enabled in System \rightarrow General on the destination computer, Remote HTTP/HTTPS Port must be set to the HTTPS port (usually 443) and Remote HTTPS must be enabled when creating the replication on the source computer.

A	dd Replication	8	
	Volume/Dataset:	alphavol/alphadata 💌 🚺	
	Remote ZFS Volume/Dataset:	betavol	
	Recursively replicate child dataset's snapshots:		
	Delete stale snapshots on remote system:		
	Replication Stream Compression:	Iz4 (fastest)	
	Limit (kB/s):	o (i)	
	Begin:	00:00:00 💌 🚺	
	End:	23:59:00 💌 🚺	
	Enabled:		
	Setup mode:	Semi-automatic	
	Remote hostname:	10.0.0.118	
	Remote HTTP/HTTPS Port:	80	
	Remote HTTPS:		
	Remote Auth Token:	On the remote host go to Storage -> Replicat Tasks, click the Temporary Auth Token button a paste the resulting value in to this field.	tion and
	Dedicated User Enabled:	paste the resulting value in to this field.	
	Dedicated User:	▼	
	Encryption Cipher:	Standard	
	OK Cancel		

Fig. 8.24: Add Replication Dialog, Semi-Automatic

The *Remote Auth Token* field expects a special token from the *Beta* computer. On *Beta*, choose *Storage* \rightarrow *Replication Tasks*, then click *Temporary Auth Token*. A dialog showing the temporary authorization token is shown as in Figure 8.25.

Highlight the temporary authorization token string with the mouse and copy it.

(i)



Fig. 8.25: Temporary Authentication Token on Destination

On the *Alpha* system, paste the copied temporary authorization token string into the *Remote Auth Token* field as shown in Figure 8.26.

Remote Auth Token:

eb8645c5-cle7-4clb-aef2-as

Fig. 8.26: Temporary Authentication Token Pasted to Source

Finally, click the *OK* button to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See *Limiting Replication Times* (page 160) for information about restricting when replication is allowed to run.

Note: The temporary authorization token is only valid for a few minutes. If a *Token is invalid* message is shown, get a new temporary authorization token from the destination system, clear the *Remote Auth Token* field, and paste in the new one.

8.3.3 Example: FreeNAS® to FreeNAS® Dedicated User Replication

A *dedicated user* can be used for replications rather than the root user. This example shows the process using the semiautomatic replication setup between two FreeNAS[®] systems with a dedicated user named *repluser*. SSH key authentication is used to allow the user to log in remotely without a password.

In this example, the periodic snapshot task has not been created yet. If the periodic snapshot shown in the *example configuration* (page 152) has already been created, go to *Storage* \rightarrow *Periodic Snapshot Tasks*, click on the task to select it, and click *Delete* to remove it before continuing.

On Alpha, select Account \rightarrow Users. Click the Add User. Enter repluser for Username, enter /mnt/alphavol/repluser in the Create Home Directory In field, enter Replication Dedicated User for the Full Name, and set the Disable password login option. Leave the other fields at their default values, but note the User ID number. Click OK to create the user.

On *Beta*, the same dedicated user must be created as was created on the sending computer. Select $Account \rightarrow Users$. Click the *Add User*. Enter the *User ID* number from *Alpha*, *repluser* for *Username*, enter */mnt/betavol/repluser* in the *Create Home Directory In* field, enter *Replication Dedicated User* for the *Full Name*, and set the *Disable password login* option. Leave the other fields at their default values. Click *OK* to create the user.

A dataset with the same name as the original must be created on the destination computer, *Beta*. Select *Storage* \rightarrow *Volumes*, click on *betavol*, then click the *Create Dataset* icon at the bottom. Enter *alphadata* as the *Dataset Name*, then click *Add Dataset*.

The replication user must be given permissions to the destination dataset. Still on *Beta*, open a *Shell* (page 289) and enter this command:

zfs allow -ldu repluser create,destroy,diff,mount,readonly,receive,release,send,userprop betavol/ →alphadata

The destination dataset must also be set to read-only. Enter this command in the Shell (page 289):

zfs set readonly=on betavol/alphadata

Close the Shell (page 289) by typing exit and pressing Enter.

The replication user must also be able to mount datasets. Still on *Beta*, go to *System* \rightarrow *Tunables*. Click *Add Tunable*. Enter *vfs.usermount* for the *Variable*, 1 for the *Value*, and choose *Sysctl* from the *Type* drop-down. Click *OK* to save the tunable settings.

Back on Alpha, create a periodic snapshot of the source dataset by selecting Storage \rightarrow Periodic Snapshot Tasks. Click the alphavol/alphadata dataset to highlight it. Create a periodic snapshot (page 150) of it by clicking Periodic Snapshot Tasks, then Add Periodic Snapshot as shown in Figure 8.23.

Still on *Alpha*, create the replication task by clicking *Replication Tasks* and *Add Replication*. *alphavol/alphadata* is selected as the dataset to replicate. *betavol/alphadata* is the destination volume and dataset where *alphadata* snapshots are replicated.

The *Setup mode* dropdown is set to *Semi-automatic* as shown in Figure 8.24. The IP address of *Beta* is entered in the *Remote hostname* field. A hostname can be entered here if local DNS resolves for that hostname.

Note: If WebGUI HTTP -> HTTPS Redirect has been enabled in System \rightarrow General on the destination computer, Remote HTTP/HTTPS Port must be set to the HTTPS port (usually 443) and Remote HTTPS must be enabled when creating the replication on the source computer.

The *Remote Auth Token* field expects a special token from the *Beta* computer. On *Beta*, choose *Storage* \rightarrow *Replication Tasks*, then click *Temporary Auth Token*. A dialog showing the temporary authorization token is shown as in Figure 8.25.

Highlight the temporary authorization token string with the mouse and copy it.

On the *Alpha* system, paste the copied temporary authorization token string into the *Remote Auth Token* field as shown in Figure 8.26.

Set the *Dedicated User* option. Choose *repluser* in the *Dedicated User* drop-down.

Click the OK button to create the replication task.

Note: The temporary authorization token is only valid for a few minutes. If a *Token is invalid* message is shown, get a new temporary authorization token from the destination system, clear the *Remote Auth Token* field, and paste in the new one.

Replication will begin when the periodic snapshot task runs.

Additional replications can use the same dedicated user that has already been set up. The permissions and read only settings made through the *Shell* (page 289) must be set on each new destination dataset.

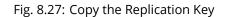
8.3.4 Example: FreeNAS® to FreeNAS® or Other Systems, Manual Setup

This example uses the same basic configuration of source and destination computers shown above, but the destination computer is not required to be a FreeNAS[®] system. Other operating systems can receive the replication if they support SSH, ZFS, and the same features that are in use on the source system. The details of creating volumes and datasets, enabling SSH, and copying encryption keys will vary when the destination computer is not a FreeNAS[®] system.

8.3.4.1 Encryption Keys

A public encryption key must be copied from *Alpha* to *Beta* to allow a secure connection without a password prompt. On *Alpha*, select *Storage* \rightarrow *Replication Tasks* \rightarrow *View Public Key*, producing the window shown in Figure 8.27. Use the mouse to highlight the key data shown in the window, then copy it.





On *Beta*, select *Account* \rightarrow *Users* \rightarrow *View Users*. Click the *root* account to select it, then click *Modify User*. Paste the copied key into the *SSH Public Key* field and click *OK* as shown in Figure 8.28.

expand all collapse all	Account			Permit Sudo:		•
Account	Groups	Users		Microsoft Account:		
🛨 💐 Groups	Add User					A
- 🔒 Users				SSH Public Key:	ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDouCSaPolD	
Add User					XMmkP8PmZ094aKdBzHNsNEkSS8p/VUNuS8+o0BI2	
🛣 View Users	User ID	Username	Primary Group ID		+Braub0kEDL6x0nr9J3c904eElmTTXvJaGtuTu08 HxFHaKL9MI0nbtAu6+shKH0+GJv2vtKZjhIzzNRz	Microsoft Account
📧 🌃 System					8o1xSQkkjv7R6MypGbVe9xcv0YgrKCzpCn2DuBzu	
🛨 🔯 Tasks	0	root	0		YhOzyB3u /fQ90gkJSCBwngyKUoiabJ95WhW34CopaPh8I	false 🔺
🛨 👥 Network	1	daemon	T		/aZQBUm/rx0HoTbEl+e35BpAw7bApHkA+s1	false
🖃 🚰 Storage					/moGzOzZGGKpd7vVgakqjAdwh45E1h6qQtKML0PH	
- 🛃 Volumes	2	operator	5		AuuAQ2/IsQOkFmzeGhlg /1pxTCAdfV3Lmvi+a5sjzcREYtiNOayz Key	false
Volume Manager	3	bin	7		for replication	false
🛃 Import Disk	, i	5				
🛃 Import Volume	4	tty	65533	Home Directory	Owner Group Other	false
View Disks	·	,		Mode:	Read 🖉 🧭	
Yiew Volumes	5	kmem	2	Hode.	Write 🔽	false
📧 💕 Periodic Snapshot Tasks	7	games	13		Execute 🗸 🗸	false
🖭 💼 Replication Tasks				Auxiliary groups:	Available Selected	
📧 🚟 Scrubs	8	news	8	tty	tty	false
Snapshots	9	man	9		uucp >>	false
🛨 📷 VMware-Snapshots 🔤	14	ftp	14		webdav	false
📧 🚺 Directory Service	22	sshd	22		wheel	false
🗄 🚱 Sharing						
🗄 💣 Services	25	smmsp	25	OK Cancel		false
representation of the second s			_			<u>✓</u>
📧 🎹 Jails	Modify Use	r Change E	mail			· .
Reporting						v

Fig. 8.28: Paste the Replication Key

Back on *Alpha*, create the replication task by clicking *Replication Tasks* and *Add Replication*. *alphavol/alphadata* is selected as the dataset to replicate. The destination volume is *betavol*. The *alphadata* dataset and snapshots are replicated there. The IP address of *Beta* is entered in the *Remote hostname* field as shown in Figure 8.29. A hostname can be entered here if local DNS resolves for that hostname.

Click the *SSH Key Scan* button to retrieve the SSH host keys from *Beta* and fill the *Remote hostkey* field. Finally, click *OK* to create the replication task. After each periodic snapshot is created, a replication task will copy it to the destination system. See *Limiting Replication Times* (page 160) for information about restricting when replication is allowed to run.

A	dd Replication	8
	Volume/Dataset:	alphavol/alphadata 💌 🛈
	Remote ZFS Volume/Dataset:	betavol
	Recursively replicate child dataset's snapshots:	
	Delete stale snapshots on remote system:	
	Replication Stream Compression:	Iz4 (fastest)
	Limit (kB/s):	o (i)
	Begin:	00:00:00 👻 (i)
	End:	23:59:00 💌 🚺
	Enabled:	
	Setup mode:	Manual
	Remote hostname:	10.0.0.118
	Remote port:	22
	Dedicated User Enabled:	
	Dedicated User:	.
	Encryption Cipher:	Standard
	Remote hostkey:	10.0.0.118 ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAABAQC4WnS+kfJa CDL1SnPWEqHwuVjEOk8pl+kU8JlS8yyfOALP1/aB c82DdZoNGwtJjn14xTyxA1XJKXio1YYkTnTiLj7M R+S905HLt+vwSUhkfs3EdD8/oOCFmeiw /OOdzjT9oiCrqqnHiL+dySqBjAEOyfoQyTGfzbsy FYG9BZ6aLSzA+oEd7i+aJlE++n6oRCENUCopeFGF m9gADtWwETiHxJkY292JRqhY02k7JrhyzYPSLZvL Yy3mwObSG1Xjf8D2xGgxs7qdiai3r6aKl+TRA4Bi /d8GxVAKwzJPgv /K/aWiibmaUcVBavUbM6OyaRFg9uuhn43HYMHbJa 4fE/r1 10.0.0.118 ecdsa-sha2-nistp256 AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlz dHAyNTYAAABBBANGLOmMyTZl/Fp1aScYX /8s/b3nvXibX /levDCDwJecuD1ASWY5Xx+Wp8YkraJzLv9bonf1w yc2fCL4gzFsOAg= 10.0.0.118 ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIOZtUTtc59hv90WH 7nDoD4li3GdRKaZR/V70gzT8t7GE

OK Cancel SSH Key Scan

8.3.5 Replication Options

Table 8.8 describes the options in the replication task dialog.

	I	
Setting	Value	Description
Volume/Dataset	drop-down menu	On the source computer with snapshots to replicate, choose an existing ZFS pool or dataset with an active periodic snapshot task.
Remote ZFS Vol-	string	Enter the ZFS volume or dataset on the remote or destination computer
ume/Dataset		which will store the snapshots. Example: poolname/datasetname, not the mount point or filesystem path.
Recursively replicate child dataset snapshots	checkbox	When enabled, include snapshots of child datasets from the primary dataset.
Delete stale snapshots	checkbox	Set to delete previous snapshots from the remote or destination system which are no longer present on the source computer.
Replication Stream Com-	drop-down	Choices are <i>lz4</i> (fastest), pigz (all rounder), plzip (best compression), or Off
pression	menu	(no compression). Selecting a compression algorithm can reduce the size of the data being replicated.
Limit (kbps)	integer	Limit replication speed to the specified value in kilobits/second. Default of 0 is unlimited.
Begin	drop-down menu	Define a time to start the replication task.
End	drop-down menu	Define the point in time by which replication must start. A started replica- tion task conitinues until it is finished.
Enabled	checkbox	Deselect to disable the scheduled replication task without deleting it.
Setup mode	drop-down menu	Choose the configuration mode for the remote. Choices are <i>Manual</i> or <i>Semi-automatic</i> . Note semi-automatic only works with remote version 9.10.2 or later.
Remote hostname	string	Enter the IP address or DNS name of remote system to receive the replica- tion data.
Remote port	string	Enter the port number used by the SSH server on the remote or destina- tion computer.
Dedicated User Enabled	checkbox	Select the user account other than root to be used for replication.
Dedicated User	drop-down menu	Only available if <i>Dedicated User Enabled</i> is enabled. Select the user account to be used for replication.
Encryption Cipher	drop-down menu	Standard, Fast, or Disabled.
Remote hostkey	string	Click <i>SSH Key Scan</i> to retrieve the public host key of the remote or destina- tion computer and populate this field with that key.

Table 8.8: F	Replication	Task Options
--------------	-------------	--------------

The replication task runs after a new periodic snapshot is created. The periodic snapshot and any new manual snapshots of the same dataset are replicated onto the destination computer.

When multiple replications have been created, replication tasks run serially, one after another. Completion time depends on the number and size of snapshots and the bandwidth available between the source and destination computers.

The first time a replication runs, it must duplicate data structures from the source to the destination computer. This can take much longer to complete than subsequent replications, which only send differences in data.

Warning: Snapshots record incremental changes in data. If the receiving system does not have at least one snapshot that can be used as a basis for the incremental changes in the snapshots from the sending system, there is no way to identify only the data that has changed. In this situation, the snapshots in the receiving system target dataset are removed so a complete initial copy of the new replicated data can be created.

Selecting Storage \rightarrow Replication Tasks displays Figure 8.30, the list of replication tasks. The Last snapshot sent to remote side column shows the name of the last snapshot that was successfully replicated, and Status shows the current status of each

replication task. The display is updated every five seconds, always showing the latest status.

Storage	ige									
Volumes Perio	Volumes Periodic Snapshot Tasks Replication Tasks Resilver Priority Scrubs Snapshots VMware-Snapshot									
Add Replication ATTENTION: A periodic snapshot of a given ZFS Volume/Dataset is required to create a replication task										
Volume/Dataset Last snapshot sent to remote side Status Remote Hostname Status Remote ZFS Delete stale snapshots on remote system remote syste										
volume1/smb- storage	auto-20170116.0950	beta	Succeeded	betavol	true	lz4	0	00:00:00	23:59:00	true



Note: The encryption key that was copied from the source computer (*Alpha*) to the destination computer (*Beta*) is an RSA public key located in the /data/ssh/replication.pub file on the source computer. The host public key used to identify the destination computer (*Beta*) is from the /etc/ssh/ssh_host_rsa_key.pub file on the destination computer.

8.3.6 Replication Encryption

The default *Encryption Cipher Standard* setting provides good security. *Fast* is less secure than *Standard* but can give reasonable transfer rates for devices with limited cryptographic speed. For networks where the entire path between source and destination computers is trusted, the *Disabled* option can be chosen to send replicated data without encryption.

8.3.7 Limiting Replication Times

The *Begin* and *End* times in a replication task make it possible to restrict when replication is allowed. These times can be set to only allow replication after business hours, or at other times when disk or network activity will not slow down other operations like snapshots or *Scrubs* (page 162). The default settings allow replication to occur at any time.

These times control when replication task are allowed to start, but will not stop a replication task that is already running. Once a replication task has begun, it will run until finished.

8.3.8 Troubleshooting Replication

Replication depends on SSH, disks, network, compression, and encryption to work. A failure or misconfiguration of any of these can prevent successful replication.

8.3.8.1 SSH

SSH (page 245) must be able to connect from the source system to the destination system with an encryption key. This can be tested from *Shell* (page 289) by making an *SSH* (page 245) connection from the source system to the destination system. From the previous example, this is a connection from *Alpha* to *Beta* at *10.0.0.118*. Start the *Shell* (page 289) on the source machine (*Alpha*), then enter this command:

```
ssh -vv -i /data/ssh/replication 10.0.0.118
```

On the first connection, the system might say

No matching host key fingerprint found in DNS. Are you sure you want to continue connecting (yes/no)?

Verify that this is the correct destination computer from the preceding information on the screen and type yes. At this point, an *SSH* (page 245) shell connection is open to the destination system, *Beta*.

If a password is requested, SSH authentication is not working. See Figure 8.27 above. This key value must be present in the /root/.ssh/authorized_keys file on *Beta*, the destination computer. The /var/log/auth.log file can show diagnostic errors for login problems on the destination computer also.

8.3.8.2 Compression

Matching compression and decompression programs must be available on both the source and destination computers. This is not a problem when both computers are running FreeNAS[®], but other operating systems might not have *lz4*, *pigz*, or *plzip* compression programs installed by default. An easy way to diagnose the problem is to set *Replication Stream Compression* to *Off*. If the replication runs, select the preferred compression method and check /var/log/debug.log on the FreeNAS[®] system for errors.

8.3.8.3 Manual Testing

On Alpha, the source computer, the /var/log/messages file can also show helpful messages to locate the problem.

On the source computer, *Alpha*, open a *Shell* (page 289) and manually send a single snapshot to the destination computer, *Beta*. The snapshot used in this example is named auto-20161206.1110-2w. As before, it is located in the *alphavol/alphadata* dataset. A @ symbol separates the name of the dataset from the name of the snapshot in the command.

zfs send alphavol/alphadata@auto-20161206.1110-2w | ssh -i /data/ssh/replication 10.0.0.118 zfs recv_ →betavol

If a snapshot of that name already exists on the destination computer, the system will refuse to overwrite it with the new snapshot. The existing snapshot on the destination computer can be deleted by opening a *Shell* (page 289) on *Beta* and running this command:

zfs destroy -R betavol/alphadata@auto-20161206.1110-2w

Then send the snapshot manually again. Snapshots on the destination system, *Beta*, can be listed from the *Shell* (page 289) with zfs list -t snapshot or by going to *Storage* \rightarrow *Snapshots*.

Error messages here can indicate any remaining problems.

8.4 Resilver Priority

Resilvering, or the process of copying data to a replacement disk, is best completed as quickly as possible. Increasing the priority of resilvers can help them to complete more quickly. The *Resilver Priority* tab makes it possible to increase the priority of resilvering at times where the additional I/O or CPU usage will not affect normal usage. Select *Storage* \rightarrow *Resilver Priority* to display the screen shown in Figure 8.31. Table 8.9 describes the fields on this screen.

Storage						
Volumes	Periodic Snapshot Tasks	Replication Tasks	Resilver Priority	Scrubs	Snapshots	VMware-Snapshot
Enabled	:					
Begin hi	gher priority resilvering	g at this time:	6:00 PM 👻			
End hig	ner priority resilvering a	at this time:	9:00 AM 👻			
Weekda	y:		 Monday Tuesday Wednesday Hursday Thursday Friday Saturday Sunday 	,		
Save						

Fig. 8.31: Resilver Priority

Table 8.9: Resilver	Priority	Options
---------------------	----------	---------

Setting	Value	Description
Enabled	checkbox	Set to enable higher-priority resilvering.
Begin higher priority resilvering at this time	drop-down	Start time to begin higher-priority resilvering.
End higher priority resilvering at this time	drop-down	End time to begin higher-priority resilvering.
Weekday	checkboxes	Use higher-priority resilvering on these days of the week.

8.5 Scrubs

A scrub is the process of ZFS scanning through the data on a volume. Scrubs help to identify data integrity problems, detect silent data corruptions caused by transient hardware issues, and provide early alerts of impending disk failures. FreeNAS[®] makes it easy to schedule periodic automatic scrubs.

Each volume should be scrubbed at least once a month. Bit errors in critical data can be detected by ZFS, but only when that data is read. Scheduled scrubs can find bit errors in rarely-read data. The amount of time needed for a scrub is proportional to the quantity of data on the volume. Typical scrubs take several hours or longer.

The scrub process is I/O intensive and can negatively impact performance. Schedule scrubs for evenings or weekends to minimize impact to users. Make certain that scrubs and other disk-intensive activity like *S.M.A.R.T. Tests* (page 110) are scheduled to run on different days to avoid disk contention and extreme performance impacts.

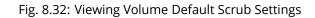
Scrubs only check used disk space. To check unused disk space, schedule *S.M.A.R.T. Tests* (page 110) of *Type Long Self-Test* to run once or twice a month.

Scrubs are scheduled and managed with Storage \rightarrow Scrubs.

When a volume is created, a ZFS scrub is automatically scheduled. An entry with the same volume name is added to Storage

 \rightarrow Scrubs. A summary of this entry can be viewed with Storage \rightarrow Scrubs \rightarrow View Scrubs. Figure 8.32 displays the default settings for the volume named volume1. In this example, the entry has been highlighted and the *Edit* button clicked to display the *Edit* screen. Table 8.10 summarizes the options in this screen.

Edit	8
Volume:	volumel
Threshold days:	35 <i>(i)</i>
Description:	
Minute:	Every N minute Each selected minute 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59
Hour:	<i>Every N hour</i> Each selected hour 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Day of month:	Every N day of month Each selected day of month I I I I I I I I I I I I I I I I I I
Month:	• 🕎 January



Setting	Value	Description
Volume	drop-down	Choose a volume to be scrubbed.
	menu	
Threshold days	integer	Define the number of days to prevent a scrub from running after the last has
		completed. This ignores any other calendar schedule. The default is a multi-
		ple of 7 to ensure that the scrub always occurs on the same day of the week.
Description	string	Optional text description of scrub.
Minute	slider or minute	If the slider is used, a scrub occurs every N minutes. If specific minutes are
	selections	chosen, a scrub runs only at the selected minute values.
Hour	slider or hour	If the slider is used, a scrub occurs every N hours. If specific hours are chosen,
	selections	a scrub runs only at the selected hour values.
Day of Month	slider or month	If the slider is used, a scrub occurs every N days. If specific days of the month
	selections	are chosen, a scrub runs only on the selected days of the selected months.
Month	checkboxes	Define the day of the month to run the scrub.
Day of week	checkboxes	A scrub occurs on the selected days. The default is <i>Sunday</i> to least impact
		users. Note that this field and the <i>Day of Month</i> field are ORed together: set-
		ting <i>Day of Month</i> to <i>01,15</i> and <i>Day of week</i> to <i>Thursday</i> will cause scrubs to run
		on the 1st and 15th days of the month, but also on any Thursday.
Enabled	checkbox	Unset to disable the scheduled scrub without deleting it.

Tahle	8 10.	7FS	Scruh	Options
Table	0.10.	213	Sciub	Options

Review the default selections and, if necessary, modify them to meet the needs of the environment. Note that the *Threshold* field is used to prevent scrubs from running too often, and overrides the schedule chosen in the other fields. Also, if a pool is locked or unmounted when a scrub is scheduled to occur, it will not be scrubbed.

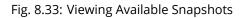
Scheduled scrubs can be deleted with the *Delete* button, but this is not recommended. **Scrubs can provide an early indication of disk issues before a disk failure.** If a scrub is too intensive for the hardware, consider temporarily deselecting the *Enabled* button for the scrub until the hardware can be upgraded.

8.6 Snapshots

Snapshots are scheduled using Storage \rightarrow Periodic Snapshot Tasks. To view and manage the listing of created snapshots, use Storage \rightarrow Snapshots. An example listing is shown in Figure 8.33.

Note: If snapshots do not appear, check that the current time configured in *Periodic Snapshot Tasks* (page 150) does not conflict with the *Begin, End*, and *Interval* settings. If the snapshot was attempted but failed, an entry is added to /var/log/ messages. This log file can be viewed in *Shell* (page 289).

Stora	ge				
Volum	es Periodic Snapshot Tasks Replication Tasks	Resilver Priority Scrubs Snapshots VMware	e-Snapshot		
	Volume/Dataset	Snapshot Name	Used	Refer	Available Actions
≫	No filter applied				
	volumel	auto-20171018.0840-2w	0	88.0 KiB	
	volumel	auto-20171018.0850-2w	0	88.0 KiB	
	volumel	auto-20171018.0900-2w	0	88.0 KiB	
	volumel	auto-20171018.0910-2w	0	88.0 KiB	



The listing includes the name of the volume or dataset, the name of each snapshot, and the amount of used and referenced data.

Used is the amount of space consumed by this dataset and all of its descendants. This value is checked against the dataset quota and reservation. The space used does not include the dataset reservation, but does take into account the reservations of any descendent datasets. The amount of space that a dataset consumes from its parent, as well as the amount of space freed if this dataset is recursively deleted, is the greater of its space used and its reservation. When a snapshot is created, the space is initially shared between the snapshot and the filesystem, and possibly with previous snapshots. As the filesystem changes, space that was previously shared becomes unique to the snapshot, and is counted in the used space of the snapshot. Additionally, deleting snapshots can increase the amount of space unique to (and used by) other snapshots. The amount of space used, available, or referenced does not take into account pending changes. While pending changes are generally accounted for within a few seconds, disk changes do not necessarily guarantee that the space usage information is updated immediately.

Tip: Space used by individual snapshots can be seen by running zfs list -t snapshot from Shell (page 289).

Refer indicates the amount of data accessible by this dataset, which may or may not be shared with other datasets in the pool. When a snapshot or clone is created, it initially references the same amount of space as the filesystem or snapshot it was created from, since its contents are identical.

Snapshots have icons on the right side for several actions.

Clone Snapshot prompts for the name of the clone to create. A clone is a writable copy of the snapshot. Since a clone is actually a dataset which can be mounted, it appears in the *Volumes* tab rather than the *Snapshots* tab. By default, -clone is added to the name of a snapshot when a clone is created.

Destroy Snapshot a pop-up message asks for confirmation. Child clones must be deleted before their parent snapshot can be deleted. While creating a snapshot is instantaneous, deleting a snapshot can be I/O intensive and can take a long time, especially when deduplication is enabled. To delete a block in a snapshot, ZFS has to walk all the allocated blocks to see if that block is used anywhere else. If it is not used, it can be freed.

The most recent snapshot also has a **Rollback Snapshot** icon. Clicking the icon asks for confirmation before rolling back to this snapshot state. Confirming by clicking *Yes* causes any files that have changed since the snapshot was taken to be reverted back to their state at the time of the snapshot.

Note: Rollback is a potentially dangerous operation and causes any configured replication tasks to fail as the replication system uses the existing snapshot when doing an incremental backup. To restore the data within a snapshot, the recommended steps are:

- 1. Clone the desired snapshot.
- 2. Share the clone with the share type or service running on the FreeNAS[®] system.
- 3. After users have recovered the needed data, destroy the clone in the Active Volumes tab.

This approach does not destroy any on-disk data and has no impact on replication.

A range of snapshots can be selected with the mouse. Click on the option in the left column of the first snapshot, then press and hold *Shift* and click on the option for the end snapshot. This can be used to select a range of obsolete snapshots to be deleted with the *Destroy* icon at the bottom. Be cautious and careful when deleting ranges of snapshots.

Periodic snapshots can be configured to appear as shadow copies in newer versions of Windows Explorer, as described in *Configuring Shadow Copies* (page 203). Users can access the files in the shadow copy using Explorer without requiring any interaction with the FreeNAS[®] graphical administrative interface.

The ZFS Snapshots screen allows the creation of filters to view snapshots by selected criteria. To create a filter, click the *Define filter* icon (near the text *No filter applied*). When creating a filter:

- Select the column or leave the default of Any Column.
- Select the condition. Possible conditions are: contains (default), is, starts with, ends with, does not contain, is not, does not start with, does not end with, and is empty.
- Enter a value that meets the view criteria.

• Click the *Filter* button to save the filter and exit the define filter screen. Alternately, click the + button to add another filter.

When creating multiple filters, select the filter to use before leaving the define filter screen. After a filter is selected, the *No filter applied* text changes to *Clear filter*. Clicking *Clear filter* produces a pop-up message indicates that this removes the filter and all available snapshots are listed.

Warning: A snapshot and any files it contains will not be accessible or searchable if the mount path of the snapshot is longer than 88 ascii characters. The data within the snapshot will be safe, and the snapshot will become accessible again when the mount path is shortened. For details of this limitation, and how to shorten a long mount path, see *Path and Name Lengths* (page 12).

8.6.1 Browsing a snapshot collection

All snapshots for a dataset are accessible as an ordinary hierarchical filesystem, which can be reached from a hidden .zfs file located at the root of every dataset. A user with permission to access that file can view and explore all snapshots for a dataset like any other files - from the CLI or via *File Sharing* services such as *Samba*, *NFS* and *FTP*. This is an advanced capability which requires some command line actions to achieve. In summary, the main changes to settings that are required are:

- Snapshot visibility must be manually enabled in the ZFS properties of the dataset.
- In Samba auxiliary settings, the veto files command must be modified to not hide the .zfs file, and the setting zfsacl:expose_snapdir=true must be added.

The effect will be that any user who can access the dataset contents, will also be able to view the list of snapshots by navigating to the .zfs directory of the dataset, and to browse and search any files they have permission to access throughout the entire snapshot collection of the dataset. A user's ability to view files within a snapshot will be limited by any permissions or ACLs set on the files when the snapshot was taken. Snapshots are fixed as "read-only", so this access does not permit the user to change any files in the snapshots, or to modify or delete any snapshot, even if they had write permission at the time when the snapshot was taken.

Note: ZFS has a zfs diff command which can list the files that have changed between any two snapshot versions within a dataset, or between any snapshot and the current data.

8.7 VMware-Snapshot

Storage \rightarrow VMware-Snapshot is used to coordinate ZFS snapshots when using FreeNAS[®] as a VMware datastore. Once this type of snapshot is created, FreeNAS[®] will automatically snapshot any running VMware virtual machines before taking a scheduled or manual ZFS snapshot of the dataset or zvol backing that VMware datastore. The temporary VMware snapshots are then deleted on the VMware side but still exist in the ZFS snapshot and can be used as stable resurrection points in that snapshot. These coordinated snapshots will be listed in *Snapshots* (page 165).

Figure 8.34 shows the menu for adding a VMware snapshot and Table 8.11 summarizes the available options.

Add VMware-Snap	oshot	36
Hostname:		
Username:		i
Password:		
ZFS Filesystem:	volume1	
Datastore:	v	i
OK Cancel	Fetch Datastores	

Fig. 8.34: Adding a VMware Snapshot

Table 8.11:	VMware	Snapshot	Options
-------------	--------	----------	---------

Setting	Value	Description
Hostname	string	Enter the IP address or hostname of VMware host. When clustering, this is
		the vCenter server for the cluster.
Username	string	Enter the username on the VMware host with permission to snapshot vir-
		tual machines.
Password	string	Enter the password associated with Username.
ZFS Filesystem	drop-down menu	Select the filesystem to snapshot.
Datastore	drop-down menu	Enter the Hostname, Username, and Password. Click Fetch Datastores to
		populate the menu and select the datastore with which to synchronize.

DIRECTORY SERVICES

FreeNAS[®] supports integration with these directory services:

- Active Directory (page 169) (for Windows 2000 and higher networks)
- *LDAP* (page 174)
- *NIS* (page 176)

It also supports *Kerberos Realms* (page 177), *Kerberos Keytabs* (page 178), and the ability to add more parameters to *Kerberos Settings* (page 179).

This section summarizes each of these services and their available configurations within the FreeNAS[®] GUI.

9.1 Active Directory

Active Directory (AD) is a service for sharing resources in a Windows network. AD can be configured on a Windows server that is running Windows Server 2000 or higher or on a Unix-like operating system that is running Samba version 4 (https://wiki.samba.org/index.php/Setting_up_Samba_as_an_Active_Directory_Domain_Controller#Provisioning_a_Samba_Active_Di Since AD provides authentication and authorization services for the users in a network, it is not necessary to recreate these user accounts on the FreeNAS[®] system. Instead, configure the Active Directory service so that it can import the account information and imported users can be authorized to access the SMB shares on the FreeNAS[®] system.

Many changes and improvements have been made to Active Directory support within FreeNAS[®]. It is strongly recommended to update the system to the latest FreeNAS[®] 11.2 before attempting Active Directory integration.

Ensure name resolution is properly configured before configuring the Active Directory service. ping the domain name of the Active Directory domain controller from *Shell* (page 289) on the FreeNAS[®] system. If the ping fails, check the DNS server and default gateway settings in *Network* \rightarrow *Global Configuration* on the FreeNAS[®] system.

Add a DNS record for the FreeNAS[®] system on the Windows server and verify the hostname of the FreeNAS[®] system can be pinged from the domain controller.

Active Directory relies on Kerberos, which is a time-sensitive protocol. The time on both the FreeNAS[®] system and the Active Directory Domain Controller cannot be out of sync by more than a few minutes. The best way to ensure the same time is running on both systems is to configure both systems to:

- use the same NTP server (set in *System* \rightarrow *NTP Servers* on the FreeNAS[®] system)
- have the same timezone
- be set to either localtime or universal time at the BIOS level

Figure 9.1 shows the screen that appears when *Directory Service* \rightarrow *Active Directory* is chosen. Table 9.1 describes the configurable options. Some settings are only available in Advanced Mode. To see these settings, either click *Advanced Mode* or configure the system to always display these settings by checking *Show advanced fields by default* in *System* \rightarrow *Advanced*.

۵	Directory Service							
1	Active Directory	LDAP	NIS	Kerberos F	lealms	Kerberos Keytabs	Kerbe	ros Settings
	Domain Name	(DNS/Real	m-Name)	:				(i)
	Domain Accou	nt Name:						(i)
	Domain Accou	nt Passwo	rd:					
	AD check conr	nectivity fro	equency	(seconds):			60	Ì
	How many rec	overy atte	mpts:				10	ì
	Enable Monito	oring:			<u> </u>)		
	Enable:							
	Save	ced Mode	Rebuild Dir	rectory Service	Cache)		

Fig. 9.1: Configuring Active Directory

Table 0.1. Active	Diroctory	Configuration	Ontions
Table 9.1: Active	Directory	Comguration	options

Setting	Value	Advanced Mode	Description
Domain Name (DNS/Realm-Name)	string		Name of Active Directory domain (<i>example.com</i>) or child do- main (<i>sales.example.com</i>). This setting is mandatory and the GUI will refuse to save the settings if the domain controller for the specified domain cannot be found.
Domain Account Name	string		Name of the Active Directory administrator account. This set- ting is mandatory and the GUI will refuse to save the settings if it cannot connect to the domain controller using this account name.
Domain Account Password	string		Password for the Active Directory administrator account. This setting is mandatory and the GUI will refuse to save the set- tings if it cannot connect to the domain controller using this password.
AD check connectivity frequency (seconds)	integer		How often to verify that Active Directory services are active.
How many recovery attempts	integer		Number of times to attempt reconnecting to the Active Direc- tory server. Tries forever when set to <i>0</i> .
Enable Monitoring	checkbox		Restart Active Directory automatically if the service is discon- nected.
Encryption Mode	drop-down menu	\checkmark	Choices are <i>Off</i> , <i>SSL</i> , or <i>TLS</i> .
Certificate	drop-down menu	✓	Select the certificate of the Active Directory server if SSL con- nections are used. If a certificate does not exist yet, create a <i>CA</i> (page 87), then create a certificate on the Active Directory server and import it to the FreeNAS [®] system with <i>Certificates</i> (page 89).
Verbose logging	checkbox	\checkmark	Set to log attempts to join the domain to /var/log/ messages.

Continued on next page

Setting	Value	Advanced	ntinued from previous page Description
-		Mode	
UNIX extensions	checkbox	√	Only set if the AD server is explicitly configured to map permissions for UNIX users. Enabling provides persistent UIDs and GUIDs, otherwise, users/groups are mapped to the UID/GUID
Allow Trusted Do-	checkbox		range configured in Samba. Only enable if the network has active domain/forest
mains		, v	trusts (https://docs.microsoft.com/en-us/previous- versions/windows/it-pro/windows-server- 2003/cc757352(v=ws.10)) and files need to be managed on multiple domains. Use with caution as it will generate more
			winbindd traffic, slowing down the ability to filter through user and group information.
Use Default Domain	checkbox	\checkmark	Unset to prepend the domain name to the username. If <i>Al-low Trusted Domains</i> is set and multiple domains use the same usernames, unset to prevent name collisions.
Allow DNS updates	checkbox	\checkmark	Unset to disable Samba from doing DNS updates when joining a domain.
Disable Active Di- rectory user/group cache	checkbox	\checkmark	Set to disable caching of AD users and groups. This is useful if the system cannot bind to a domain with a large number of users or groups.
Site Name	string	√	The relative distinguished name of the site object in Active Directory.
Domain Controller	string	√	Automatically be added to the SRV record for the domain and, when multiple controllers are specified, FreeNAS [®] selects the closest DC which responds. Uses the short form of the FQDN. An example is <i>sampleserver</i> .
Global Catalog Server	string	√	If the hostname of the global catalog server to use is specified, make sure it is resolvable.
Kerberos Realm	drop-down menu	\checkmark	Select the realm created using the instructions in <i>Kerberos Realms</i> (page 177).
Kerberos Principal	drop-down menu	~	Browse to the location of the keytab created using the instruc- tions in <i>Kerberos Keytabs</i> (page 178).
AD timeout	integer	~	In seconds, increase if the AD service does not start after con- necting to the domain.
DNS timeout	integer	 ✓ 	In seconds, increase if AD DNS queries timeout.
ldmap backend	drop-down menu and Edit	✓	Select the backend to use to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs. See Table 9.2 for a summary of the available backends. Click the <i>Edit</i> link to configure the backend.
Windbind NSS Info	drop-down menu	✓	Defines the schema to use when querying AD for user/group info. <i>rfc2307</i> uses the RFC2307 schema included in Windows 2003 R2, <i>sfu20</i> is for Services For Unix 3.0 or 3.5, and <i>sfu</i> is for Services For Unix 2.0.
SASL wrapping	drop-down menu	√ 	Defines how LDAP traffic is transmitted. Choices are <i>plain</i> (plain text), <i>sign</i> (signed only), or <i>seal</i> (signed and encrypted). Windows 2000 SP3 and newer can be configured to enforce signed LDAP connections.
Enable	checkbox		Enable the Active Directory service.
NetBIOS name	string	\checkmark	Limited to 15 characters. Automatically populated with the original hostname of the system. This must be different from the <i>Workgroup</i> name.
NetBIOS alias	string	\checkmark	Limited to 15 characters.

Table 9.1 – continued from previous page

Table 9.2 summarizes the backends which are available in the *Idmap backend* drop-down menu. Each backend has its own

man page (http://samba.org.ru/samba/docs/man/manpages/) which gives implementation details. Since selecting the wrong backend will break Active Directory integration, a pop-up menu will appear whenever changes are made to this setting.

Value	Description
ad	AD server uses RFC2307 or Services For Unix schema extensions. Mappings
	must be provided in advance by adding the uidNumber attributes for users
	and gidNumber attributes for groups in the AD.
autorid	Similar to <i>rid</i> , but automatically configures the range to be used for each do-
	main, so there is no need to specify a specific range for each domain in the
	forest. The only needed configuration is the range of UID/GIDs to use for
	user/group mappings and an optional size for the ranges.
fruit	Generate IDs the way Apple Mac OS X does, so UID and GID can be identical
	on all FreeNAS [®] servers on the network. For use in <i>LDAP</i> (page 174) environ-
	ments where Apple's Open Directory is the authoritative LDAP server.
ldap	Stores and retrieves mapping tables in an LDAP directory service. Default for
	LDAP directory service.
nss	Provides a simple means of ensuring that the SID for a Unix user is reported
	as the one assigned to the corresponding domain user.
rfc2307	An AD server is required to provide the mapping between the name and SID
	and an LDAP server is required to provide the mapping between the name
	and the UID/GID.
rid	Default for AD. Requires an explicit idmap configuration for each domain, us-
	ing disjoint ranges where a writeable default idmap range is to be defined,
	using a backend like tdb or ldap.
script	Stores mapping tables for clustered environments in the winbind_cache tdb.
tdb	Default backend used by winbindd for storing mapping tables.
tdb2	Substitute for tdb used by winbindd in clustered environments.

Table 9.2	ID Mapping	Backends
	in mapping	Dackenus

Click *Rebuild Directory Service Cache* if a new Active Directory user needs immediate access to FreeNAS[®]. This occurs automatically once a day as a cron job.

Note: Active Directory places restrictions on which characters are allowed in Domain and NetBIOS names, a limits the length of those names to 15 characters. If there are problems connecting to the realm, verify (https://support.microsoft.com/en-us/help/909264/naming-conventions-in-active-directory-for-computers-domains-sites-and) the settings do not include any disallowed characters. The Administrator account password cannot contain the \$ character. If a \$ exists in the domain administrator's password, kinit will report a "Password Incorrect" error and ldap_bind will report an "Invalid credentials (49)" error.

It can take a few minutes after configuring the Active Directory service for the AD information to be populated to the FreeNAS[®] system. Once populated, the AD users and groups will be available in the drop-down menus of the *Permissions* screen of a volume/dataset. For performance reasons, every available user may not show in the listing. However, it will autocomplete all applicable users when typing in a username.

The Active Directory users and groups that are imported to the FreeNAS[®] system are shown by typing commands in the FreeNAS[®] *Shell* (page 289):

- View users: wbinfo -u
- View groups: wbinfo -g

In addition, wbinfo -t tests the connection and, if successful, shows a message similar to:

checking the trust secret for domain YOURDOMAIN via RPC calls succeeded

To manually check that a specified user can authenticate, use net ads join -S dcname -U username.

getent passwd and getent group can provide more troubleshooting information if no users or groups are listed in the output.

Tip: Sometimes network users do not appear in the drop-down menu of a *Permissions* screen but the wbinfo commands display these users. This is typically due to the FreeNAS[®] system taking longer than the default ten seconds to join Active Directory. Increase the value of *AD timeout* to 60 seconds.

To change a certificate, set the *Encryption Mode* to *Off*, then disable AD by unchecking *Enable*. Click *Save*. Select the new *Certificate*, set the *Encryption Mode* as desired, check *Enable* to re-enable AD, and click *Save* to restart AD.

9.1.1 Troubleshooting Tips

When running AD in a 2003/2008 mixed domain, see this posting (https://forums.freenas.org/index.php?threads/2008r2-2003-mixed-domain.1931/) for instructions to prevent the secure channel key from becoming corrupt.

Active Directory uses DNS to determine the location of the domain controllers and global catalog servers in the network. Use host -t srv _ldap._tcp.domainname.com to determine the SRV records of the network and change the weight and/or priority of the SRV record to reflect the fastest server. More information about SRV records can be found in the Technet article How DNS Support for Active Directory Works (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc759550(v=ws.10)).

The realm used depends upon the priority in the SRV DNS record. DNS can override the system Active Directory settings. When unable to connect to the correct realm, check the SRV records on the DNS server.

If the cache becomes out of sync due to an AD server being taken off and back online, resync the cache using *Directory Service* \rightarrow *Active Directory* \rightarrow *Rebuild Directory Service Cache*.

An expired password for the administrator account will cause kinit to fail. Ensure the password is still valid. Also, doublecheck the password on the AD account being used does not include any spaces, special symbols, and is not unusually long.

If the Windows server version is lower than 2008 R2, try creating a *Computer* entry on the Windows server's OU. When creating this entry, enter the FreeNAS[®] hostname in the *name* field. Make sure it is under 15 characters, the same name as the one set in the *Hostname* field in *Network* \rightarrow *Global Configuration*, and the same *NetBIOS Name* in *Directory Service* \rightarrow *Active Directory* settings. Make sure the hostname of the domain controller is set in the *Domain Controller* field of *Directory Service* \rightarrow *Active Directory*.

9.1.2 If the System Does not Join the Domain

If the system will not join the Active Directory domain, run these commands in the order listed. echo commands will return a value of 0 and klist will show a Kerberos ticket:

```
sqlite3 /data/freenas-v1.db "update directoryservice_activedirectory set ad_enable=1;"
echo $?
service ix-kerberos start
service ix-nsswitch start
service ix-kinit start
service ix-kinit status
echo $?
klist
```

If the cache becomes out of sync due to an AD server being taken off and back online, resync the cache using *Directory Service* \rightarrow *Active Directory* \rightarrow *Rebuild Directory Service Cache*.

Note: If any of the commands fail or result in a traceback, create a bug report at https://redmine.ixsystems.com/projects/ freenas/issues that includes the commands in the order in which they were run and the exact wording of the error message or traceback.

Next, only run these two commands **if** *Unix extensions* is set in *Advanced Mode* and a keytab has been uploaded using *Kerberos Keytabs* (page 178):

service ix-sssd start service sssd start

Finally, run these commands. echo returns a 0 unless something has gone wrong:

```
python /usr/local/www/freenasUI/middleware/notifier.py start cifs
service ix-activedirectory start
service ix-activedirectory status
echo $?
python /usr/local/www/freenasUI/middleware/notifier.py restart cifs
service ix-pam start
service ix-cache start &
```

9.2 LDAP

FreeNAS[®] includes an OpenLDAP (http://www.openldap.org/) client for accessing information from an LDAP server. An LDAP server provides directory services for finding network resources such as users and their associated permissions. Examples of LDAP servers include Microsoft Server (2000 and newer), Mac OS X Server, Novell eDirectory, and OpenLDAP running on a BSD or Linux system. If an LDAP server is running on the network, configure the FreeNAS[®] LDAP service so network users can authenticate to the LDAP server and have authorized access to the data stored on the FreeNAS[®] system.

Note: LDAP authentication for SMB shares is disabled unless the LDAP directory has been configured for and populated with Samba attributes. The most popular script for performing this task is smbldap-tools (https://wiki.samba.org/index.php/4.1_smbldap-tools). In addition, the LDAP server must support SSL/TLS and the certificate for the LDAP server CA must be imported with *System* \rightarrow *CAs* \rightarrow *Import CA*. Note that non-CA certificates are not supported at this time.

Tip: Apple's Open Directory (https://manuals.info.apple.com/MANUALS/0/MA954/en_US/Open_Directory_Admin_v10.5_3rd_Ed.pdf) is an LDAP-compatible directory service into which FreeNAS[®] can be integrated. See FreeNAS with Open Directory in Mac OS X environments (https://forums.freenas.org/index.php?threads/howto-freenas-with-open-directory-in-mac-os-x-environments.46493/).

Figure 9.2 shows the LDAP Configuration screen that is seen after clicking *Directory Service* \rightarrow *LDAP*.

D	irectory Service					
A	ctive Directory	LDAP	NIS	Kerberos Realms	Kerberos Keytabs	Kerberos Settings
	Hostname:			i		
	Base DN:			i		
	Bind DN:			i		
	Bind password	1:		i		
	Enable:					
	Save	anced Mode	Rebuild	Directory Service	Cache	

Fig. 9.2: Configuring LDAP

Table 9.3 summarizes the available configuration options. Some settings are only available in Advanced Mode. To see these settings, either click the *Advanced Mode* button or configure the system to always display these settings by checking the box *Show advanced fields by default* in *System* \rightarrow *Advanced*.

Those new to LDAP terminology should read the OpenLDAP Software 2.4 Administrator's Guide (http://www.openIdap.org/doc/admin24/).

Setting	Value	Advanced Mode	Description
Hostname	string		Hostname or IP address of the LDAP server.
Base DN	string		Top level of the LDAP directory tree to be used when searching for resources. Example: <i>dc=test,dc=org</i> .
Bind DN	string		Name of administrative account on the LDAP server. Example: <i>cn=Manager,dc=test,dc=org</i> .
Bind password	string		Password for Root bind DN.
Allow Anonymous	checkbox	✓	Instructs the LDAP server to not provide authentication and to
Binding			allow read and write access to any client.
User Suffix	string	~	Optional. Can be added to the name when the user account is added to the LDAP directory. Example: dept. or company name.
Group Suffix	string	✓	Optional. Can be added to the name when the group is added to the LDAP directory. Example: dept. or company name.
Password Suffix	string	✓	Optional. Can be added to the password when the password is added to LDAP directory.
Machine Suffix	string	\checkmark	Optional. Can be added to the name when the system added to the LDAP directory. Example: server, accounting.
SUDO Suffix	string	\checkmark	Use if LDAP-based users need superuser access.

Table 9.3: LDAP Configuration Options

Continued on next page

Setting	Value	Advanced Mode	Description
Kerberos Realm	drop-down menu	✓	Select the realm created using the instructions in <i>Kerberos Realms</i> (page 177).
Kerberos Principal	drop-down menu	\checkmark	Browse to the location of the principal in the keytab created as described in <i>Kerberos Keytabs</i> (page 178).
Encryption Mode	drop-down menu	✓	Choices are <i>Off</i> , <i>SSL</i> , or <i>TLS</i> . Note that either <i>SSL</i> or <i>TLS</i> and a <i>Certificate</i> must be selected for authentication to work.
Certificate	drop-down menu	✓	Select the certificate of the LDAP CA (required if authentication is used). The certificate for the LDAP server CA must first be imported with <i>System</i> \rightarrow <i>Certificates</i> \rightarrow <i>Import Certificate</i> .
LDAP timeout	integer	✓	Increase this value (in seconds) if obtaining a Kerberos ticket times out.
DNS timeout	integer	\checkmark	Increase this value (in seconds) if DNS queries timeout.
ldmap backend	drop-down menu and Edit	✓	Select the backend to use to map Windows security identifiers (SIDs) to UNIX UIDs and GIDs. See Table 9.2 for a summary of the available backends. Click the <i>Edit</i> link to configure the selected backend.
Samba Schema	checkbox	\checkmark	Set if LDAP authentication for SMB shares is needed and the LDAP server is already configured with Samba attributes.
Auxiliary Parameters	string	✓	Additional options for sssd.conf(5) (https://jhrozek.fedorapeople.org/sssd/1.11.6/man/sssd.conf.5.html
Schema	drop-down menu	✓	If <i>Samba Schema</i> is set, select the schema to use. Choices are <i>rfc2307</i> and <i>rfc2307bis</i> .
Enable	checkbox		Unset to disable the configuration without deleting it.
NetBIOS Name	string	✓	Limited to 15 characters. Automatically populated with the original hostname of the system. This must be different from the <i>Workgroup</i> name
NetBIOS Alias	string	\checkmark	Limited to 15 characters.

Table 9.3 – continued from previous page

Click the *Rebuild Directory Service Cache* button after adding a user to LDAP who needs immediate access to FreeNAS[®]. Otherwise this occurs automatically once a day as a cron job.

Note: FreeNAS[®] automatically appends the root DN. This means the scope and root DN are not to be included when configuring the user, group, password, and machine suffixes.

LDAP users and groups appear in the drop-down menus of the guilabel: *Permissions* screen of a dataset after configuring the LDAP service. Type getent passwd from *Shell* (page 289) to verify the users have been imported. Type getent group to verify the groups have been imported.

If the users and groups are not listed, refer to Common errors encountered when using OpenLDAP Software (http://www.openIdap.org/doc/admin24/appendix-common-errors.html) for common errors and how to fix them. When troubleshooting LDAP, open *Shell* (page 289) and look for error messages in /var/log/auth.log.

To clear LDAP users and groups from FreeNAS[®], go to *Directory Services* \rightarrow *LDAP*, clear the *Hostname* field, unset *Enable*, and click *Save*. Confirm LDAP users and groups are cleared by going to the *Shell* and viewing the output of the getent passwd and getent group commands.

9.3 NIS

The Network Information Service (NIS) maintains and distributes a central directory of Unix user and group information, hostnames, email aliases, and other text-based tables of information. If an NIS server is running on the network, the FreeNAS[®] system can be configured to import the users and groups from the NIS directory. **Note:** In Windows Server 2016, Microsoft removed the Identity Management for Unix (IDMU) and NIS Server Role. See Clarification regarding the status of Identity Management for Unix (IDMU) & NIS Server Role in Windows Server 2016 Technical Preview and beyond (https://blogs.technet.microsoft.com/activedirectoryua/2016/02/09/identity-management-for-unix-idmu-is-deprecated-in-windows-server/).

Figure 9.3 shows the configuration screen which opens after navigating *Directory Service* \rightarrow *NIS*. Table 9.4 summarizes the configuration options.

Directory Service					
Active Directory	LDAP	NIS	Kerberos Realms	Kerberos Keytabs	Kerberos Settings
NIS domain:			i		
NIS servers:			ì		
Secure mode:	i				
Manycast:	i				
Enable:					
Save	uild Director	ry Service Cac	he		

Fig. 9.3: NIS Configuration

Table 9.4: NIS Configuration Options

Setting	Value	Description
NIS domain	string	Name of NIS domain.
NIS servers	string	Comma-delimited list of hostnames or IP addresses.
Secure mode	checkbox	If set, ypbind(8) (https://www.freebsd.org/cgi/man.cgi?query=ypbind) will refuse to bind to any NIS server that is not running as root on a TCP port number over 1024.
Manycast	checkbox	If set, ypbind will bind to the server that responds the fastest. This is useful when no local NIS server is available on the same subnet
Enable	checkbox	Unset to disable the configuration without deleting it.

Click the *Rebuild Directory Service Cache* button after adding a user to NIS who needs immediate access to FreeNAS[®]. Otherwise this occurs automatically once a day as a cron job.

9.4 Kerberos Realms

A default Kerberos realm is created for the local system in FreeNAS[®]. *Directory Service* \rightarrow *Kerberos Realms* can be used to view and add Kerberos realms. If the network contains a KDC, click *Add kerberos realm* to add the realm. This configuration

screen is shown in Figure 9.4.

Add kerberos	realm	8	
Realm:	<	Kerbe	eros realm.
ОК Са	Advanced Mode)	

Fig. 9.4: Adding a Kerberos Realm

Table 9.5 summarizes the configurable options. Some settings are only available in Advanced Mode. To see these settings, either click *Advanced Mode* or configure the system to always display these settings by checking the box *Show advanced fields* by default in *System* \rightarrow *Advanced*.

Setting	Value	Advanced Mode	Description
Realm	string		Mandatory. Name of the realm.
KDC	string	\checkmark	Name of the Key Distribution Center.
Admin Server	string	\checkmark	Server where all changes to the database are performed.
Password Server	string	\checkmark	Server where all password changes are performed.

Table 9.5:	Kerheros	Realm	Ontions
Table 9.5.	Keibei 05	кеаши	Options

9.5 Kerberos Keytabs

Kerberos keytabs are used to do Active Directory or LDAP joins without a password. This means the password for the Active Directory or LDAP administrator account does not need to be saved into the FreeNAS[®] configuration database, which is a security risk in some environments.

When using a keytab, it is recommended to create and use a less privileged account for performing the required queries as the password for that account will be stored in the FreeNAS[®] configuration database. To create the keytab on a Windows system, use the ktpass (https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass) command:

where:

- freenas.keytab is the file to upload to the $\ensuremath{\mathsf{FreeNAS}}^{\ensuremath{\$}}$ server.
- useraccount is the name of the user account for the FreeNAS[®] server generated in Active Directory Users and Computers (https://technet.microsoft.com/en-us/library/aa998508(v=exchg.65).aspx).
- http/useraccount@EXAMPLE.COM is the principal name written in the format *host/user.account@KERBEROS.REALM*. By convention, the kerberos realm is written in all caps, but make sure the case used for the *Kerberos Realm* (page 177) matches the realm name. See this note (https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/ktpass#BKMK_remarks) about using /princ for more details.
- userpass is the password associated with useraccount.

Setting /crypto to ALL allows using all supported cryptographic types. These keys can be specified instead of ALL:

• DES-CBC-CRC is used for compatibility.

- DES-CBC-MD5 adheres more closely to the MIT implementation and is used for compatibility.
- *RC4-HMAC-NT* uses 128-bit encryption.
- AES256-SHA1 uses AES256-CTS-HMAC-SHA1-96 encryption.
- AES128-SHA1 uses AES128-CTS-HMAC-SHA1-96 encryption.

This will create a keytab with sufficient privileges to grant tickets.

After the keytab is generated, use *Directory Service* \rightarrow *Kerberos Keytabs* \rightarrow *Add kerberos keytab* to add it to the FreeNAS[®] system.

To instruct the Active Directory service to use the keytab, select the installed keytab using the drop-down *Kerberos keytab* menu in *Directory Service* \rightarrow *Active Directory*. When using a keytab with Active Directory, make sure that the "username" and "userpass" in the keytab matches the "Domain Account Name" and "Domain Account Password" fields in *Directory Service* \rightarrow *Active Directory*.

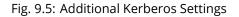
To instruct LDAP to use a principal from the keytab, select the principal from the drop-down *Kerberos Principal* menu in *Directory Service* \rightarrow *LDAP*.

9.6 Kerberos Settings

To configure additional Kerberos parameters, use *Directory Service* \rightarrow *Kerberos Settings*. Figure 9.5 shows the fields available:

- **Appdefaults auxiliary parameters:** contains settings used by some Kerberos applications. The available settings and their syntax are listed in the [appdefaults] section of krb.conf(5) (http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#appdefaults).
- Libdefaults auxiliary parameters: contains settings used by the Kerberos library. The available settings and their syntax are listed in the [libdefaults] section of krb.conf(5) (http://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html#libdefaults).

Directory Service					
Active Directory	LDAP	NIS	Kerberos Realms	Kerberos Keytabs	Kerberos Settings
Appdefaults au	uxiliary param	neters:			
Libdefaults au	xiliary parame	eters:			
Save					



SHARING

Shares are created to make part or all of a volume accessible to other computers on the network. The type of share to create depends on factors like which operating systems are being used by computers on the network, security requirements, and expectations for network transfer speeds.

FreeNAS[®] provides a *Wizard* (page 281) for creating shares. The *Wizard* (page 281) automatically creates the correct type of dataset and permissions for the type of share, sets the default permissions for the share type, and starts the service needed by the share. It is recommended to use the Wizard to create shares, fine-tune the share settings using the instructions in the rest of this chapter if needed, then fine-tune the default permissions from the client operating system to meet the requirements of the network.

Note: Shares are created to provide and control access to an area of storage. Before creating shares, making a list of the users that need access to storage data, which operating systems these users are using, whether all users should have the same permissions to the stored data, and whether these users should authenticate before accessing the data is recommended. This information can help determine which type of shares are needed, whether multiple datasets are needed to divide the storage into areas with different access and permissions, and how complex it will be to set up those permission requirements. Note that shares are used to provide access to data. When a share is deleted, it removes access to data but does not delete the data itself.

These types of shares and services are available:

- *AFP* (page 181): Apple File Protocol shares are often used when the client computers all run macOS. Apple has slowly shifted to preferring *SMB* (page 195) for modern networks, although Time Machine still requires AFP.
- Unix (NFS) (page 188): Network File System shares are accessible from Mac OS X, Linux, BSD, and the professional and enterprise versions (but not the home editions) of Windows. This can be are a good choice when the client computers do not all run the same operating system but NFS client software is available for all of them.
- *WebDAV* (page 194): WebDAV shares are accessible using an authenticated web browser (read-only) or WebDAV client (https://en.wikipedia.org/wiki/WebDAV#Client_support) running on any operating system.
- SMB (page 195): Server Message Block shares, also known as Common Internet File System (CIFS) shares, are accessible by Windows, Mac OS X, Linux, and BSD computers. Access is slower than an NFS share due to the single-threaded design of Samba. SMB provides more configuration options than NFS and is a good choice on a network for Windows systems. However, it is a poor choice if the CPU on the FreeNAS[®] system is limited; if the CPU is maxed out, upgrade the CPU or consider another type of share.
- *Block (iSCSI)* (page 205): block or iSCSI shares appear as an unformatted disk to clients running iSCSI initiator software or a virtualization solution such as VMware. These are usually used as virtual drives.

Fast access from any operating system can be obtained by configuring the *FTP* (page 225) service instead of a share and using a cross-platform FTP file manager application such as Filezilla (https://filezilla-project.org/). Secure FTP can be configured if the data needs to be encrypted.

When data security is a concern and the network users are familiar with SSH command line utilities or WinSCP (https://winscp.net/eng/index.php), consider using the *SSH* (page 245) service instead of a share. It is slower than unencrypted FTP due to the encryption overhead, but the data passing through the network is encrypted. **Note:** It is generally a mistake to share a volume or dataset with more than one share type or access method. Different types of shares and services use different file locking methods. For example, if the same volume is configured to use both NFS and FTP, NFS will lock a file for editing by an NFS user, but an FTP user can simultaneously edit or delete that file. This results in lost edits and confused users. Another example: if a volume is configured for both AFP and SMB, Windows users can be confused by the "extra" filenames used by Mac files and delete them. This corrupts the files on the AFP share. Pick the one type of share or service that makes the most sense for the types of clients accessing that volume, and use that single type of share or service. To support multiple types of shares, divide the volume into datasets and use one dataset per share.

This section demonstrates configuration and fine-tuning of AFP, NFS, SMB, WebDAV, and iSCSI shares. FTP and SSH configurations are described in *Services* (page 219).

10.1 Apple (AFP) Shares

FreeNAS[®] uses the Netatalk (http://netatalk.sourceforge.net/) AFP server to share data with Apple systems. This section describes the configuration screen for fine-tuning AFP shares created using the *Wizard* (page 281). It then provides configuration examples for using the *Wizard* (page 281) to create a guest share, configuring Time Machine to back up to a dataset on the FreeNAS[®] system, and for connecting to the share from a macOS client.

To view the AFP share created by the Wizard, click *Sharing* \rightarrow *Apple (AFP)* and highlight the name of the share. Click its *Edit* button to see the configuration options shown in Figure 10.1. The values showing for these options will vary, depending upon the information given when the share was created.

Sharing					
Apple (AFP) UNIX (NFS) WebDAV Windows ((SMB) Block (iSCSI)				
Add Apple (AFP) Share					
Path No entry has been found		Name			Share Comment
_	Add Apple (AFP) Share Path: /mnt/v Use as home share: Name: afp1 Time Machine: Auxiliary Parameters: OK Cancel)	Browse	∞	

Fig. 10.1: Creating an AFP Share

Note: Table 10.1 summarizes the options available to fine-tune an AFP share. Leaving these options at the default settings is recommended as changing them can cause unexpected behavior. Most settings are only available with *Advanced Mode*. Do **not** change an advanced option without fully understanding the function of that option. Refer to Setting up Netatalk (http://netatalk.sourceforge.net/2.2/htmldocs/configuration.html) for a more detailed explanation of these options.

Setting	Value	Advanced Mode	Description
Path	browse but- ton		<i>Browse</i> to the volume/dataset to share. Do not nest additional volumes, datasets, or symbolic links beneath this path. Ne-tatalk does not fully support nesting functionality.
Use as home share	checkbox		Set to allow the share to host user home directories. Only one share can be used as the home share.
Name	string		Enter the volume name that appears in in macOS after select- ing $Go \rightarrow Connect$ to server in the Finder menu. Limited to 27 characters and cannot contain a period.
Share Comment	string	 ✓ 	Enter an optional comment.
Allow List	string	√	Comma-delimited list of allowed users and/or groups where groupname begins with a @. Note that adding an entry will deny any user/group that is not specified.
Deny List	string	\checkmark	Comma-delimited list of denied users and/or groups where groupname begins with a @. Note that adding an entry will al- low all users/groups that are not specified.
Read-only Access	string	\checkmark	Comma-delimited list of users and/or groups who only have read access where groupname begins with a @.
Read-write Access	string	\checkmark	Comma-delimited list of users and/or groups who have read and write access where groupname begins with a @.
Time Machine	checkbox		Set to advertise FreeNAS [®] as a Time Machine disk so it can be found by Macs. Setting multiple shares for Time Machine use is not recommended. When multiple Macs share the same pool, low diskspace issues and intermittently failed backups can oc- cur.
Zero Device Numbers	checkbox	\checkmark	Enable when the device number is not constant across a reboot.
No Stat	checkbox	~	If enabled, AFP does not stat the volume path when enumerat- ing the volumes list. Useful for automounting or volumes cre- ated by a preexec script.
AFP3 UNIX Privs	checkbox	√	Set to enable Unix privileges supported by OSX 10.5 and higher. Do not enable this if the network contains macOS 10.4 clients or lower. Those systems do not support this feature.
Default file permis- sion	checkboxes	\checkmark	Only works with Unix ACLs. New files created on the share are set with the selected permissions.
Default directory per- mission	checkboxes	\checkmark	Only works with Unix ACLs. New directories created on the share are set with the selected permissions.
Default umask	integer	\checkmark	Umask is used for newly created files. Default is 000 (anyone can read, write, and execute).
Hosts Allow	string	\checkmark	Enter a list of allowed hostnames or IP addresses. Separate entries with a comma, space, or tab.
Hosts Deny	string	\checkmark	Enter a list of denied hostnames or IP addresses. Separate en- tries with a comma, space, or tab.
Auxiliary Parameters	string		Additional afp.conf (https://www.freebsd.org/cgi/man.cgi?query=afp.conf) parameters not covered by other option fields.

Table 10.1: AFP Share Configuration Options

10.1.1 Creating AFP Guest Shares

AFP supports guest logins, meaning that macOS users can access the AFP share without requiring their user accounts to first be created on or imported into the FreeNAS[®] system.

Note: When a guest share is created along with a share that requires authentication, AFP only maps users who log in as *guest* to the guest share. If a user logs in to the share that requires authentication, permissions on the guest share can prevent that user from writing to the guest share. The only way to allow both guest and authenticated users to write to a guest share is to set the permissions on the guest share to 777 or to add the authenticated users to a guest group and set the permissions to 77*x*.

Before creating a guest share, go to Services \rightarrow AFP and make sure that the Guest Access option is enabled.

To create the AFP guest share, click *Wizard*, then click the *Next* button twice to display the screen shown in Figure 10.2. Complete these fields in this screen:

- 1. **Share name:** enter a name for the share that is identifiable but less than 27 characters long. This name cannot contain a period. In this example, the share is named *afp_guest*.
- 2. Click the button for Mac OS X (AFP).
- 3. Click the *Ownership* button. Click the drop-down *User* menu and select *nobody*. Click the *Return* button to return to the previous screen.
- 4. Click the *Add* button. **The share is not created until the button is clicked**. Clicking the *Add* button adds an entry to the *Name* frame with the name that was entered in *Share name*.

Wizard	
Share name: afp_guest Purpose Windows (SMB) Allow Guest Mac OS X (AFP) Time Machine Generic Unix (NFS) Block Storage (iSCSI) Size:	
Add Delete Update	1
Name	
afp_guest	
Previous Next Exit	

Fig. 10.2: Creating a Guest AFP Share

Click the *Next* button twice, then the *Confirm* button to create the share. The Wizard automatically creates a dataset for the share that contains the correct default permissions and starts the AFP service so the share is immediately available. The new share is also added as an entry to *Sharing* \rightarrow *Apple (AFP)*.

macOS users can connect to the guest AFP share by navigating $Go \rightarrow Connect$ to Server. In the example shown in Figure 10.3, the user entered *afp://* followed by the IP address of the FreeNAS[®] system.

Click the *Connect* button. Once connected, Finder opens automatically. The name of the AFP share is displayed in the SHARED section in the left frame and the contents of any data saved in the share is displayed in the right frame.

🛒 Finde	er File Edit View Go Window Help	
	O Connect to Server	
	Server Address:	
	afp://192.168.2.2	+ 💽 🕴
	Favorite Servers:	
	(?) Remove Browse	Connect
•		
		A DECEMBER OF

Fig. 10.3: Connect to Server Dialogue

To disconnect from the volume, click the *eject* button in the *Shared* sidebar.

10.1.2 Creating Authenticated and Time Machine Shares

macOS includes the Time Machine application which can be used to schedule automatic backups. In this configuration example, a Time Machine user will be configured to backup to an AFP share on a FreeNAS[®] system. Creating a separate Time Machine share for each user that will be using Time Machine to backup their macOS system to FreeNAS[®] is recommended. The process for creating an authenticated share for a user is the same as creating a Time Machine share for that user.

To use the Wizard to create an authenticated or Time Machine share, enter the following information, as seen in the example in Figure 10.4.

- 1. **Share name:** enter a name for the share that is identifiable but less than 27 characters long. The name cannot contain a period. In this example, the share is named *backup_user1*.
- 2. Click the button for *Mac OS X (AFP)* and enable the *Time Machine* option.
- 3. Click the *Ownership* button. If the user already exists on the FreeNAS[®] system, click the drop-down *User* menu to select their user account. If the user does not yet exist on the FreeNAS[®] system, type their name into the *User* field and enable the *Create User* option. If the user will be a member of a group that already exists on the FreeNAS[®] system, click the drop-down *Group* menu to select the group name. To create a new group to be used by Time Machine users, enter the name in the *Group* field and set the *Create Group* option. Otherwise, enter the same name as the user. In the example shown in Figure 10.5, both a new *user1* user and a new *tm_backups* group will be created. Since a new user is being created, this screen prompts for the user password to be used when accessing the share. It also provides an

opportunity to change the default permissions on the share. When finished, click *Return* to return to the screen shown in Figure 10.4.

4. Click the *Add* button. **Remember to do this or the share will not be created**. Clicking the *Add* button adds an entry to the *Name* frame with the name that was entered in *Share name*.

To configure multiple authenticated or Time Machine shares, repeat for each user, giving each user their own *Share name* and *Ownership*. When finished, click the *Next* button twice, then the *Confirm* button to create the shares. The Wizard automatically creates a dataset for each share with the correct ownership and starts the AFP service so the shares are immediately available. The new shares will appear in *Sharing* \rightarrow *Apple (AFP)*.

Wizard	8
Share name: backup_user1	
Purpose Windows (SMB) Allow Guest Mac OS X (AFP) Time Machine Generic Unix (NFS) Ownership)
Block Storage (iSCSI) Size:	
Add Delete Update	
Name	
backup_user1	-
Previous Next Exit	Y

Fig. 10.4: Creating a Time Machine Share

Wizard		38
User:	user1 💌	🔽 Create User (i
User Password:	•••••	
Confirm User Password:	•••••	
Group:	tm_backups	🔽 Create Group 👔
Mode:	Owner Group Other Read Image: Comparison of the comparison of	
Return		

Fig. 10.5: Creating an Authenticated User

At this point, it may be desirable to configure a quota for each Time Machine share, to restrict backups from using all of the available space on the FreeNAS[®] system. The first time Time Machine makes a backup, it will create a full backup after waiting two minutes. It will then create a one hour incremental backup for the next 24 hours, and then one backup each day, each week and each month. **Since the oldest backups are deleted when a Time Machine share becomes full, make sure that the quota size is sufficient to hold the desired number of backups.** Note that a default installation of macOS is ~21 GiB in size.

To configure a quota, go to *Storage* \rightarrow *Volumes* and highlight the entry for the share. In the example shown in Figure 10.6, the Time Machine share name is *backup_user1*. Click the *Edit Options* button for the share, then *Advanced Mode*. Enter a value in the *Quota for this dataset* field, then click *Edit Dataset* to save the change. In this example, the Time Machine share is restricted to 200 GiB.

Storage	Edit Options		8
Volumes Periodic Snapsh			
Volume Manager Import Dis	Dataset: volume1/backup_user1		
	Compression level:	Inherit (lz4)	
Name volume1	Share type:	Mac	
✓ volume1 backup_user1	Enable atime:	 Inherit (on) On Off 	
	Quota for this dataset:	200GiB	
	Quota for this dataset and all children:	0	
	Reserved space for this dataset:	0	
	Reserved space for this dataset and all children:	0	
	ZFS Deduplication:	Enabling dedup may have drastic performance implications,	Ļ

Fig. 10.6: Setting a Quota

Note: An alternative is to create a global quota using the instructions in Set up Time Machine for multiple machines with OSX Server-Style Quotas (https://forums.freenas.org/index.php?threads/how-to-set-up-time-machine-for-multiple-machines-with-osx-server-style-quotas.47173/).

To configure Time Machine on the macOS client, go to *System Preferences* \rightarrow *Time Machine* which opens the screen shown in Figure 10.7. Click ON and a pop-up menu shows the FreeNAS[®] system as a backup option. In this example, it is listed as *backup_user1 on "freenas"*. Highlight the FreeNAS[®] system and click *Use Backup Disk*. A connection bar opens and prompts for the user account's password–in this example, the password that was set for the *user1* account.



Fig. 10.7: Configuring Time Machine on Mac OS X Lion

If Time Machine could not complete the backup. The backup disk image could not be created (error 45) is shown when backing up to the FreeNAS[®] system, a sparsebundle image must be created using these instructions (https://community.netgear.com/t5/Stora-Legacy/Solution-to-quot-Time-Machine-could-not-complete-the-backup/td-p/294697).

If Time Machine completed a verification of your backups. To improve reliability, Time Machine must create a new backup for you. is shown, follow the instructions in this post (http://www.garth.org/archives/2011,08,27,169,fix-time-machine-sparsebundle-nas-based-backup-errors.html) to avoid making another backup or losing past backups.

10.2 Unix (NFS) Shares

FreeNAS[®] supports sharing over the Network File System (NFS). Clients use the mount command to mount the share. Once mounted, the NFS share appears as just another directory on the client system. Some Linux distros require the installation of additional software in order to mount an NFS share. On Windows systems, enable Services for NFS in the Ultimate or Enterprise editions or install an NFS client application.

Note: For performance reasons, iSCSI is preferred to NFS shares when FreeNAS[®] is installed on ESXi. When considering creating NFS shares on ESXi, read through the performance analysis presented in Running ZFS over NFS as a VMware Store (https://tinyurl.com/archive-zfs-over-nfs-vmware).

To create an NFS share using the *Wizard* (page 281), click the *Next* button twice to display the screen shown in Figure 10.8. Enter a *Share name*. Spaces are not allowed in these names. Click the button for *Generic Unix (NFS)*, then click *Add* so the share name appears in the *Name* frame. When finished, click the *Next* button twice, then the *Confirm* button to create the share. Creating an NFS share using the wizard automatically creates a new dataset for the share, starts the services required for NFS, and adds an entry in *Sharing* \rightarrow *Unix (NFS) Shares*. Depending on the requirements, the IP addresses that are allowed to access the NFS share can be restricted, or the permissions adjusted.

Wizard	36
Share name: nfs_share1 Purpose Windows (SMB) Allow Guest Mac OS X (AFP) Time Machine Generic Unix (NFS) Ownership	
O Block Storage (iSCSI) Size: Add Delete	
Name nfs_share1	-
	*
Previous Next Exit	

Fig. 10.8: NFS Share Wizard

NFS shares are edited by clicking *Sharing* \rightarrow *Unix (NFS)*, highlighting the entry for the share, and clicking the *Edit* button. In the example shown in Figure 10.9, the configuration screen is open for the *nfs_share1* share.

Sharing					
Apple (AF	P) UNIX (NFS) We	bDAV Windo	ws (CIFS) Block	(iSCSI)	
Add Unix (NFS) Share				
Paths		Comment	All Directories	Read Only	
/mnt/volu	me1/nfs_share1	nfs_share1	false	false	
. E		_	_		
	dit		_	28	
	Path			_	
	Path: /mn	t/volume1/nfs	_share1	Browse	
	Delete:				
	Add extra Path				
	Comment:	nfs_share1			
	All Directories:	Ì		_	
	Read Only:	i			
Edit	ОК Cancel	Delete	anced Mode		

Fig. 10.9: NFS Share Settings

Table 10.2 summarizes the available configuration options in this screen. Some settings are only available by clicking the *Advanced Mode* button.

Table	10.2.	NFS	Share	Options
TUDIC	10.2.		Jun	Options

Setting	Value	Advanced	Description
		Mode	
Path	browse but-		Browse to the volume or dataset to be shared. Click Add extra
	ton		<i>path</i> to select multiple paths.
Comment	string		Set the share name. If left empty, share name is the list of se-
			lected Path entries.
Authorized networks	string	\checkmark	Space-delimited list of allowed networks in network/mask CIDR
			notation. Example: 1.2.3.0/24. Leave empty to allow all.
Authorized IP ad-	string	 ✓ 	Space-delimited list of allowed IP addresses or hostnames.
dresses or hosts			Leave empty to allow all.
All directories	checkbox		Set to allow the client to mount any subdirectory within the
			Path.
Read only	checkbox		Set to prohibit writing to the share.
Quiet	checkbox	\checkmark	Set to inhibit some syslog diagnostics to
			avoid some error messages. See exports(5)
			(https://www.freebsd.org/cgi/man.cgi?query=exports) for
			examples.
Maproot User	drop-down	\checkmark	When a user is selected, the <i>root</i> user is limited to permissions
	menu		of that user.
Maproot Group	drop-down	\checkmark	When a group is selected, the <i>root</i> user is also limited to per-
	menu		missions of that group.

Continued on next page

Setting	Value	Advanced	Description
		Mode	
Mapall User	drop-down	\checkmark	All clients use the permissions of the specified user.
	menu		
Mapall Group	drop-down	\checkmark	All clients use the permissions of the specified group.
	menu		
Security	selection	V	Only appears if <i>Enable NFSv4</i> is enabled in <i>Services</i> \rightarrow <i>NFS</i> . Choices are <i>sys</i> or these Kerberos options: <i>krb5</i> (authentication only), <i>krb5i</i> (authentication and integrity), or <i>krb5p</i> (authentication and privacy). If multiple security mechanisms are added to the <i>Selected</i> column using the arrows, use the <i>Up</i> or <i>Down</i> buttons to list in order of preference.

Table 10.2 – continued from previous page

When creating NFS shares, keep these points in mind:

- 1. Clients will specify the Path when mounting the share.
- 2. The *Maproot* and *Mapall* options are exclusive, meaning only one can be used was the GUI does not allow both. The *Mapall* options supersede the *Maproot* options. To restrict only the *root* user's permissions, set the *Maproot* option. To restrict permissions of all users, set the *Mapall* options.
- 3. Each volume or dataset is considered to be its own filesystem and NFS is not able to cross filesystem boundaries.
- 4. The network and host must be unique per share and per filesystem or directory. Since /etc/exports does not act like an ACL, the rule to apply is undefined among overlapping networks or when using the same share with multiple hosts.
- 5. The *All directories* option can only be used once per share per filesystem.

To better understand these restrictions, consider a scenario where there are:

- two networks, 10.0.0.0/8 and 20.0.0.0/8
- a ZFS volume named volume1 with 2 datasets named dataset1 and dataset2
- dataset1 contains a directory named directory1

Because of restriction #3, an error is shown when trying to create one NFS share like this:

- Authorized networks set to 10.0.0.0/8 20.0.0.0/8
- Path set to /mnt/volume1/dataset1 and /mnt/volume1/dataset1/directory1

Instead, set a Path of /mnt/volume1/dataset1 and check the All directories box.

That directory could also be restricted to one of the networks by creating two shares instead:

First NFS share:

- Authorized networks set to 10.0.0/8
- Path set to /mnt/volume1/dataset1

Second NFS share:

- Authorized networks set to 20.0.0.0/8
- Path set to /mnt/volume1/dataset1/directory1

Note that this requires the creation of two shares. It cannot be done with only one share.

10.2.1 Example Configuration

By default, the *Mapall* fields are not set. This means that when a user connects to the NFS share, the user has the permissions associated with their user account. This is a security risk if a user is able to connect as *root* as they will have complete access to the share.

A better option is to do this:

- 1. Specify the built-in *nobody* account to be used for NFS access.
- 2. In the *Change Permissions* screen of the volume/dataset that is being shared, change the owner and group to *nobody* and set the permissions according to the desired requirements.
- 3. Select *nobody* in the *Mapall User* and *Mapall Group* drop-down menus for the share in *Sharing* \rightarrow *Unix (NFS) Shares*.

With this configuration, it does not matter which user account connects to the NFS share, as it will be mapped to the *nobody* user account and will only have the permissions that were specified on the volume/dataset. For example, even if the *root* user is able to connect, it will not gain *root* access to the share.

10.2.2 Connecting to the Share

The following examples share this configuration:

- 1. The FreeNAS[®] system is at IP address *192.168.2.2*.
- 2. A dataset named /mnt/volume1/nfs_share1 is created and the permissions set to the *nobody* user account and the *nobody* group.
- 3. An NFS share is created with these attributes:
 - Path: /mnt/volume1/nfs_share1
 - Authorized Networks: 192.168.2.0/24
 - All Directories option is enabled
 - MapAll User is set to nobody
 - MapAll Group is set to nobody

10.2.2.1 From BSD or Linux

NFS shares are mounted on BSD or Linux clients with this command executed as the superuser (root) or with sudo:

mount -t nfs 192.168.2.2:/mnt/volume1/nfs_share1 /mnt

- **-t nfs** specifies the filesystem type of the share
- 192.168.2.2 is the IP address of the FreeNAS® system
- · /mnt/volume/nfs_share1 is the name of the directory to be shared, a dataset in this case
- **/mnt** is the mountpoint on the client system. This must be an existing, *empty* directory. The data in the NFS share appears in this directory on the client computer.

A successful mounting of the share returns to the command prompt without any status or error messages.

Note: If this command fails on a Linux system, make sure that the nfs-utils (https://sourceforge.net/projects/nfs/files/nfs-utils/) package is installed.

This configuration allows users on the client system to copy files to and from /mnt (the mount point). All files are owned by *nobody:nobody*. Changes to any files or directories in /mnt are written to the FreeNAS[®] system's /mnt/volume1/nfs_share1 dataset.

Settings cannot be changed on the NFS share if it is mounted on any client computers. The umount command is used to unmount the share on BSD and Linux clients. Run it as the superuser or with sudo on each client computer:

umount /mnt

10.2.2.2 From Microsoft

Windows NFS client support varies with versions and releases. For best results, use Windows (SMB) Shares (page 195).

10.2.2.3 From macOS

To mount the NFS volume from a macOS client, go to $Go \rightarrow Connect$ to Server. In the Server Address field, enter nfs:// followed by the IP address of the FreeNAS[®] system and the name of the volume/dataset being shared by NFS. The example shown in Figure 10.10 continues with our example of 192.168.2.2:/mnt/volume1/nfs_share1.

Finder opens automatically after connecting. The IP address of the FreeNAS[®] system is displayed in the SHARED section in the left frame and the contents of the share are displayed in the right frame. In the example shown in Figure 10.11, /mnt/data has one folder named images. The user can now copy files to and from the share.

	Connect to Serv	/er	
Server Address:			
nfs://192.168.2.2:/mnt/vol	ume1/nfs_share1		+ @~
Favorite Servers:			_
? Remove		Browse	Connect
? Remove		Browse	Connect

Fig. 10.10: Mounting the NFS Share from macOS

• • •	mfs_share1
$\langle \rangle$	
Favorites	Earlier
All My Files	
MirDrop	
Applications	
Desktop	Images
Downloads	
Movies	
Music	
Dictures	
Creative Clou	
iCloud	
Cloud Drive	
Desktop	
Documents	
Devices	
Remote Disc	
Shared	
□ 192.168.2.2 ≜	

Fig. 10.11: Viewing the NFS Share in Finder

10.2.3 Troubleshooting NFS

Some NFS clients do not support the NLM (Network Lock Manager) protocol used by NFS. This is the case if the client receives an error that all or part of the file may be locked when a file transfer is attempted. To resolve this error, add the option **-o nolock** when running the mount command on the client to allow write access to the NFS share.

If a "time out giving up" error is shown when trying to mount the share from a Linux system, make sure that the portmapper service is running on the Linux client. If portmapper is running and timeouts are still shown, force the use of TCP by including **-o tcp** in the mount command.

If a RPC: Program not registered error is shown, upgrade to the latest version of FreeNAS[®] and restart the NFS service after the upgrade to clear the NFS cache.

If clients see "reverse DNS" errors, add the FreeNAS[®] IP address in the Host name data base field of Network \rightarrow Global Configuration.

If clients receive timeout errors when trying to mount the share, add the client IP address and hostname to the Host name data base field in Network \rightarrow Global Configuration.

Some older versions of NFS clients default to UDP instead of TCP and do not auto-negotiate for TCP. By default, FreeNAS[®] uses TCP. To support UDP connections, go to Services \rightarrow NFS and enable the Serve UDP NFS clients option.

The nfsstat -c or nfsstat -s commands can be helpful to detect problems from the *Shell* (page 289). A high proportion of retries and timeouts compared to reads usually indicates network problems.

10.3 WebDAV Shares

In FreeNAS[®], WebDAV shares can be created so that authenticated users can browse the contents of the specified volume, dataset, or directory from a web browser.

Configuring WebDAV shares is a two step process. First, create the WebDAV shares to specify which data can be accessed. Then, configure the WebDAV service by specifying the port, authentication type, and authentication password. Once the configuration is complete, the share can be accessed using a URL in the format:

protocol://IP_address:port_number/share_name

where:

- **protocol:** is either *http* or *https*, depending upon the *Protocol* configured in *Services* \rightarrow *WebDAV*.
- **IP address:** is the IP address or hostname of the FreeNAS[®] system. Take care when configuring a public IP address to ensure that the network firewall only allows access to authorized systems.
- port_number: is configured in Services → WebDAV. If the FreeNAS[®] system is to be accessed using a public IP address, consider changing the default port number and ensure that the network's firewall only allows access to authorized systems.
- **share_name:** is configured in *Sharing* → *WebDAV Shares*.

Entering the URL in a web browser brings up an authentication pop-up message. Enter a username of *webdav* and the password configured in *Services* \rightarrow *WebDAV*.

Warning: At this time, only the *webdav* user is supported. For this reason, it is important to set a good password for this account and to only give the password to users which should have access to the WebDAV share.

To create a WebDAV share, click *Sharing* \rightarrow *WebDAV Shares* \rightarrow *Add WebDAV Share* which will open the screen shown in Figure 10.12.

A	dd WebDAV Share		88
			_
	Share Name:	!	(i)
	Comment:		
	Path:		Browse
	Read Only:		
	Change User & Group Ownership:	i	
	OK Cancel		

Fig. 10.12: Adding a WebDAV Share

Table 10.3 summarizes the available options.

Setting	Value	Description
Share Path Name	string	Enter a name for the share.
Comment	string	Optional.
Path	browse button	<i>Browse</i> to the volume/dataset to share.
Read Only	checkbox	Set to prohibit users from writing to the share.
Change User &	checkbox	Enable to automatically set the share contents to the <i>webdav</i> user and
Group Ownership		group.

Table 10.3: WebDAV Share Options

After clicking *OK*, a pop-up asks about enabling the service. Once the service starts, review the settings in *Services* \rightarrow *WebDAV* as they are used to determine which URL is used to access the WebDAV share and whether or not authentication is required to access the share. These settings are described in *WebDAV* (page 252).

10.4 Windows (SMB) Shares

FreeNAS[®] uses Samba (https://www.samba.org/) to share volumes using Microsoft's SMB protocol. SMB is built into the Windows and macOS operating systems and most Linux and BSD systems pre-install the Samba client in order to provide support for SMB. If the distro did not, install the Samba client using the distro software repository.

The SMB protocol supports many different types of configuration scenarios, ranging from the simple to complex. The complexity of the scenario depends upon the types and versions of the client operating systems that will connect to the share, whether the network has a Windows server, and whether Active Directory is being used. Depending on the authentication requirements, it might be necessary to create or import users and groups.

Samba supports server-side copy of files on the same share with clients from Windows 8 and higher. Copying between two different shares is not server-side. Windows 7 clients support server-side copying with Robocopy (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/cc733145(v=ws.11)).

This chapter starts by summarizing the available configuration options. It demonstrates some common configuration scenarios as well as offering some troubleshooting tips. Reading through this entire chapter before creating any SMB shares is recommended to gain a better understanding of the configuration scenario that meets the specific network requirements.

Warning: SMB1 is disabled by default for security reasons. If legacy clients are unable to connect to the share, open *Shell* (page 289), type sysctl freenas.services.smb.config.server_min_protocol=NT1, then restart the *SMB* (page 238) service. If that resolves the issue, go to *Tunables* (page 77) and creating a tunable with a *Variable* of *freenas.services.smb.config.server_min_protocol*, a *Value* of *NT1*, and a *Type* of *Sysctl*.

Tip: SMB Tips and Tricks (https://forums.freenas.org/index.php?resources/smb-tips-and-tricks.15/) shows helpful hints for configuring and managing SMB networking. The FreeNAS and Samba (CIFS) permissions (https://www.youtube.com/watch?v=RxggaE935PM) and Advanced Samba (CIFS) permissions on FreeNAS (https://www.youtube.com/watch?v=QhwOyLtArw0) videos clarify setting up permissions on SMB shares. Another helpful reference is Methods For Fine-Tuning Samba Permissions (https://forums.freenas.org/index.php?threads/methods-for-fine-tuning-samba-permissions.50739/).

Tip: Run smbstatus from the Shell (page 289) for a list of active connections and users.

Figure 10.13 shows the configuration screen that appears after clicking Sharing \rightarrow Windows (SMB Shares) \rightarrow Add Windows (SMB) Share.

Add Windows (SMB) Share	8
Path:	Browse
Use as home share:	
Name:	
Apply Default Permissions:	i
Allow Guest Access:	i
OK Cancel Advanced M	lode

Fig. 10.13: Adding an SMB Share

Table 10.4 summarizes the options when creating a SMB share. Some settings are only available after clicking the *Advanced Mode* button. For simple sharing scenarios, *Advanced Mode* options are not needed. For more complex sharing scenarios, only change an *Advanced Mode* option after fully understanding the function of that option. smb.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=smb.conf) provides more details for each configurable option.

Setting	Value	Advanced Mode	Description
Path	browse but- ton		Select volume/dataset/directory to share.
Use as home share	checkbox		Set to allow this share to hold user home directories. Only one share can be the home share. Note that lower case names for user home directories are strongly recommended, as Samba maps usernames to all lower case. For example, the username John will be mapped to a home directory named john. If the <i>Path</i> to the home share includes an upper case username, delete the existing user and recreate it in <i>Accounts</i> \rightarrow <i>Users</i> with an all lower case <i>Username</i> . Return to <i>Sharing</i> \rightarrow <i>SMB</i> to create the home share, and select the <i>Path</i> that contains the new lower case username.
Name	string		Enter a mandatory name for the share.
Comment	string	\checkmark	Optional description.
Apply Default Permis- sions	checkbox		When enabled, the ACLs grant read and write for owner or group and read-only for others. Only leave unset when creat- ing a share on a system that already has custom ACLs set.
Export Read Only	checkbox	\checkmark	Set to prohibit write access to the share.
Browsable to Net- work Clients	checkbox	\checkmark	Set for users to see the contents of /home. This includes other home directories of other users. When unset, users see only their own home directory.

Table 10.4: Options for a SMB Share

Continued on next page

Setting	Value	Advanced Mode	Description	
Export Recycle Bin	checkbox	✓	When set, deleted files are moved to a hidden .recycle in the root folder of the share. The .recycle directory can be deleted to reclaim space and is automatically recreated when a file is deleted.	
Show Hidden Files	checkbox	✓	Set to disable the Windows <i>hidden</i> attribute on a new Unix hid- den file. Unix hidden filenames start with a dot: .foo. Existing files are not affected.	
Allow Guest Access	checkbox		Set to allow access to this share without a password. See <i>SMB</i> (page 238) service for more information about guest user permissions.	
Only Allow Guest Ac- cess	checkbox	~	Requires <i>Allow guest access</i> to also be enabled. Forces guest access for all connections.	
Access Based Share Enumeration	checkbox	✓	When enabled, users can only see the shares they have per- mission to access. To change the default that grants everyone access, use the computer management MMC on Windows or the sharesec command-line utility.	
Hosts Allow	string	✓	Enter a list of allowed hostnames or IP addresses. Separate entries with a space, comma, or tab.	
Hosts Deny	string	✓	Enter a list of denied hostnames or IP addresses. Separate entries with a space, comma, or tab. Specify ALL and list any hosts from <i>Hosts Allow</i> to have those hosts take precedence.	
VFS Objects	selection		Adds virtual file system modules to enhance functionality. Ta- ble 10.5 summarizes the available modules.	
Periodic Snapshot Task	drop-down menu	\checkmark	Used to configure directory shadow copies on a per-share ba- sis. Select the pre-configured periodic snapshot task to use for the shadow copies of the share. Periodic snapshot must be re- cursive.	
Auxiliary Parameters	string	✓	Additional smb4.conf parameters not covered by other option fields.	

Table 10.4 – continued from previous page

Here are some notes about *ADVANCED MODE* settings:

- Hostname lookups add some time to accessing the SMB share. If only using IP addresses, unset the *Hostnames lookups* option in *Services* → *SMB*.
- When the *Browsable to Network Clients* option is enabled (the default), the share is visible through Windows File Explorer or through net view. When the *Use as a home share* option is selected, deselecting the *Browsable to Network Clients* option hides the share named *homes* so that only the dynamically generated share containing the authenticated user home directory will be visible. By default, the *homes* share and the user home directory are both visible. Users are not automatically granted read or write permissions on browsable shares. This option provides no real security because shares that are not visible in Windows File Explorer can still be accessed with a *UNC* path.
- If some files on a shared volume should be hidden and inaccessible to users, put a *veto files*= line in the *Auxiliary Parameters* field. The syntax for the *veto files* option and some examples can be found in the smb.conf manual page (https://www.freebsd.org/cgi/man.cgi?query=smb.conf).

Samba disables NTLMv1 authentication by default for security. Standard configurations of Windows XP and some configurations of later clients like Windows 7 will not be able to connect with NTLMv1 disabled. Security guidance for NTLMv1 and LM network authentication (https://support.microsoft.com/en-us/help/2793313/security-guidance-for-ntlmv1-and-lm-networkauthentication) has information about the security implications and ways to enable NTLMv2 on those clients. If changing the client configuration is not possible, NTLMv1 authentication can be enabled by enabling the *NTLMv1 auth* option in *Services* \rightarrow *SMB*.

Table 10.5 provides an overview of the available VFS modules. Be sure to research each module **before** adding or deleting it from the *Selected* column of the *VFS Objects* field of the share. Some modules need additional configuration after they are added. Refer to Stackable VFS modules (https://www.samba.org/samba/docs/old/Samba3-HOWTO/VFS.html) and the vfs_*

man pages (https://www.samba.org/samba/docs/current/man-html/) for more details.

	Table 10.5: Available VFS Modules
Value	Description
acl_tdb	Stores NTFS ACLs in a tdb file to enable full mapping of Windows ACLs.
acl_xattr	Stores NTFS ACLs in Extended Attributes (EAs) to en- able the full mapping of Windows ACLs.
aio_fork	Enables async I/O.
 aio_pthread	Implements async I/O in Samba vfs using a pthread pool instead of the internal Posix AlO interface.
audit	Logs share access, connects/disconnects, directory opens/creates/removes, and file opens/closes/renames/unlinks/chmods to syslog.
cacheprime	Primes the kernel file data cache.
сар	Translates filenames to and from the CAP encoding format, commonly used in Japanese language environ- ments.
catia	Improves Mac interoperability by translating charac- ters that are unsupported by Windows.
commit	Tracks the amount of data written to a file and syn- chronizes it to disk when a specified amount accumu- lates.
crossrename	Allows server side rename operations even if source and target are on different physical devices.
default_quota	Stores the default quotas that are reported to a win- dows client in the quota record of a user.
dfs_samba4	Distributed file system for providing an alternative name space, load balancing, and automatic failover.
dirsort	Sorts directory entries alphabetically before sending them to the client.
expand_msdfs	Enables support for Microsoft Distributed File System (DFS).
extd_audit	Sends <i>audit</i> logs to both syslog and the Samba log files.
fake_acls	Stores file ownership and ACLs as extended attributes.
fake_perms	Allows roaming profile files and directories to be set as read-only.
fruit	Enhances macOS support by providing the SMB2 AAPL extension and Netatalk interoperability. Automatically loads <i>catia</i> and <i>streams_xattr</i> but read the caveat in NOTE below table.
full_audit	Record selected client operations to the system log.
linux_xfs_sgid	Used to work around an old Linux XFS bug.
media_harmony	Allows Avid editorial workstations to share a network drive.
netatalk	Eases the co-existence of SMB and AFP shares.
offline	Marks all files in the share with the DOS <i>offline</i> at- tribute. This can prevent Windows Explorer from read- ing files just to make thumbnail images.
posix_eadb	Provides Extended Attributes (EAs) support so they can be used on filesystems which do not provide native support for EAs.
preopen	Useful for video streaming applications that want to read one file per frame.

Table 10.5: Available VFS Modules

Continued on next page

Table 10.5 – continued from previous page			
Value	Description		
readahead	Useful for Windows Vista clients reading data using		
	Windows Explorer.		
readonly	Marks a share as read-only for all clients connecting		
	within the configured time period.		
shadow_copy	Allows Microsoft shadow copy clients to browse		
	shadow copies on Windows shares.		
shadow_copy_test	Shadow copy testing.		
shell_snap	Provides shell-script callouts for snapshot creation and		
	deletion operations issued by remote clients using the		
	File Server Remote VSS Protocol (FSRVP).		
skel_opaque	Implements dummy versions of all VFS modules (use-		
	ful to VFS module developers).		
skel_transparent	Implements dummy passthrough functions of all VFS		
	modules (useful to VFS module developers).		
snapper	Provides the ability for remote SMB clients to access		
	shadow copies of FSRVP snapshots using Windows		
	Explorer.		
streams_depot	Experimental module to store alternate data streams		
	in a central directory. The association with the primary		
	file can be lost due to inode numbers changing when a		
	directory is copied to a new location (see https://marc.		
	info/?l=samba&m=132542069802160&w=2) .		
streams_xattr	Enables storing of NTFS alternate data streams in the		
	file system.		
syncops	Ensures metadata operations are performed syn-		
	chronously.		
time_audit	Logs system calls that take longer than the number of		
	defined milliseconds.		
unityed_media	Allows multiple Avid clients to share a network drive.		
winmsa	Emulate Microsoft's MoveSecurityAttributes=0 registry		
	option, setting the ACL for file and directory hierar-		
	chies to inherit from the parent directory into which		
	they are moved.		
worm	Controls the writability of files and folders depending		
vattr tdb	on their change time and an adjustable grace period.		
xattr_tdb	Stores Extended Attributes (EAs) in a tdb file so they		
	can be used on filesystems which do not provide sup- port for EAs.		
zfs_space	Correctly calculates ZFS space used by the share, in-		
zis_space	cluding space used by ZFS snapshots, quotas, and re-		
	sevations. Enabled by default.		
zfsacl	Provide ACL extensions for proper integration with		
2130(1	ZFS. Enabled by default.		
	ZI J. LHUNCU DY UEIGUIL.		

Table 10.5 – continued from previous page

Note: Be careful when using multiple SMB shares, some with and some without *fruit*. macOS clients negotiate SMB2 AAPL protocol extensions on the first connection to the server, so mixing shares with and without fruit will globally disable AAPL if the first connection occurs without fruit. To resolve this, all macOS clients need to disconnect from all SMB shares and the first reconnection to the server has to be to a fruit-enabled share.

These VFS objects do not appear in the selection box:

• **recycle:** moves deleted files to the recycle directory instead of deleting them. Controlled by *Export Recycle Bin* in the *SMB share options* (page 196).

• **shadow_copy2:** a more recent implementation of *shadow_copy* with some additional features. *shadow_copy2* and the associated parameters are automatically added to the smb4.conf when a *Periodic Snapshot Task* is selected.

10.4.1 Configuring Unauthenticated Access

SMB supports guest logins, meaning that users can access the SMB share without needing to provide a username or password. This type of share is convenient as it is easy to configure, easy to access, and does not require any users to be configured on the FreeNAS[®] system. This type of configuration is also the least secure as anyone on the network can access the contents of the share. Additionally, since all access is as the guest user, even if the user inputs a username or password, there is no way to differentiate which users accessed or modified the data on the share. This type of configuration is best suited for small networks where quick and easy access to the share is more important than the security of the data on the share.

To configure an unauthenticated SMB share, click *Wizard*, then click the *Next* button twice to display the screen shown in Figure 10.14. Complete the following fields in this screen:

- 1. **Share name:** enter a name for the share that is useful. In this example, the share is named *smb_insecure*.
- 2. Click the button for Windows (SMB) and enable the Allow Guest option.
- 3. Click the *Ownership* button. Click the drop-down *User* menu and select *nobody*. Click the *Return* button to return to the previous screen.
- 4. Click the *Add* button. **If this step is forgotten, the share will not be created**. Clicking the *Add* button adds an entry to the *Name* frame with the name that was entered in *Share name*.

Wizard	8
Share name: smb_insecure Purpose Windows (SMB) Allow Guest Mac OS X (AFP) Time Machine Generic Unix (NFS) Block Storage (iSCSI) Size:	
Add Delete Update	
smb_insecure	
Previous Next Exit	

Fig. 10.14: Creating an Unauthenticated SMB Share

Click the *Next* button twice, then the *Confirm* button to create the share. The Wizard automatically creates a dataset for the share and starts the SMB service so the share is immediately available. The new share will appear in *Sharing* \rightarrow *Windows (SMB)*.

Users can now access the share from any SMB client and will not be prompted for their username or password. For example, to access the share from a Windows system, open Explorer and click on *Network*. For this configuration example, a system named *FREENAS* appears with a share named *insecure_smb*. The user can copy data to and from the unauthenticated SMB share.

10.4.2 Configuring Authenticated Access Without a Domain Controller

Most configuration scenarios require each user to have their own user account and to authenticate before accessing the share. This allows the administrator to control access to data, provide appropriate permissions to that data, and to determine who accesses and modifies stored data. A Windows domain controller is not needed for authenticated SMB shares, which means that additional licensing costs are not required. However, since there is no domain controller to provide authentication for the network, each user account needs to be created on the FreeNAS[®] system. This type of configuration scenario is often used in home and small networks as it does not scale well if many users accounts are needed.

Before configuring this scenario, determine which users will need authenticated access. While not required for the configuration, it eases troubleshooting if the username and password that will be created on the FreeNAS[®] system matches that information on the client system. Next, determine if each user should have their own share to store their own data or if several users will be using the same share. The simpler configuration is to make one share per user as it does not require the creation of groups, adding the correct users to the groups, and ensuring that group permissions are set correctly.

To use the Wizard to create an authenticated SMB share, enter the following information, as shown in the example in Figure 10.15.

- 1. **Share name:** enter a name for the share that is useful. In this example, the share is named *smb_user1*.
- 2. Click the button for Windows (SMB).
- 3. Click the *Ownership* button. To create the user account on the FreeNAS[®] system, type their name into the *User* field and enable the *Create User* option. The user's password is then entered and confirmed. **If the user will not be sharing this share with other users**, type their name into the *Group* field and click *Create Group*. **If, however, the share will be used by several users**, instead type in a group name and enable the *Create Group* option. In the example shown in Figure 10.16, *user1* has been used for both the user and group name, meaning that this share will only be used by *user1*. When finished, click *Return* to return to the screen shown in Figure 10.15.
- 4. Click the *Add* button. **If this step is forgotten, the share will not be created**. Clicking the *Add* button adds an entry to the *Name* frame with the name that was entered in *Share name*.

When configuring multiple authenticated shares, repeat for each user, giving each user their own *Share name* and *Ownership*. When finished, click *Next* twice, then *Confirm* to create the shares. The Wizard automatically creates a dataset with the correct ownership for each share and starts the SMB service so the shares are available immediately. The new shares are also added to *Sharing* \rightarrow *Windows* (*SMB*).

Wizard 🛞
Share name: smb_user1 Purpose Windows (SMB) Allow Guest Mac OS X (AFP) Time Machine Generic Unix (NFS) Block Storage (iSCSI) Size:
Add Delete Update
Name
smb_user1
Previous Next Exit

Fig. 10.15: Creating an Authenticated SMB Share

Wi	zard		X
	User:	user1 💌	Create User (i)
	User Password:	•••••	
	Confirm User Password:	•••••	
	Group:	user1 💌	🔽 Create Group 🧃
	Mode:	Owner Group Other Read I I I I I I I I I I I I I I I I I I I	
	Return		

Fig. 10.16: Creating the User and Group

The authenticated share can now be tested from any SMB client. For example, to test an authenticated share from a Windows system, open Explorer and click on *Network*. For this configuration example, a system named *FREENAS* appears with a share named *smb_user1*. After clicking *smb_user1*, a Windows Security pop-up screen prompts for that user's username and password. Enter the values that were configured for that share, in this case user *user1*. After authentication, the user can copy data to and from the SMB share.

To prevent Windows Explorer from hanging when accessing the share, map the share as a network drive. To do this, rightclick the share and select *Map network drive...*. Choose a drive letter from the drop-down menu and click the *Finish* button.

Note that Windows systems cache a user's credentials. This can cause issues when testing or accessing multiple authenticated shares as only one authentication is allowed at a time. When authenticating to a share, if problems occur and the username and password are correct, type cmd in the *Search programs and files* box and use the following command to see if the share is already authenticated. In this example, the user has already authenticated to the smb_user1 share:

To clear the cache:

```
net use * /DELETE
You have these remote connections:
                \\FREENAS\smb_user1
Continuing will cancel the connections.
Do you want to continue this operation? <Y/N> [N]: y
```

An additional warning is shown if the share is currently open in Explorer:

```
There are open files and/or incomplete directory searches pending on the connection
to \\FREENAS|smb_user1.
Is it OK to continue disconnecting and force them closed? <Y/N> [N]: y
The command completed successfully.
```

The next time a share is accessed with Explorer, a prompt to authenticate will occur.

10.4.3 Configuring Shadow Copies

Shadow Copies (https://en.wikipedia.org/wiki/Shadow_copy), also known as the Volume Shadow Copy Service (VSS) or Previous Versions, is a Microsoft service for creating volume snapshots. Shadow copies can be used to restore previous versions of files from within Windows Explorer. Shadow Copy support is built into Vista and Windows 7. Windows XP or 2000 users need to install the Shadow Copy client (http://www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=16220).

When a periodic snapshot task is created on a ZFS volume that is configured as a SMB share in FreeNAS[®], it is automatically configured to support shadow copies.

Before using shadow copies with FreeNAS[®], be aware of the following caveats:

- If the Windows system is not fully patched to the latest service pack, Shadow Copies may not work. If no previous versions of files to restore are visible, use Windows Update to make sure that the system is fully up-to-date.
- Shadow copy support only works for ZFS pools or datasets. This means that the SMB share must be configured on a
 volume or dataset, not on a directory.
- Datasets are filesystems and shadow copies cannot traverse filesystems. To see the shadow copies in the child datasets, create separate shares for them.
- Shadow copies will not work with a manual snapshot. Creating a periodic snapshot task for the pool or dataset being shared by SMB or a recursive task for a parent dataset is recommended.
- The periodic snapshot task should be created and at least one snapshot should exist **before** creating the SMB share. If the SMB share was created first, restart the SMB service in *Services* \rightarrow *Control Services*.

- Appropriate permissions must be configured on the volume/dataset being shared by SMB.
- Users cannot delete shadow copies on the Windows system due to the way Samba works. Instead, the administrator can remove snapshots from the FreeNAS[®] administrative GUI. The only way to disable shadow copies completely is to remove the periodic snapshot task and delete all snapshots associated with the SMB share.

To configure shadow copy support, use the instructions in *Configuring Authenticated Access Without a Domain Controller* (page 201) to create the desired number of shares. In this configuration example, a Windows 7 computer has two users: *user1* and *user2*. For this example, two authenticated shares are created so that each user account has their own share. The first share is named *user1* and the second share is named *user2*. Then:

- 1. Use Storage → Periodic Snapshot Tasks → Add Periodic Snapshot to create at least one periodic snapshot task. There are two options for snapshot tasks. One is to create a snapshot task for each user's dataset. In this example the datasets are /mnt/volume1/user1 and /mnt/volume1/user2. Another option is to create one periodic snapshot task for the entire volume; file:/mnt/volume1 in this case. Before continuing to the next step, confirm that at least one snapshot tasks, keep in mind how often the users need to access modified files and during which days and time of day they are likely to make changes.
- 2. Go to Sharing → Windows (SMB) Shares. Highlight a share and click Edit, then Advanced Mode. Click the Periodic Snapshot Task drop-down menu and select the periodic snapshot task to use for that share. Repeat for each share being configured as a shadow copy. For this example, the share named /mnt/volume1/user1 is configured to use a periodic snapshot task that was configured to take snapshots of the /mnt/volume1/user1 dataset and the share named /mnt/volume1/user2 is configured to use a periodic snapshot task that was configured to use a periodic snapshot task that was configured to use a periodic snapshot task that was configured to take snapshots of the /mnt/volume1/user2 is configured to use a periodic snapshot task that was configured to take snapshots of the /mnt/volume1/user2 dataset.
- 3. Verify that the SMB service is set to ON in Services \rightarrow Control Services.

Figure 10.17 provides an example of using shadow copies while logged in as *user1* on the Windows system. In this example, the user right-clicked *modified file* and selected *Restore previous versions* from the menu. This particular file has three versions: the current version, plus two previous versions stored on the FreeNAS[®] system. The user can choose to open one of the previous versions, copy a previous version to the current folder, or restore one of the previous versions, overwriting the existing file on the Windows system.

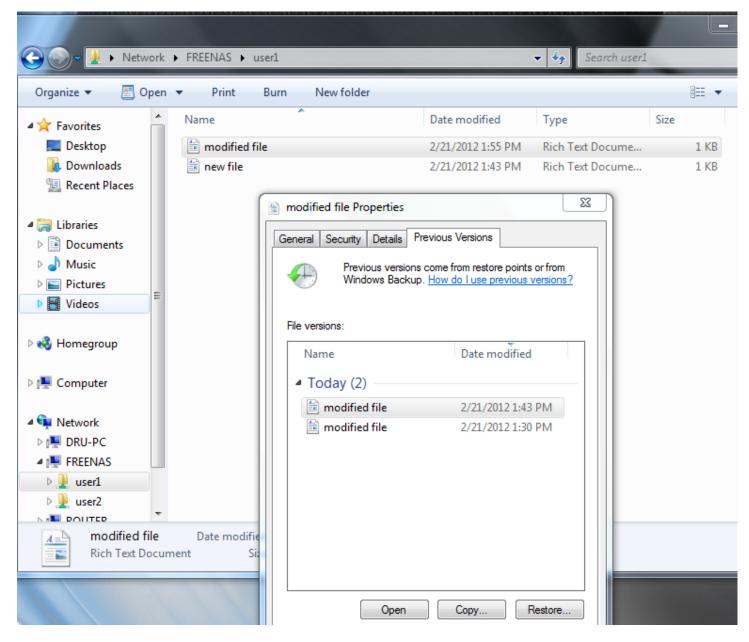


Fig. 10.17: Viewing Previous Versions within Explorer

10.5 Block (iSCSI)

iSCSI is a protocol standard for the consolidation of storage data. iSCSI allows FreeNAS[®] to act like a storage area network (SAN) over an existing Ethernet network. Specifically, it exports disk devices over an Ethernet network that iSCSI clients (called initiators) can attach to and mount. Traditional SANs operate over fibre channel networks which require a fibre channel infrastructure such as fibre channel HBAs, fibre channel switches, and discrete cabling. iSCSI can be used over an existing Ethernet network, although dedicated networks can be built for iSCSI traffic in an effort to boost performance. iSCSI also provides an advantage in an environment that uses Windows shell programs; these programs tend to filter "Network Location" but iSCSI mounts are not filtered.

Before configuring the iSCSI service, be familiar with this iSCSI terminology:

CHAP: an authentication method which uses a shared secret and three-way authentication to determine if a system is authorized to access the storage device and to periodically confirm that the session has not been hijacked by another system.

In iSCSI, the initiator (client) performs the CHAP authentication.

Mutual CHAP: a superset of CHAP in that both ends of the communication authenticate to each other.

Initiator: a client which has authorized access to the storage data on the FreeNAS[®] system. The client requires initiator software to initiate the connection to the iSCSI share.

Target: a storage resource on the FreeNAS[®] system. Every target has a unique name known as an iSCSI Qualified Name (IQN).

Internet Storage Name Service (iSNS): protocol for the automated discovery of iSCSI devices on a TCP/IP network.

Extent: the storage unit to be shared. It can either be a file or a device.

Portal: indicates which IP addresses and ports to listen on for connection requests.

LUN: *Logical Unit Number* representing a logical SCSI device. An initiator negotiates with a target to establish connectivity to a LUN. The result is an iSCSI connection that emulates a connection to a SCSI hard disk. Initiators treat iSCSI LUNs as if they were a raw SCSI or SATA hard drive. Rather than mounting remote directories, initiators format and directly manage filesystems on iSCSI LUNs. When configuring multiple iSCSI LUNs, create a new target for each LUN. Since iSCSI multiplexes a target with multiple LUNs over the same TCP connection, there can be TCP contention when more than one target accesses the same LUN. FreeNAS[®] supports up to 1024 LUNs.

In FreeNAS[®], iSCSI is built into the kernel. This version of iSCSI supports Microsoft Offloaded Data Transfer (ODX) (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11)), meaning that file copies happen locally, rather than over the network. It also supports the *VAAI* (page 324) (vStorage APIs for Array Integration) primitives for efficient operation of storage tasks directly on the NAS. To take advantage of the VAAI primitives, create a zvol using the instructions in *Create zvol* (page 135) and use it to create a device extent, as described in *Extents* (page 213).

To configure iSCSI:

- 1. Review the target global configuration parameters.
- 2. Create at least one portal.
- 3. Determine which hosts are allowed to connect using iSCSI and create an initiator.
- 4. Decide if authentication will be used, and if so, whether it will be CHAP or mutual CHAP. If using authentication, create an authorized access.
- 5. Create a target.
- 6. Create either a device or a file extent to be used as storage.
- 7. Associate a target with an extent.
- 8. Start the iSCSI service in Services \rightarrow Control Services.

The rest of this section describes these steps in more detail.

10.5.1 Target Global Configuration

Sharing \rightarrow Block (iSCSI) \rightarrow Target Global Configuration, shown in Figure 10.18, contains settings that apply to all iSCSI shares. Table 10.6 summarizes the settings that are configured in the Target Global Configuration screen.

Some built-in values affect iSNS usage. Fetching of allowed initiators from iSNS is not implemented, so target ACLs must be configured manually. To make iSNS registration useful, iSCSI targets should have explicitly configured port IP addresses. This avoids initiators attempting to discover unconfigured target portal addresses like 0.0.0.0.

The iSNS registration period is *900* seconds. Registered Network Entities not updated during this period are unregistered. The timeout for iSNS requests is *5* seconds.

Sharing			
Apple (AFP) UNIX (NFS) WebD/	AV Windows (SMB) Block (iSCSI)		
Target Global Configuration Porta	ls Initiators Authorized Access	Targets Extents	Associated Targets
Base Name:	iqn.2005-10.org.freenas.ctl	(i)	
ISNS Servers:			ì
Pool Available Space Threshol	ld (%):	(i)	
Save			



Table 10.6:	Target Global	Configuration	Settings
-------------	---------------	---------------	----------

Setting	Value	Description
Base Name	string	See the "Constructing iSCSI names using the iqn. format" section of RFC 3721 (https://tools.ietf.org/html/rfc3721.html) if unfamiliar with this for- mat.
ISNS Servers	string	Enter the hostnames or IP addresses of ISNS servers to be registered iSCSI targets and portals of the system.
Pool Available Space Threshold	integer	Enter the percentage of free space to remain in the pool. When this per- centage is reached, the system issues an alert, but only if zvols are used. See VAAI (page 324) Threshold Warning for more information.

10.5.2 Portals

A portal specifies the IP address and port number to be used for iSCSI connections. Sharing \rightarrow Block (iSCSI) \rightarrow Portals \rightarrow Add Portal brings up the screen shown in Figure 10.19.

Table 10.19 summarizes the settings that can be configured when adding a portal. To assign additional IP addresses to the portal, click the link *Add extra Portal IP*.

Sharing	
Apple (AFP) UNIX (NFS) WebDAV	Windows (SMB) Block (iSCSI)
Target Global Configuration Portals	Initiators Authorized Access Targets Extents Associated Targets
Add Portal	
Portal Group ID Listen	Comment Discovery Auth Method Discovery Auth Group
No entry has been found	Add Portal 🔀
	Add Portal 🛛 🕅
	Comment:
	Discovery Auth Method: None
	Discovery Auth Group: None
	Portal IP
	IP Address: 0.0.0.0 -
	Port: 3260
	Add extra Portal IP
	OK Cancel

Fig. 10.19: Adding an iSCSI Portal

Setting	Value	Description
Comment	string	Optional description. Portals are automatically assigned a numeric group ID.
Discovery Auth Method	drop-down menu	<i>iSCSI</i> (page 230) supports multiple authentication methods that are used by the target to discover valid devices. <i>None</i> allows anonymous discovery while <i>CHAP</i> and <i>Mutual CHAP</i> both require authentication.
Discovery Auth Group	drop-down menu	Select a user created in <i>Authorized Access</i> if the <i>Discovery Auth Method</i> is set to <i>CHAP</i> or <i>Mutual CHAP</i> .
IP address	drop-down menu	Select the IPv4 or IPv6 address associated with an interface or the wild- card address of 0.0.0.0 (any interface).
Port	integer	TCP port used to access the iSCSI target. Default is 3260.

FreeNAS[®] systems with multiple IP addresses or interfaces can use a portal to provide services on different interfaces or subnets. This can be used to configure multi-path I/O (MPIO). MPIO is more efficient than a link aggregation.

If the FreeNAS[®] system has multiple configured interfaces, portals can also be used to provide network access control. For

example, consider a system with four interfaces configured with these addresses:

192.168.1.1/24

192.168.2.1/24

192.168.3.1/24

192.168.4.1/24

A portal containing the first two IP addresses (group ID 1) and a portal containing the remaining two IP addresses (group ID 2) could be created. Then, a target named A with a Portal Group ID of 1 and a second target named B with a Portal Group ID of 2 could be created. In this scenario, the iSCSI service would listen on all four interfaces, but connections to target A would be limited to the first two networks and connections to target B would be limited to the last two networks.

Another scenario would be to create a portal which includes every IP address **except** for the one used by a management interface. This would prevent iSCSI connections to the management interface.

10.5.3 Initiators

The next step is to configure authorized initiators, or the systems which are allowed to connect to the iSCSI targets on the FreeNAS[®] system. To configure which systems can connect, use *Sharing* \rightarrow *Block (iSCSI)* \rightarrow *Initiators* \rightarrow *Add Initiator*, shown in Figure 10.20.

Add Initiator		38
Initiators	ALL	Ì
Authorized network	ALL	i
Comment	<i>(i)</i>	
OK Cancel		

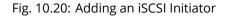


Table 10.8 summarizes the settings that can be configured when adding an initiator.

Table 10.8: Initiator	Configuration Settings
-----------------------	------------------------

Setting	Value	Description
Initiators	string	Use ALL keyword or a list of initiator hostnames separated by spaces.
Authorized network	string	Network addresses that can use this initiator. Use ALL or list network addresses with a CIDR (https://en.wikipedia.org/wiki/Classless_Inter- Domain_Routing) mask. Separate multiple addresses with a space: 192. 168.2.0/24 192.168.2.1/12.
Comment	string	Notes or a description of the initiator.

In the example shown in Figure 10.21, two groups are created. Group 1 allows connections from any initiator on any network. Group 2 allows connections from any initiator on the *10.10.1.0/24* network. Click an initiator's entry to display its *Edit* and *Delete* buttons.

Note: Attempting to delete an initiator causes a warning that indicates if any targets or target/extent mappings depend

upon the initiator. Confirming the delete causes these to be deleted also.

Apple (AFP) UNIX (NFS)	WebDAV	Windows (SN	1B) Block (iSCSI)			
Target Global Configuration	n Portals	Initiators	Authorized Access	Targets	Extents	Associated Targets
Add Initiator	Initiatoro		therized potwork	Common		
Group ID	Initiators		uthorized network	Comment	t	
	Initiators ALL		uthorized network LL	Comment	:	

Fig. 10.21: Sample iSCSI Initiator Configuration

10.5.4 Authorized Accesses

When using CHAP or mutual CHAP to provide authentication, creating an authorized access in Sharing \rightarrow Block (iSCSI) \rightarrow Authorized Accesses \rightarrow Add Authorized Access is recommended. This screen is shown in Figure 10.22.

Note: This screen sets login authentication. This is different from discovery authentication which is set in *Target Global Configuration* (page 206).

A	dd Authorized Access		88
	Group ID	1	
	User		ì
	Secret		ì
	Secret (Confirm)		ì
	Peer User		ì
	Peer Secret		ì
	Peer Secret (Confirm)		ì
	OK Cancel		

Fig. 10.22: Adding an iSCSI Authorized Access

Table 10.9 summarizes the settings that can be configured when adding an authorized access:

cation pro-
nentication
vith the
initiator
be be-
need to be
n the Se-

Note: CHAP does not work with GlobalSAN initiators on macOS.

As authorized accesses are added, they will be listed under *View Authorized Accesses*. In the example shown in Figure 10.23, three users (*test1*, *test2*, and *test3*) and two groups (*1* and *2*) are created, with group 1 consisting of one CHAP user and group 2 consisting of one mutual CHAP user and one CHAP user. Click an authorized access entry to display its *Edit* and *Delete* buttons.

Apple (AFP)	UNIX (NFS)	WebDAV	Windows (SN	1B) Block (iSCSI)			
Target Global	Configuration	Portals	Initiators	Authorized Access	Targets	Extents	Associated Targets
Add Author	ized Access						
ria a ria citor							
			llsor		Deerliser		
Group ID			User tost1		Peer User		
Group ID 1			testl				
Group ID					Peer User test2		

Fig. 10.23: Viewing Authorized Accesses

10.5.5 Targets

Next, create a Target using *Sharing* \rightarrow *Block* (*iSCSI*) \rightarrow *Targets* \rightarrow *Add Target*, as shown in Figure 10.24. A target combines a portal ID, allowed initiator ID, and an authentication method. Table 10.10 summarizes the settings that can be configured when creating a Target.

Note: An iSCSI target creates a block device that may be accessible to multiple initiators. A clustered filesystem is required on the block device, such as VMFS used by VMware ESX/ESXi, in order for multiple initiators to mount the block device read/write. If a traditional filesystem such as EXT, XFS, FAT, NTFS, UFS, or ZFS is placed on the block device, care must be taken that only one initiator at a time has read/write access or the result will be filesystem corruption. If multiple clients need access to the same data on a non-clustered filesystem, use SMB or NFS instead of iSCSI, or create multiple iSCSI targets (one per client).

Add Target	88
Target Name:	Base Name will be appended automatically when starting without 'iqn.', 'eui.' or 'naa.'.
Target Alias:	(1)
iSCSI Group	
Portal Group ID:	
Initiator Group ID:	
Auth Method:	None 🔽 (i)
Authentication Group number:	None
Add extra iSCSI Group	
OK	

Fig. 10.24: Adding an iSCSI Target

		Table 10.10. Target Settings
Setting	Value	Description
Target Name	string	Required value. Base name will be appended automatically if it does not
		start with <i>iqn</i> .
Target Alias	string	Enter an optional user-friendly name.
Portal Group ID	drop-down	Leave empty or select number of existing portal to use.
	menu	
Initiator Group ID	drop-down	Select which existing initiator group has access to the target.
	menu	
Auth Method	drop-down	Choices are None, Auto, CHAP, or Mutual CHAP.
	menu	
Authentication Group num-	drop-down	Select <i>None</i> or an integer representing number of existing authorized ac-
ber	menu	Cess.

Table 10.10: Target Settings

10.5.6 Extents

iSCSI targets provide virtual access to resources on the FreeNAS[®] system. *Extents* are used to define resources to share with clients. There are two types of extents: *device* and *file*.

Device extents provide virtual storage access to zvols, zvol snapshots, or physical devices like a disk, an SSD, a hardware RAID volume, or a HAST device (https://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/disks-hast.html).

File extents provide virtual storage access to an individual file.

Tip: For typical use as storage for virtual machines where the virtualization software is the iSCSI initiator, device extents with zvols provide the best performance and most features. For other applications, device extents sharing a raw device can be appropriate. File extents do not have the performance or features of device extents, but do allow creating

multiple extents on a single filesystem.

Virtualized zvols support all the FreeNAS[®] VAAI (page 324) primitives and are recommended for use with virtualization software as the iSCSI initiator.

The ATS, WRITE SAME, XCOPY and STUN, primitives are supported by both file and device extents. The UNMAP primitive is supported by zvols and raw SSDs. The threshold warnings primitive is fully supported by zvols and partially supported by file extents.

Virtualizing a raw device like a single disk or hardware RAID volume limits performance to the abilities of the device. Because this bypasses ZFS, such devices do not benefit from ZFS caching or provide features like block checksums or snapshots.

Virtualizing a zvol adds the benefits of ZFS, such as read and write cache. Even if the client formats a device extent with a different filesystem, the data still resides on a ZFS volume and benefits from ZFS features like block checksums and snapshots.

Warning: For performance reasons and to avoid excessive fragmentation, keep the used space of the pool below 50% when using iSCSI. The capacity of an existing extent can be increased as shown in *Growing LUNs* (page 217).

To add an extent, go to Sharing \rightarrow Block (iSCSI) \rightarrow Extents \rightarrow Add Extent. In the example shown in Figure 10.25, the device extent is using the export zvol that was previously created from the /mnt/volume1 volume.

Table 10.11 summarizes the settings that can be configured when creating an extent. Note that **file extent creation fails when the name of the file to be created to the volume/dataset name.** is not appended.

A	dd Extent		88
	Extent Name:		String identifier of the extent.
	Extent Type:	Device	
	Device:	adal (10.0 GiB) 🔻	
	Serial:	0800274e099600	Ð
	Logical Block Size:	512 - 1	
	Disable Physical Block Size Reporting:	i)	
	Comment:		Ð
	Enable TPC:	i	
	Xen initiator compat mode:		
	LUN RPM:	ssd – i	
	Read-only:		
	OK Cancel		

Fig. 10.25: Adding an iSCSI Extent

Setting	Value	Description
Extent Name	string	Enter the extent name. If the <i>Extent size</i> is not 0, it cannot be an existing
		file within the volume/dataset.
Extent Type	drop-down	Select from <i>File</i> or <i>Device</i> .
	menu	
Device	drop-down	Only appears if <i>Device</i> is selected. Select the unformatted disk, controller,
	menu	zvol, zvol snapshot, or HAST device.
Serial	string	Unique LUN ID. The default is generated from the system MAC address.
Path to the extent	browse	Only appears if <i>File</i> is selected. Browse to an existing file and use 0 as the
	button	<i>Extent size</i> , or browse to the volume or dataset, click <i>Close</i> , append the <i>Ex</i> -
		<i>tent Name</i> to the path, and specify a value in <i>Extent size</i> . Extents cannot be
		created inside the jail root directory.
Extent size	integer	Only appears if <i>File</i> is selected. If the size is specified as <i>0</i> , the file must
		already exist and the actual file size will be used. Otherwise, specify the
		size of the file to create.
Logical Block Size	drop-down	Only override the default if the initiator requires a different block size.
	menu	

Table 10.11:	Extent Configuration Settings
	Externe configuration settings

Continued on next page

Setting	Value	Description
Disable Physical Block Size	checkbox	Set if the initiator does not support physical block size values over
Reporting		4K (MS SQL). Setting can also prevent constant block size warnings
		(https://www.virten.net/2016/12/the-physical-block-size-reported-by-the-
		device-is-not-supported/) when using this share with ESXi.
Available Space Threshold	string	Only appears if <i>File</i> or a zvol is selected. When the specified percentage
		of free space is reached, the system issues an alert. See VAAI (page 324)
		Threshold Warning for more information.
Comment	string	Enter an optional comment.
Enable TPC	checkbox	If enabled, an initiator can bypass normal access control and access any
		scannable target. This allows xcopy operations otherwise blocked by ac-
		cess control.
Xen initiator compat mode	checkbox	Set this option when using Xen as the iSCSI initiator.
LUN RPM	drop-down	Do NOT change this setting when using Windows as the initiator. Only
	menu	needs to be changed in large environments where the number of systems
		using a specific RPM is needed for accurate reporting statistics.
Read-only	checkbox	Set to prevent the initiator from initializing this LUN .

T 1 1 1 1 1 1		<i>c</i>		
Table 10.11 -	 continued 	trom	previous	page
10010 10.11	continueu		previous	Pape

10.5.7 Target/Extents

The last step is associating an extent to a target within *Sharing* \rightarrow *Block* (*iSCSI*) \rightarrow *Associated Targets* \rightarrow *Add Target/Extent*. This screen is shown in Figure 10.26. Use the drop-down menus to select the existing target and extent. Click *OK* to add an entry for the LUN.

Add Targe	t / Extent	8
Target:		i
LUN ID:	0 💌	
Extent:		
ок	Cancel	

Fig. 10.26: Associating a Target With an Extent

Table 10.12 summarizes the settings that can be configured when associating targets and extents.

Setting	Value	Description
Target	drop-down menu	Select an existing target.
LUN ID	integer	Type a value between 0 and 1023. Note that some initiators expect a value below 256. Enter 0 to statically assign the next available ID.
Extent	drop-down menu	Select an existing extent.

Table 10.12: Target/Extents Configuration Settings

Always associating extents to targets in a one-to-one manner is recommended, even though the GUI will allow multiple extents to be associated with the same target.

Note: Each LUN entry has *Edit* and *Delete* buttons for modifying the settings or deleting the LUN entirely. A verification popup appears when the *Delete* button is clicked. If an initiator has an active connection to the LUN, it is indicated in red text. Clearing initiator connections to a LUN before deleting it is recommended.

After iSCSI has been configured, remember to start it in *Services* \rightarrow *Control Services*. Click the red *OFF* button next to iSCSI. After a second or so, it will change to a blue *ON*, indicating that the service has started.

10.5.8 Connecting to iSCSI

To access the iSCSI target, clients must use iSCSI initiator software.

An iSCSI Initiator client is pre-installed with Windows 7. A detailed how-to for this client can be found here (http://techgenix.com/Connecting-Windows-7-iSCSI-SAN/). A client for Windows 2000, XP, and 2003 can be found here (http://www.microsoft.com/en-us/download/details.aspx?id=18986). This how-to (https://www.pluralsight.com/blog/software-development/freenas-8-iscsi-target-windows-7) shows how to create an iSCSI target for a Windows 7 system.

Mac OS X does not include an initiator. globalSAN (http://www.studionetworksolutions.com/globalsan-iscsi-initiator/) is a commercial, easy-to-use Mac initiator.

BSD systems provide command line initiators: iscontrol(8) (https://www.freebsd.org/cgi/man.cgi?query=iscontrol) comes with FreeBSD versions 9.x and lower, iscsictl(8) (https://www.freebsd.org/cgi/man.cgi?query=iscsictl) comes with FreeBSD versions 10.0 and higher, iscsi-initiator(8) (http://netbsd.gw.com/cgi-bin/man-cgi?iscsi-initiator++NetBSD-current) comes with NetBSD, and iscsid(8) (http://man.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man8/iscsid.8?query=iscsid) comes with OpenBSD.

Some Linux distros provide the command line utility iscsiadm from Open-iSCSI (http://www.open-iscsi.com/). Use a web search to see if a package exists for the distribution should the command not exist on the Linux system.

If a LUN is added while iscsiadm is already connected, it will not see the new LUN until rescanned with iscsiadm -m node -R. Alternately, use iscsiadm -m discovery -t st -p portal_IP to find the new LUN and iscsiadm -m node -T LUN_Name -1 to log into the LUN.

Instructions for connecting from a VMware ESXi Server can be found at How to configure FreeNAS 8 for iSCSI and connect to ESX(i) (https://www.vladan.fr/how-to-configure-freenas-8-for-iscsi-and-connect-to-esxi/). Note that the requirements for booting vSphere 4.x off iSCSI differ between ESX and ESXi. ESX requires a hardware iSCSI adapter while ESXi requires specific iSCSI boot firmware support. The magic is on the booting host side, meaning that there is no difference to the FreeNAS[®] configuration. See the iSCSI SAN Configuration Guide (https://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf) for details.

The VMware firewall only allows iSCSI connections on port *3260* by default. If a different port has been selected, outgoing connections to that port must be manually added to the firewall before those connections will work.

If the target can be seen but does not connect, check the Discovery Auth settings in Target Global Configuration.

If the LUN is not discovered by ESXi, make sure that promiscuous mode is set to *Accept* in the vSwitch.

10.5.9 Growing LUNs

The method used to grow the size of an existing iSCSI LUN depends on whether the LUN is backed by a file extent or a zvol. Both methods are described in this section.

Enlarging a LUN with one of the methods below gives it more unallocated space, but does not automatically resize filesystems or other data on the LUN. This is the same as binary-copying a smaller disk onto a larger one. More space is available on the new disk, but the partitions and filesystems on it must be expanded to use this new space. Resizing virtual disk images is usually done from virtual machine management software. Application software to resize filesystems is dependent on the type of filesystem and client, but is often run from within the virtual machine. For instance, consider a Windows VM with the last partition on the disk holding an NTFS filesystem. The LUN is expanded and the partition table edited to add the new space to the last partition. The Windows disk manager must still be used to resize the NTFS filesystem on that last partition to use the new space.

10.5.9.1 Zvol Based LUN

To grow a zvol based LUN, go to *Storage* \rightarrow *Volumes* \rightarrow *View Volumes*, highlight the zvol to be grown, and click *Edit zvol*. In the example shown in Figure 10.27, the current size of the zvol named *zvol1* is 4 GiB.

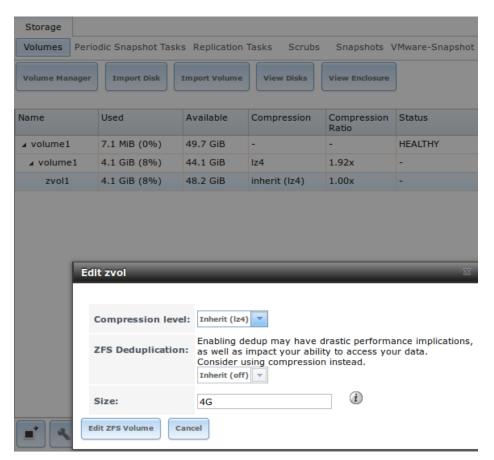


Fig. 10.27: Editing an Existing Zvol

Enter the new size for the zvol in the *Size* field and click *Edit ZFS Volume*. This menu closes and the new size for the zvol is immediately shown in the *Used* column of the *View Volumes* screen.

Note: The GUI does not allow reducing (shrinking) the size of the zvol, as doing so could result in loss of data. It also does not allow increasing the size of the zvol past 80% of the volume size.

10.5.9.2 File Extent Based LUN

To grow a file extent based LUN, go to Services \rightarrow iSCSI \rightarrow File Extents \rightarrow View File Extents to determine the path of the file extent to grow. Open Shell to grow the extent. This example grows /mnt/volume1/data by 2 G:

truncate -s +2g /mnt/volume1/data

Go back to Services \rightarrow iSCSI \rightarrow File Extents \rightarrow View File Extents and click the Edit button for the file extent. Set the size to 0 as this causes the iSCSI target to use the new size of the file.

CHAPTER ELEVEN

SERVICES

Services that ship with FreeNAS[®] are configured, started, or stopped in *Services*. FreeNAS[®] includes these built-in services:

- AFP (page 221)
- Domain Controller (page 222)
- Dynamic DNS (page 224)
- FTP (page 225)
- *iSCSI* (page 230)
- *LLDP* (page 230)
- Netdata (page 231)
- NFS (page 232)
- *Rsync* (page 234)
- 53 (page 236)
- *S.M.A.R.T.* (page 237)
- SMB (page 238)
- SNMP (page 243)
- SSH (page 245)
- *TFTP* (page 247)
- UPS (page 248)
- WebDAV (page 252)

This section demonstrates starting a FreeNAS[®] service and the available configuration options for each FreeNAS[®] service.

11.1 Control Services

 $Services \rightarrow Control Services$, shown in Figure 11.1, lists all services. It also shows where to start, stop, or configure the available services. The S.M.A.R.T. service is enabled by default, but only runs if the storage devices support S.M.A.R.T. data (https://en.wikipedia.org/wiki/S.M.A.R.T.) Other services default to off until started.

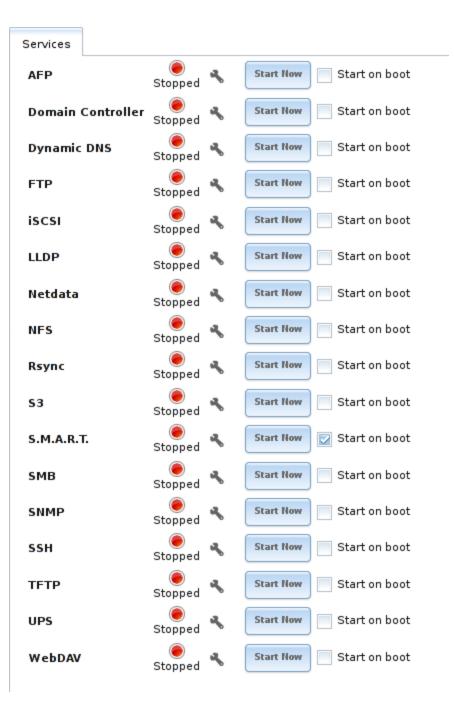


Fig. 11.1: Control Services

Stopped services show a red stop symbol and a *Start Now* button. Running services show a green light with a *Stop Now* button.

Tip: Using a proxy server can prevent the list of services from being displayed. If a proxy server is used, do not configure it to proxy local network connections or websocket connections. VPN software can also cause problems. If the list of services is displayed when connecting on the local network but not when connecting through the VPN, check the VPN software configuration.

Services are configured by clicking the wrench icon or the name of the service in the Services section of the tree menu.

If a service does not start, go to System \rightarrow Advanced and enable Show console messages in the footer. Console messages appear

at the bottom of the browser. Clicking the console message area makes it into a pop-up window, allowing scrolling through or copying the messages. Watch these messages for errors when stopping or starting the problematic service.

To read the system logs for more information about a service failure, open *Shell* (page 289) and type more /var/log/ messages.

11.2 AFP

The settings that are configured when creating AFP Shares in *Sharing* \rightarrow *Apple (AFP) Shares* \rightarrow *Add Apple (AFP) Share* are specific to each configured AFP Share. In contrast, global settings which apply to all AFP shares are configured in Services \rightarrow *AFP*.

Figure 11.2 shows the available global AFP configuration options which are described in Table 11.1.

Settings		_	38
Guest Access:			
Guest account:	nobody 👻	٢	
Max. Connections:	50	(i)	
Database Path:	(i)	Browse	
Global auxiliary parameters:			ì
Map ACLs:	Rights 💌 (i		
Chmod Request:	Preserve 💌 (i)		
Bind IP Addresses:	10.231.1.3		
OK Cancel			

Fig. 11.2: Global AFP Configuration

Setting	Value	Description
Guest Access	checkbox	Set to disable the password prompt that appears before clients access AFP shares.
Guest account	drop-down menu	Select an account to use for guest access. The account must have permissions to the volume or dataset being shared.
Max Connec- tions	integer	Maximum number of simultaneous connections.
Database Path	browse button	Sets the database information to be stored in the path. Default is the root of the volume. The path must be writable even if the volume is read only.
Global auxiliary parameters	string	Add any additional afp.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=afp.conf) parameters not covered elsewhere in this screen.
Map ACLs	drop-down menu	Choose mapping of effective permissions for authenticated users. Choices are: <i>Rights</i> (default, Unix-style permissions), <i>Mode</i> (ACLs), or <i>None</i>
Chmod Request	drop-down menu	Sets how Access Control Lists are handled. <i>Ignore</i> : ignores requests and gives the parent directory ACL inheritance full control over new items. <i>Preserve</i> : preserves ZFS Access Control Entries for named users and groups or the POSIX ACL group mask. <i>Simple</i> : is set to chmod() as requested without any extra steps.
Bind IP Ad- dresses	selection	Specify the IP addresses to listen for FTP connections. Highlight the de- sired IP addresses in the <i>Available</i> list and use the >> button to add to the <i>Selected</i> list.

Table 11.1:	Global AF	P Configuration	Options
	01000171	Configuration	options

11.2.1 Troubleshooting AFP

Check for error messages in /var/log/afp.log.

Determine which users are connected to an AFP share by typing afpusers.

If *Something wrong with the volume's CNID DB* is shown, run this command from *Shell* (page 289), replacing the path to the problematic AFP share:

dbd -rf /path/to/share

This command can take some time, depending upon the size of the pool or dataset being shared. The CNID database is wiped and rebuilt from the CNIDs stored in the AppleDouble files.

11.3 Domain Controller

FreeNAS[®] can be configured to act either as the domain controller for a network or to join an existing *Active Directory* (page 169) network as a domain controller.

Note: This section demonstrates how to configure the FreeNAS[®] system to act as a domain controller. If the goal is to integrate with an existing *Active Directory* (page 169) network to access its authentication and authorization services, configure *Active Directory* (page 169) instead.

Note that configuring a domain controller is a complex process that requires a good understanding of how *Active Directory* (page 169) works. While *Services* \rightarrow *Domain Controller* makes it easy to enter the needed settings into the administrative graphical interface, it is important to understand what those settings should be. Before beginning configuration, read through the Samba AD DC HOWTO (https://wiki.samba.org/index.php/Samba_AD_DC_HOWTO). After FreeNAS[®] is configured, use the RSAT utility from a Windows system to manage the domain controller. The Samba AD DC HOWTO includes instructions for installing and configuring RSAT.

Figure 11.3 shows the configuration screen for creating a domain controller and Table 11.2 summarizes the available options.

Se	ettings	
	Realm:	[(i)
	Domain:	i
	Server Role:	active directory domain controller 💌 🚺
	DNS Forwarder:	i
	Domain Forest Level:	2003 💌 (i)
	Administrator Password:	i
	Confirm Administrator Password:	
	Kerberos Realm:	
	OK Cancel Delete	

Fig. 11.3: Domain Controller Settings

Setting	Value	Description
Realm	string	Enter a capitalized DNS realm name.
Domain	string	Enter a capitalized domain name.
Server Role	drop-down menu	At this time, the only supported role is as the domain controller for a new
		domain.
DNS Forwarder	string	Enter the IP address of the DNS forwarder. Required for recursive queries
		when SAMBA_INTERNAL is selected.
Domain Forest	drop-down menu	Choices are 2000, 2003, 2008, 2008_R2, 2012, or 2012_R2. Refer to Un-
Level		derstanding Active Directory Domain Services (AD DS) Functional Lev-
		els (https://docs.microsoft.com/en-us/previous-versions/windows/it-
		pro/windows-server-2008-R2-and-2008/cc754918(v=ws.10)).
Administrator	string	Enter the password to be used for the <i>Active Directory</i> (page 169) adminis-
password		trator account.
Kerberos Realm	drop-down menu	Auto-populates with information from the <i>Realm</i> when the settings in this
		screen are saved.

11.3.1 Samba Domain Controller Backup

A samba_backup script is available to back up Samba4 domain controller settings is available. From the *Shell* (page 289), run /usr/local/bin/samba_backup --usage to show the input options.

11.4 Dynamic DNS

Dynamic DNS (DDNS) is useful if the FreeNAS[®] system is connected to an ISP that periodically changes the IP address of the system. With dynamic DNS, the system can automatically associate its current IP address with a domain name, allowing access to the FreeNAS[®] system even if the IP address changes. DDNS requires registration with a DDNS service such as DynDNS (https://dyn.com/dns/).

Figure 11.4 shows the DDNS configuration screen and Table 11.3 summarizes the configuration options. The values for these fields are provided by the DDNS provider. After configuring DDNS, remember to start the DDNS service in Services \rightarrow Control Services.

S	ettings		86
	Provider:		
	CheckIP Server SSL:		
	CheckIP Server:		i
	CheckIP Path:		i
	Use SSL:		
	Domain name:		i
	Username:	admin	
	Password:		
	Confirm Password:		
	Update Period:	300	i
	OK Cancel		

Fig. 11.4: Configuring DDNS

Setting	Value	Description
Provider	drop-down menu	Several providers are supported. If a specific provider is not listed, select <i>Custom Provider</i> and enter the information in the <i>Custom Server</i> and <i>Custom Path</i> fields.
CheckIP Server SSL	string	Set to use HTTPS for the connection to the <i>CheckIP Server</i> .
CheckIP Server	string	Enter the name and port of the server that reports the external IP ad- dress. Example: <i>server.name.org:port</i> .
CheckIP Path	string	Enter the path that is requested by the <i>CheckIP Server</i> to determine the user IP address.
Use SSL	checkbox	Set to use HTTPS for the connection to the server that updates the DNS record.
Domain name	string	Enter a fully qualified domain name. Example: <i>yourname.dyndns.org</i> .
Username	string	Enter the username used to log in to the provider and update the record.
Password	string	Enter the password used to log in to the provider and update the record.
Update period	integer	How often the IP is checked in seconds.

Table 11.3: DDNS Configuration Options

When using he.net, enter the domain name for *Username* and enter the DDNS key generated for that domain's A entry at the he.net (https://he.net) website for *Password*.

11.5 FTP

FreeNAS[®] uses the proftpd (http://www.proftpd.org/) FTP server to provide FTP services. Once the FTP service is configured and started, clients can browse and download data using a web browser or FTP client software. The advantage of FTP is that easy-to-use cross-platform utilities are available to manage uploads to and downloads from the FreeNAS[®] system. The disadvantage of FTP is that it is considered to be an insecure protocol, meaning that it should not be used to transfer sensitive files. If concerned about sensitive data, see *Encrypting FTP* (page 230).

This section provides an overview of the FTP configuration options. It then provides examples for configuring anonymous FTP, specified user access within a chroot environment, encrypting FTP connections, and troubleshooting tips.

Figure 11.5 shows the configuration screen for Services \rightarrow FTP. Some settings are only available in Advanced Mode. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by enabling the Show advanced fields by default setting in System \rightarrow Advanced.

FTP Settings	Doporting Wissed	33
Port:	21	æ
Port:	μ1	·
Clients:	5	<i>(</i>)
Connections:	2	ì
Login Attempts:	1	i
Timeout:	600	i
Allow Root Login:		
Allow Anonymous Login:		
Path:		Browse
Allow Local User Login:		
Display Login:		
	1	
Allow Transfer Resumption:		
	~	•

Fig. 11.5: Configuring FTP

Table 11.4 summarizes the available options when configuring the FTP server.

Setting	Value	Advanced	Description
		Mode	
Port	integer		Set the port the FTP service listens on.
Clients	integer		Set the maximum number of simultaneous clients.
Connections	integer		Set the maximum number of connections per IP address where 0 means unlimited.
Login Attempts	integer		Enter the maximum number of attempts before client is dis- connected. Increase this if users are prone to typos.
Timeout	integer		Enter the maximum client idle time in seconds before client is disconnected.
Allow Root Login	checkbox		Enabling this option is discouraged as increases security risk.
Allow Anonymous	checkbox		Set to enable anonymous FTP logins with access to the direc-
Login			tory specified in <i>Path</i> .
Path	browse but- ton		Set the root directory for anonymous FTP connections.
Allow Local User Lo- gin	checkbox		Required if Anonymous Login is disabled.
Display Login	string		Specify the message displayed to local login users after authen- tication. Not displayed to anonymous login users.
File Permission	checkboxes	\checkmark	Set the default permissions for newly created files.
Directory Permission	checkboxes	\checkmark	Set the default permissions for newly created directories.
Directory remission	спескоолез	v	Set the delidate permissions for newly created directories.

			ntinued from previous page	
Setting	Value	Advanced Mode	Description	
Enable FXP	checkbox	./	Set to enable the File eXchange Protocol. This setting makes	
(https://en.wikipedia.or		nge_Protocol)	the server vulnerable to FTP bounce attacks so it is not recom- mended	
Allow Transfer Re- sumption	checkbox		Set to allow FTP clients to resume interrupted transfers.	
Always Chroot	checkbox		When set, a local user is only allowed access to their home di- rectory unless the user is a member of group <i>wheel</i> .	
Require IDENT Au- thentication	checkbox	\checkmark	Setting this option results in timeouts if identd is not running on the client.	
Perform Reverse DNS Lookups	checkbox		Set to perform reverse DNS lookups on client IPs. Can cause long delays if reverse DNS is not configured.	
Masquerade address	string		Public IP address or hostname. Set if FTP clients cannot con- nect through a NAT device.	
Minimum passive port	integer	√	Used by clients in PASV mode, default of <i>0</i> means any port above 1023.	
Maximum passive port	integer	\checkmark	Used by clients in PASV mode, default of <i>0</i> means any port above 1023.	
Local user upload bandwidth	integer	\checkmark	Defined in KiB/s, default of 0 means unlimited.	
Local user download bandwidth	integer	\checkmark	Defined in KiB/s, default of 0 means unlimited.	
Anonymous user up- load bandwidth	integer	\checkmark	Defined in KiB/s, default of 0 means unlimited.	
Anonymous user download bandwidth	integer	\checkmark	Defined in KiB/s, default of 0 means unlimited.	
Enable TLS	checkbox	\checkmark	Set to enable encrypted connections. Requires a certificate to be created or imported using <i>Certificates</i> (page 89).	
TLS policy	drop-down menu	V	The selected policy defines whether the control channel, data channel, both channels, or neither channel of an FTP ses- sion must occur over SSL/TLS. The policies are described here (http://www.proftpd.org/docs/directives/linked/config_ref_TLSRequi	ired
TLS allow client rene- gotiations	checkbox	\checkmark	Enabling this option is not recommended as it breaks several security measures. For this and the rest of the TLS fields, refer to mod_tls (http://www.proftpd.org/docs/contrib/mod_tls.html) for more details.	
TLS allow dot login	checkbox	\checkmark	If set, the user home directory is checked for a .tlslogin file which contains one or more PEM-encoded certificates. If not found, the user is prompted for password authentication.	
TLS allow per user	checkbox	\checkmark	If set, the user password can be sent unencrypted.	
TLS common name required	checkbox	\checkmark	Set to require the certificate common name to match the FQDN of the host.	
TLS enable diagnos- tics	checkbox	\checkmark	If set when troubleshooting a connection, logs more verbosely.	
TLS export certificate data	checkbox	\checkmark	If set, exports the certificate environment variables.	
TLS no certificate re- quest	checkbox	\checkmark	Try enabling this option if the client cannot connect and it is suspected the client software is not properly handling server certificate requests.	
TLS no empty frag- ments	checkbox	\checkmark	Enabling this is not recommended as it bypasses a security mechanism.	
TLS no session reuse required	checkbox	\checkmark	Enabling this reduces the security of the connection. Only use this if the client does not understand reused SSL sessions.	
	·		Continued on next page	

Table 11.4 – continued from previous page

Setting	Value	Advanced Mode	Description
TLS export standard	checkbox	√ Viole	If enabled, sets several environment variables.
TLS DNS name re-	checkbox	\checkmark	If set, the client DNS name must resolve to its IP address and
quired			the cert must contain the same DNS name.
TLS IP address re- quired	checkbox	\checkmark	If set, the client certificate must contain the IP address that matches the IP address of the client.
Certificate	drop-down menu		The SSL certificate to be used for TLS FTP connections. To create a certificate, use System \rightarrow Certificates.
Auxiliary parameters	string	\checkmark	Add any additional proftpd(8) (https://www.freebsd.org/cgi/man.cgi?query=proftpd) parameters not covered elsewhere in this screen.

Table 11.4 – continued from previous page

This example demonstrates the auxiliary parameters that prevent all users from performing the FTP DELETE command:

<Limit DELE> DenyAll </Limit>

11.5.1 Anonymous FTP

Anonymous FTP may be appropriate for a small network where the FreeNAS[®] system is not accessible from the Internet and everyone in the internal network needs easy access to the stored data. Anonymous FTP does not require a user account for every user. In addition, passwords are not required so it is not necessary to manage changed passwords on the FreeNAS[®] system.

To configure anonymous FTP:

- 1. Give the built-in ftp user account permissions to the volume/dataset to be shared in *Storage* \rightarrow *Volumes* as follows:
 - Owner(user): select the built-in ftp user from the drop-down menu
 - Owner(group): select the built-in ftp group from the drop-down menu
 - Mode: review that the permissions are appropriate for the share

Note: For FTP, the type of client does not matter when it comes to the type of ACL. This means that Unix ACLs are always used, even if Windows clients are accessing FreeNAS[®] via FTP.

- 2. Configure anonymous FTP in Services \rightarrow FTP by setting these attributes:
 - Allow Anonymous Login: enable this option
 - Path: browse to the volume/dataset/directory to be shared
- 3. Start the FTP service in Services \rightarrow Control Services. Click the Start Now button next to FTP. The FTP service takes a second or so to start. The indicator changes to green when the service is running, and the button changes to Stop Now.
- 4. Test the connection from a client using a utility such as Filezilla (https://filezilla-project.org/).

In the example shown in Figure 11.6, the user has entered this information into the Filezilla client:

- IP address of the FreeNAS[®] server: 192.168.1.113
- Username: anonymous
- Password: the email address of the user

<u>F</u> ile <u>E</u> dit	<u>V</u> iew <u>T</u> ransfer <u>S</u> erver <u>B</u> ookmarks <u>H</u> elp	
111 - [🖻 🗂 🗮 í 🛇 🏦 🕸 🗽 🗊 🖓 🤗 🙈	
<u>H</u> ost: 192.		ickconnect 👻
Response: Status: Status: Command: Response:	200 UTF8 set to on Logged in Retrieving directory listing PWD 257 "/" is the current directory Directory listing of "/" successful	• •
Local site:	/usr/home/tmoore/	Remote site: /

Fig. 11.6: Connecting Using Filezilla

The messages within the client indicate the FTP connection is successful. The user can now navigate the contents of the root folder on the remote site. This is the pool or dataset specified in the FTP service configuration. The user can also transfer files between the local site (their system) and the remote site (the FreeNAS[®] system).

11.5.2 FTP in chroot

If users are required to authenticate before accessing the data on the FreeNAS[®] system, either create a user account for each user or import existing user accounts using *Active Directory* (page 169) or *LDAP* (page 174). Then create a ZFS dataset for *each* user. Next, chroot each user so they are limited to the contents of their own home directory. Datasets provide the added benefit of configuring a quota so that the size of a user home directory is limited to the size of the quota.

To configure this scenario:

- 1. Create a ZFS dataset for each user in *Storage* \rightarrow *Volumes*. Click an existing *ZFS volume* \rightarrow *Create ZFS Dataset* and set an appropriate quota for each dataset. Repeat this process to create a dataset for every user that needs access to the FTP service.
- 2. When not using AD or LDAP, create a user account for each user in *Account* \rightarrow *Users* \rightarrow *Add User*. For each user, browse to the dataset created for that user in the *Home Directory* field. Repeat this process to create a user account for every user that needs access to the FTP service, making sure to assign each user their own dataset.
- 3. Set the permissions for each dataset in *Storage* \rightarrow *Volumes*. Click the *Change Permissions* button for a dataset to assign a user account as *Owner* of that dataset and to set the desired permissions for that user. Repeat for each dataset.

Note: For FTP, the type of client does not matter when it comes to the type of ACL. This means Unix ACLs are always used, even if Windows clients will be accessing FreeNAS[®] with FTP.

- 4. Configure FTP in Services \rightarrow FTP with these attributes:
 - *Path*: browse to the parent volume containing the datasets.
 - Make sure the options for Allow Anonymous Login and Allow Root Login are unselected.
 - Select the Allow Local User Login option to enable it.
 - Enable the Always Chroot option.
- 5. Start the FTP service in Services \rightarrow Control Services. Click the Start Now button next to FTP. The FTP service takes a second or so to start. The indicator changes to green to show that the service is running, and the button changes to Stop Now.
- 6. Test the connection from a client using a utility such as Filezilla.

To test this configuration in Filezilla, use the *IP address* of the FreeNAS[®] system, the *Username* of a user that is associated with a dataset, and the *Password* for that user. The messages will indicate the authorization and the FTP connection are successful. The user can now navigate the contents of the root folder on the remote site. This time it is not the entire pool but the dataset created for that user. The user can transfer files between the local site (their system) and the remote site (their dataset on the FreeNAS[®] system).

11.5.3 Encrypting FTP

To configure any FTP scenario to use encrypted connections:

- 1. Import or create a certificate authority using the instructions in *CAs* (page 87). Then, import or create the certificate to use for encrypted connections using the instructions in *Certificates* (page 89).
- 2. In Services \rightarrow FTP, choose the certificate in the Certificate, and set the Enable TLS option.
- 3. Specify secure FTP when accessing the FreeNAS[®] system. For example, in Filezilla enter *ftps://IP_address* (for an implicit connection) or *ftpes://IP_address* (for an explicit connection) as the Host when connecting. The first time a user connects, they will be presented with the certificate of the FreeNAS[®] system. Click *OK* to accept the certificate and negotiate an encrypted connection.
- 4. To force encrypted connections, select on for the TLS Policy.

11.5.4 Troubleshooting FTP

The FTP service will not start if it cannot resolve the system hostname to an IP address with DNS. To see if the FTP service is running, open *Shell* (page 289) and issue the command:

sockstat -4p 21

If there is nothing listening on port 21, the FTP service is not running. To see the error message that occurs when FreeNAS[®] tries to start the FTP service, go to *System* \rightarrow *Advanced*, check *Show console messages in the footer*, and click *Save*. Go to *Services* \rightarrow *Control Services* and switch the FTP service off, then back on. Watch the console messages at the bottom of the browser for errors.

If the error refers to DNS, either create an entry in the local DNS server with the FreeNAS[®] system hostname and IP address, or add an entry for the IP address of the FreeNAS[®] system in the *Network* \rightarrow *Global Configuration Host name data base* field.

11.6 iSCSI

Refer to Block (iSCSI) (page 205) for instructions on configuring iSCSI. To start the iSCSI service, click its entry in Services.

Note: A warning message is shown if the iSCSI service is stopped when initiators are connected. Open the *Shell* (page 289) and type ctladm islist to determine the names of the connected initiators.

11.7 LLDP

The Link Layer Discovery Protocol (LLDP) is used by network devices to advertise their identity, capabilities, and neighbors on an Ethernet network. FreeNAS[®] uses the ladvd (https://github.com/sspans/ladvd) LLDP implementation. If the network contains managed switches, configuring and starting the LLDP service will tell the FreeNAS[®] system to advertise itself on the network.

Figure 11.7 shows the LLDP configuration screen and Table 11.5 summarizes the configuration options for the LLDP service.

L	LDP Settings		
	Interface Description:	1	
	Country Code:		i
	Location:		i
	ОК Cancel		

Fig. 11.7: Configuring LLDP

Setting	Value	Description
Interface De-	checkbox	Set to enable receive mode and to save received peer information in inter-
scription		face descriptions.
Country Code	string	Required for LLDP location support. Enter a two-letter ISO 3166 country
		code.
Location	string	Optional. Specify the physical location of the host.

11.8 Netdata

Netdata is a real-time performance and monitoring system. It displays data as web dashboards.

Start the Netdata service from the *Services* (page 219) screen. Click the wrench icon to display the Netdata settings dialog shown in Figure 11.8.

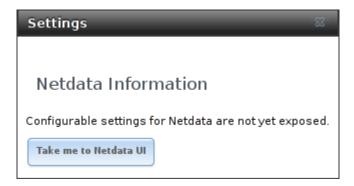


Fig. 11.8: Netdata Settings Dialog

Click the *Take me to the Netdata UI* button to view the web dashboard as shown in Figure 11.9.

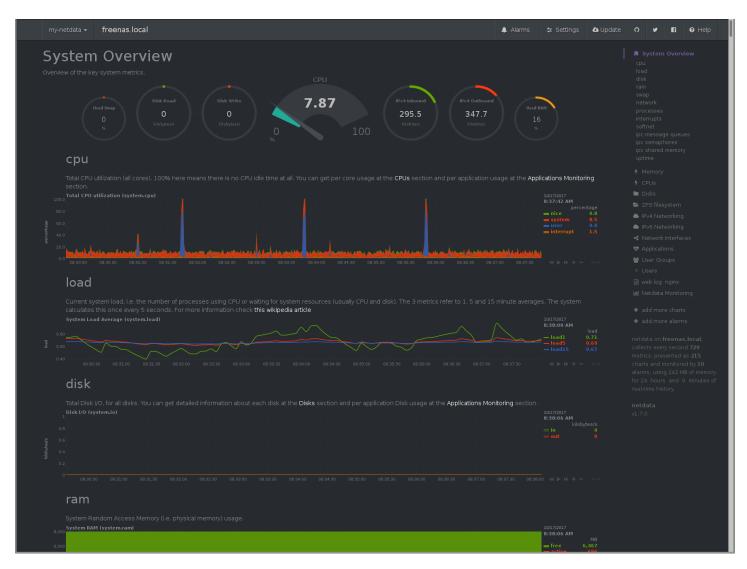


Fig. 11.9: Netdata Web Dashboard

More information on configuring and using Netdata is available at the Netdata website (https://my-netdata.io/).

11.9 NFS

The settings that are configured when creating NFS Shares in *Sharing* \rightarrow *Unix (NFS) Shares* \rightarrow *Add Unix (NFS) Share* are specific to each configured NFS Share. In contrast, global settings which apply to all NFS shares are configured in *Services* \rightarrow *NFS*. Figure 11.10 shows the configuration screen and Table 11.6 summarizes the configuration options for the NFS service.

S	ettings		38
	Number of servers:	4	i
	Serve UDP NFS clients:		
	Bind IP Addresses:	10.0.0.142	
	Allow non-root mount:	(i)	
	Enable NFSv4:		
	NFSv3 ownership model for NFSv4:	(i)	
	Require Kerberos for NFSv4:		
	mountd(8) bind port:		i
	rpc.statd(8) bind port:		i
	rpc.lockd(8) bind port:		i
	Support >16 groups:		
	Log mountd(8) requests:		
	Log rpc.statd(8) and rpc.lockd(8):		
	OK Cancel		

Fig. 11.10: Configuring NFS

Setting	Value	Description
Number of	integer	Specify how many servers to create. Increase if NFS client responses are
servers		slow. To limit CPU context switching, keep this number less than or equal
		to the number of CPUs reported by sysctl -n kern.smp.cpus.
Serve UDP NFS	checkbox	Set if NFS clients need to use UDP.
clients		
Bind IP Ad-	checkboxes	Select the IP addresses to listen on for NFS requests. When unselected,
dresses		NFS listens on all available addresses.

Table 11.6: NFS Configuration Options

Setting	Value	Description
Allow non-root	checkbox	Set only if the NFS client requires it.
mount		
Enable NFSv4	checkbox	Set to switch from NFSv3 to NFSv4. The default is NFSv3.
NFSv3 owner-	checkbox	Grayed out unless Enable NFSv4 is checked and, in turn, grays out Sup-
ship model for		<i>port>16 groups</i> which is incompatible. Set this option if NFSv4 ACL support
NFSv4		is needed without requiring the client and the server to sync users and groups.
Require Ker- beros for NFSv4	checkbox	Set to force NFS shares to fail if the Kerberos ticket is unavailable.
mountd(8) bind	integer	Optional. Specify the port that mountd(8)
port		(https://www.freebsd.org/cgi/man.cgi?query=mountd) binds to.
rpc.statd(8) bind	integer	Optional. Specify the port that rpc.statd(8)
port		(https://www.freebsd.org/cgi/man.cgi?query=rpc.statd) binds to.
rpc.lockd(8) bind	integer	Optional. Specify the port that rpc.lockd(8)
port		(https://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) binds to.
Support>16	checkbox	Set this option if any users are members of more than 16 groups (useful
groups		in AD environments). Note this assumes group membership is configured
		correctly on the NFS server.
Log mountd(8)	checkbox	Enable logging of mountd(8)
requests		(https://www.freebsd.org/cgi/man.cgi?query=mountd) requests by
		syslog.
Log rpc.statd(8)	checkbox	Enable logging of rpc.statd(8)
and rpc.lockd(8)		(https://www.freebsd.org/cgi/man.cgi?query=rpc.statd) and rpc.lockd(8)
		(https://www.freebsd.org/cgi/man.cgi?query=rpc.lockd) requests by syslog.

Table 11.6 – continued from previous page

Note: NFSv4 sets all ownership to *nobody:nobody* if user and group do not match on client and server.

11.10 Rsync

Services \rightarrow Rsync is used to configure an rsync server when using rsync module mode. Refer to Rsync Module Mode (page 107) for a configuration example.

This section describes the configurable options for the rsyncd service and rsync modules.

11.10.1 Configure Rsyncd

Figure 11.11 shows the rsyncd configuration screen which is accessed from Services \rightarrow Rsync \rightarrow Configure Rsyncd.

С	onfigure Rsyncd		36
	TCP Port	873	
	Auxiliary parameters		i
	OK Cancel		

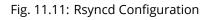


Table 11.7 summarizes the configuration options for the rsync daemon:

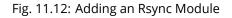
Setting	Value	Description	
TCP Port	integer	Port for rsyncd to listen on. Default is 873.	
Auxiliary param-	string	Enter any additional parameters from rsyncd.conf(5)	
eters		(https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).	

Table 11 7 [.]	Rsyncd Configuration Optic	ons
	nsynca conngaration optiv	5115

11.10.2 Rsync Modules

Figure 11.12 shows the configuration screen that appears after clicking $Services \rightarrow Rsync \rightarrow Rsync Modules \rightarrow Add Rsync Module$. Table 11.8 summarizes the configuration options available when creating a rsync module.

Add Rsync Module		8	A
Module name			
Comment			
Path		Browse	_
Access Mode	Read and Write 👻 (i)		Ш
Maximum connections	0		
User	nobody 👻	(i)	
Group	nobody 👻	(i)	
Hosts allow	i		
Hosts deny			-



Setting	Value	Description
Module name	string	Mandatory. This is required to match the setting on the rsync client.
Comment	string	Optional description.
Path	browse button	Browse to the volume or dataset to hold received data.
Access Mode	drop-down menu	Choices are Read and Write, Read-only, or Write-only.
Maximum con-	integer	0 is unlimited.
nections		
User	drop-down menu	Select the user to control file transfers to and from the module.
Group	drop-down menu	Select the group to control file transfers to and from the module.

Table 11.8: Rsync Module Configuration Options

Setting	Value	Description
Hosts allow	string	See rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf)
		Enter a list of patterns to match with the hostname and IP address of a
		connecting client. Separate patterns with whitespace or comma.
Hosts deny	string	See rsyncd.conf(5) (https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf)
		for allowed formats.
Auxiliary param-	string	Enter any additional parameters from rsyncd.conf(5)
eters		(https://www.freebsd.org/cgi/man.cgi?query=rsyncd.conf).

Table 11.8 – continued from previous page

11.11 S3

S3 is a distributed or clustered filesystem protocol compatible with Amazon S3 cloud storage. The FreeNAS[®] S3 service uses Minio (https://minio.io/) to provide S3 storage hosted on the FreeNAS[®] system itself. Minio also provides features beyond the limits of the basic Amazon S3 specifications.

Figure 11.13 shows the S3 service configuration screen and Table 11.9 summarizes the configuration options. After configuring the S3 service, start it in Services \rightarrow Control Services.

s	Settings 🛛 🕅			
	IP Address:	0.0.0.0	ì	
	Port:	9000	Ì	
	Access key of 5 to 20 characters in length:		ì	
	Secret key of 8 to 40 characters in length:		Ì	
	Confirm S3 Key:			
	Disks:	i)	Browse	
	Certificate:			
	Enable Browser:			
	OK			

Fig. 11.13: Configuring S3

Table 11.9: S3 Configuration Options

Setting	Value	Description
IP Address	drop-down menu	Enter the IP address to run the S3 service. 0.0.0.0 sets the server to listen
		on all addresses.

Setting	Value	Description	
Port	string	Enter the TCP port on which to provide the S3 service. Default is 9000.	
Access Key	string	Enter the S3 user name. This username must contain only alphanumeric	
		characters and be between 5 and 20 characters long.	
Secret Key	string	Enter the password to be used by connecting S3 systems. The key must	
		contain only alphanumeric characters and be at least 8 but no more than	
		40 characters long.	
Confirm S3 Key	string	Re-enter the S3 password to confirm.	
Disks	string	Directory where the S3 filesystem will be mounted. Ownership of this	
		directory and all subdirectories is set to <i>minio:minio</i> . <i>Create a separate</i>	
		<i>dataset</i> (page 132) for Minio to avoid issues with conflicting directory per-	
		missions or ownership.	
Certificate	drop-down menu	The SSL certificate to be used for secure S3 connections. To create a cer-	
		tificate, use System $ ightarrow$ Certificates.	
Enable Browser	checkbox	Set to enable the web user interface for the S3 service.	

Table 11.9 – continued	from	previous	page
------------------------	------	----------	------

11.12 S.M.A.R.T.

S.M.A.R.T., or Self-Monitoring, Analysis, and Reporting Technology (https://en.wikipedia.org/wiki/S.M.A.R.T.), is an industry standard for disk monitoring and testing. Drives can be monitored for status and problems, and several types of self-tests can be run to check the drive health.

Tests run internally on the drive. Most tests can run at the same time as normal disk usage. However, a running test can greatly reduce drive performance, so they should be scheduled at times when the system is not busy or in normal use. It is very important to avoid scheduling disk-intensive tests at the same time. For example, do not schedule S.M.A.R.T. tests to run at the same time, or preferably, even on the same days as *Scrubs* (page 162).

Of particular interest in a NAS environment are the *Short* and *Long* S.M.A.R.T. tests. Details vary between drive manufacturers, but a *Short* test generally does some basic tests of a drive that takes a few minutes. The *Long* test scans the entire disk surface, and can take several hours on larger drives.

FreeNAS[®] uses the smartd(8) (https://www.smartmontools.org/browser/trunk/smartmontools/smartd.8.in) service to monitor S.M.A.R.T. information. A complete configuration consists of:

- 1. Scheduling when S.M.A.R.T. tests are run in *Tasks* \rightarrow *S.M.A.R.T. Tests* \rightarrow *Add S.M.A.R.T. Test*.
- 2. Enabling or disabling S.M.A.R.T. for each disk member of a volume in *Volumes* \rightarrow *View Disks*. This setting is enabled by default for disks that support S.M.A.R.T.
- 3. Checking the configuration of the S.M.A.R.T. service as described in this section.
- 4. Starting the S.M.A.R.T. service with Services \rightarrow Control Services.

Figure 11.14 shows the configuration screen that appears after clicking Services \rightarrow S.M.A.R.T.

S	S.M.A.R.T. Settings		
	Check interval:	<u>β</u> 0	i
	Power mode:	Never - Check the device	
	Difference:	0	i
	Informational:	0	i
	Critical:	0	i
	Email to report:		ì
	OK Cancel		

Fig. 11.14: S.M.A.R.T Configuration Options

Note: smartd wakes up at the configured *Check Interval*. It checks the times configured in *Tasks* \rightarrow *S.M.A.R.T. Tests* to see if a test must begin. Since the smallest time increment for a test is an hour, it does not make sense to set a *Check Interval* value higher than 60 minutes. For example, if the *Check Interval* is set to *120* minutes and the smart test to every hour, the test will only be run every two hours because smartd only activates every two hours.

Table 11.10 summarizes the options in the S.M.A.R.T configuration screen.

	Table	e T1.10. S.M.A.R.I Configuration Options
Setting	Value	Description
Check interval	integer	Define in minutes how often smartd activates to check if any tests are
		configured to run.
Power mode	drop-down menu	Tests are not performed if the system enters the specified power mode:
		Never, Sleep, Standby, or Idle.
Difference	integer in degrees	Enter number of degrees in Celsius. S.M.A.R.T reports if the temperature
	Celsius	of a drive has changed by N degrees Celsius since the last report. Default
		of <i>0</i> disables this option.
Informational	integer in degrees	Enter a threshold temperature in Celsius. S.M.A.R.T will message with a
	Celsius	log level of LOG_INFO if the temperature is higher than specified degrees
		in Celsius. Default of 0 disables this option.
Critical	integer in degrees	Enter a threshold temperature in Celsius. S.M.A.R.T will message with a
	Celsius	log level of LOG_CRIT and send an email if the temperature is higher than
		specified degrees in Celsius. Default of 0 disables this option.
Email to report	string	Email address to receive S.M.A.R.T. alerts. Use a space to separate multi-
		ple email addresses.

Table 11.10: S.M.A.R.T Configuration Options

11.13 SMB

The settings configured when creating SMB Shares in *Sharing* \rightarrow *Windows (SMB) Shares* \rightarrow *Add Windows (SMB) Share* are specific to each configured SMB Share. In contrast, global settings which apply to all SMB shares are configured in Services \rightarrow *SMB*.

Note: After starting the SMB service, it can take several minutes for the master browser election (https://www.samba.org/samba/docs/old/Samba3-HOWTO/NetworkBrowsing.html#id2581357) to occur and for the

FreeNAS[®] system to become available in Windows Explorer.

Figure 11.15 shows the global SMB configuration options which are described in Table 11.11. This configuration screen is really a front-end to smb4.conf (https://www.freebsd.org/cgi/man.cgi?query=smb4.conf).

S	ettings			ж
	NetBIOS name:	freenas		
	NetBIOS alias:			
	Workgroup:	WORKGROUP	(i)	
	Description:	FreeNAS Server	(i)	
	DOS charset:	CP437 💌		
	UNIX charset:	UTF-8		
	Log level:	Minimum 💌		
	Use syslog only:			
	Local Master:			
	Domain logons:			
	Time Server for Domain:			
	Guest account:	nobody	(i)	
	File mask:		ì	
	Directory mask:		(i)	
	Allow Empty Password:			
	Auxiliary parameters:			
240		1		
	Unix Extensions:	(i)		

	Table I		
Setting	Value	Description	
NetBIOS Name	string	Automatically populated with the original hostname of the system. Lim-	
		ited to 15 characters. It must be different from the <i>Workgroup</i> name.	
NetBIOS Alias	string	Enter an alias. Limited to 15 characters	
Workgroup	string	Must match Windows workgroup name. This setting is ignored if the <i>Active</i>	
		<i>Directory</i> (page 169) or <i>LDAP</i> (page 174) service is running.	
Description	string	Enter an optional server description.	
DOS charset	drop-down menu	The character set Samba uses when communicating with DOS and Win-	
		dows 9x/ME clients. Default is <i>CP437</i> .	
UNIX charset	drop-down menu	Default is UTF-8 which supports all characters in all languages.	
Log level	drop-down menu	Choices are <i>Minimum</i> , <i>Normal</i> , or <i>Debug</i> .	
Use syslog only	checkbox	Set to log authentication failures to /var/log/messages instead of the	
		default of /var/log/samba4/log.smbd.	
Local Master	checkbox	Set to determine if the system will participate in a browser election. Dis-	
		able when network contains an AD or LDAP server or Vista or Windows 7	
		machines are present.	
Domain logons	checkbox	Set if it is necessary to provide the netlogin service for older Windows	
_		clients.	
Time Server for	checkbox	Determines if the system advertises itself as a time server to Windows	
Domain		clients. Disable when network contains an AD or LDAP server.	
Guest Account	drop-down menu	Select the account to be used for guest access. Default is <i>nobody</i> . Account	
		must have permission to access the shared volume/dataset. If Guest Ac-	
		count user is deleted, resets to <i>nobody</i> .	
File mask	integer	Overrides default file creation mask of 0666 which creates files with read	
		and write access for everybody.	
Directory mask	integer	Overrides default directory creation mask of 0777 which grants directory	
2		read, write and execute access for everybody.	
Allow Empty	checkbox	Set to allow users to press Enter when prompted for a password. Re-	
Password		quires the username/password to be the same as the Windows user ac-	
		count.	
Auxiliary param-	string	Add any smb.conf options not covered else-	
eters		where in this screen. See the Samba Guide	
		(http://www.oreilly.com/openbook/samba/book/appb_02.html) for	
		additional settings.	
Unix Extensions	checkbox	Set to allow non-Windows SMB clients to access symbolic links and hard	
		links, has no effect on Windows clients.	
Zeroconf share	checkbox	Enable if Mac clients will be connecting to the SMB share.	
discovery			
Hostname	checkbox	Set to allow using hostnames rather than IP addresses in the Hosts Al-	
lookups		<i>low</i> or <i>Hosts Deny</i> fields of a SMB share. Unset if IP addresses are used to	
		avoid the delay of a host lookup.	
Allow execute	checkbox	If set, Samba will allow the user to execute a file, even if that user's per-	
always		missions are not set to execute.	
Obey pam re-	checkbox	Unset this option to allow: Cross-domain authentication. Users and	
strictions		groups to be managed on another forest. Permissions to be delegated	
		from <i>Active Directory</i> (page 169) users and groups to domain admins on	
		another forest.	
	checkbox	Set to allow NTLMv1 authentication. Required by Windows XP clients and	
NTLMv1 auth		sometimes by clients in later versions of Windows.	
NTLMv1 auth			
Bind IP Ad-	checkboxes	Select the IPv4 and IPv6 addresses SMB will listen on. Always add the	
	checkboxes		
Bind IP Ad-	checkboxes	Select the IPv4 and IPv6 addresses SMB will listen on. Always add the	

Table 11.11: Global SMB C	Configuration Options
---------------------------	-----------------------

Setting	Value	Description
Idmap Range	integer	The beginning UID/GID for which this system is authoritative. Any UID/GID
Low		lower than this value is ignored, providing a way to avoid accidental
		UID/GID overlaps between local and remotely defined IDs.
Idmap Range	integer	The ending UID/GID for which this system is authoritative. Any UID/GID
High		higher than this value is ignored, providing a way to avoid accidental
		UID/GID overlaps between local and remotely defined IDs.

Table 11.11 – continued from previous page

Changes to SMB settings take effect immediately. Changes to share settings only take effect after the client and server negotiate a new session.

Note: Do not set the *directory name cache size* as an *Auxiliary parameter*. Due to differences in how Linux and BSD handle file descriptors, directory name caching is disabled on BSD systems to improve performance.

Note: SMB (page 238) cannot be disabled while Active Directory (page 169) is enabled.

11.13.1 Troubleshooting SMB

Do not connect to SMB shares as root, and do not add the root user in the SMB user database. There are security implications in attempting to do so, and Samba 4 and later take measures to prevent such actions. This can produce <code>auth_check_ntlm_password</code> and <code>FAILED</code> with error <code>NT_STATUS_WRONG_PASSWORD</code> errors.

Samba is single threaded, so CPU speed makes a big difference in SMB performance. A typical 2.5Ghz Intel quad core or greater should be capable of handling speeds in excess of GiB LAN while low power CPUs such as Intel Atoms and AMD C-30sE-350E-450 will not be able to achieve more than about 30-40 MiB/sec typically. Remember that other loads such as ZFS will also require CPU resources and may cause Samba performance to be less than optimal.

Samba's *write cache* parameter has been reported to improve write performance in some configurations and can be added to the *Auxiliary parameters* field. Use an integer value which is a multiple of _SC_PAGESIZE (typically *4096*) to avoid memory fragmentation. This will increase Samba's memory requirements and should not be used on systems with limited RAM.

Windows automatically caches file sharing information. If changes are made to an SMB share or to the permissions of a volume/dataset being shared by SMB and the share becomes inaccessible, try logging out and back in to the Windows system. Alternately, users can type net use /delete from the command line to clear their SMB sessions.

Windows also automatically caches login information. To require users to log in every time they access they system, reduce the cache settings on the client computers.

Where possible, avoid using a mix of case in filenames as this can cause confusion for Windows users. Representing and resolving filenames with Samba (http://www.oreilly.com/openbook/samba/book/ch05_04.html) explains in more detail.

If a particular user cannot connect to a SMB share, ensure their password does not contain the ? character. If it does, have the user change the password and try again.

If permissions work for Windows users but not for macOS users, try disabling *Unix Extensions* and restarting the SMB service.

If the SMB service will not start, run this command from *Shell* (page 289) to see if there is an error in the configuration:

testparm /usr/local/etc/smb4.conf

If clients have problems connecting to the SMB share, go to Services \rightarrow SMB and verify that Server maximum protocol is set to SMB2.

Using a dataset for SMB sharing is recommended. When creating the dataset, make sure that the *Share type* is set to Windows.

Do not use chmod to attempt to fix the permissions on a SMB share as it destroys the Windows ACLs. The correct way to manage permissions on a SMB share is to manage the share security from a Windows system as either the owner of the share or a member of the group that owns the share. To do so, right-click on the share, click *Properties* and navigate to the

Security tab. If the ACLs are already destroyed by using chmod, winacl can be used to fix them. Type winacl from Shell (page 289) for usage instructions.

The Common Errors (https://www.samba.org/samba/docs/old/Samba3-HOWTO/domain-member.html#id2573692) section of the Samba documentation contains additional troubleshooting tips.

The Samba Performance Tuning (https://wiki.samba.org/index.php/Performance_Tuning) page describes options to improve performance.

Directory listing speed in folders with a large number of files is sometimes a problem. A few specific changes can help improve the performance. However, changing these settings can affect other usage. In general, the defaults are adequate. **Do not change these settings unless there is a specific need.**

- *Hostname Lookups* and *Log Level* can also have a performance penalty. When not needed, they can be disabled or reduced in the *global SMB service options* (page 241).
- Make Samba datasets case insensitive by setting *Case Sensitivity* to *Insensitive* when creating them. This ZFS property
 is only available when creating a dataset. It cannot be changed on an existing dataset. To convert such datasets, back
 up the data, create a new case-insensitive dataset, create an SMB share on it, set the share level auxiliary parameter *case sensitive = true*, then copy the data from the old one onto it. After the data has been checked and verified on the
 new share, the old one can be deleted.
- If present, remove options for extended attributes and DOS attributes in Auxiliary Parameters (page 196) for the share.
- Disable as many VFS Objects as possible in the share settings (page 196). Many have performance overhead.

11.14 SNMP

SNMP (Simple Network Management Protocol) is used to monitor network-attached devices for conditions that warrant administrative attention. FreeNAS[®] uses Net-SNMP (http://net-snmp.sourceforge.net/) to provide SNMP. When starting the SNMP service, this port will be enabled on the FreeNAS[®] system:

• UDP 161 (listens here for SNMP requests)

Available MIBS are located in /usr/local/share/snmp/mibs.

Figure 11.16 shows the SNMP configuration screen. Table 11.12 summarizes the configuration options.

Settings		_	8
Location:		(i)	
Contact:		(i)	
SNMP v3 Support:			
Community:	public	(i)	
Username:			
Authentication Type:	SHA		
Password:			
Confirm Password:			
Privacy Protocol:			
Privacy Passphrase:			
Confirm Privacy Passphrase:			
Log Level:	Error		
Auxiliary parameters:			ì
OK Cancel			

Fig. 11.16: Configuring SNMP

Setting	Value	Description
Location	string	Optional description of the system location.
Contact	string	Optional. Enter the administrator email address.
SNMP v3 Sup-	checkbox	Set to enable support for SNMP version 3.
port		
Community	string	Default is <i>public</i> . Change this for security reasons! The value can only contain alphanumeric characters, underscores, dashes, periods, and spaces. This value can be empty for SNMPv3 networks.

Setting	Value	Description
Username	string	Only applies if SNMP v3 Support is set. Specify the username to
		register with this service. Refer to snmpd.conf(5) (http://net-
		snmp.sourceforge.net/docs/man/snmpd.conf.html) for more information
		about configuring this and the Authentication Type, Password, Privacy Proto-
		col, and Privacy Passphrase fields.
Authentication	drop-down menu	Only applies if SNMP v3 Support is enabled. Choices are: MD5 or SHA.
Туре		
Password	string	Only applies if SNMP v3 Support is enabled. Specify and confirm a pass-
		word of at least eight characters.
Privacy Protocol	drop-down menu	Only applies if SNMP v3 Support is enabled. Choices are: AES or DES.
Privacy	string	If not specified, <i>Password</i> is used.
Passphrase		
Log Level	drop-down menu	Choices range from fewest log entries (<i>Emergency</i>) to the most (<i>Debug</i>).
Auxiliary Param-	string	Enter additional snmpd.conf(5) (http://net-
eters		snmp.sourceforge.net/docs/man/snmpd.conf.html) options not covered in
		this screen. One option per line.

Table 11.12 – continued from previous page

11.15 SSH

Secure Shell (SSH) is used to transfer files securely over an encrypted network. When a FreeNAS[®] system is used as an SSH server, the users in the network must use SSH client software (https://en.wikipedia.org/wiki/Comparison_of_SSH_clients) to transfer files with SSH.

This section shows the FreeNAS[®] SSH configuration options, demonstrates an example configuration that restricts users to their home directory, and provides some troubleshooting tips.

Figure 11.17 shows the Services \rightarrow SSH configuration screen. After configuring SSH, remember to start it in Services \rightarrow Control Services.

s	ѕн		26
	TCP Port	22	i
	Login as Root with password	i)	
	Allow Password Authentication		
	Allow TCP Port Forwarding		
	Compress Connections		
	OK Cancel Advanced Mode	2	

Fig. 11.17: SSH Configuration

Table 11.13 summarizes the configuration options. Some settings are only available in *Advanced Mode*. To see these settings, either click the *Advanced Mode* button, or configure the system to always display these settings by enabling the *Show advanced fields by default* option in *System* \rightarrow *Advanced*.

Setting	Value	Advanced Mode	Description
Bind Interfaces	selection	✓	By default, SSH listens on all interfaces unless specific inter- faces are highlighted in the <i>Available</i> field and added to the <i>Se-</i> <i>lected</i> field.
TCP Port	integer		Port to open for SSH connection requests. 22 by default.
Login as Root with password	checkbox		As a security precaution, root logins are discouraged and disabled by default. If enabled, a password must be set for the <i>root</i> user in <i>View Users</i> .
Allow Password Au- thentication	checkbox		Unset to require key-based authentica- tion for all users. Requires additional setup (http://the.earth.li/~sgtatham/putty/0.55/htmldoc/Chapter8.html) on both the SSH client and server.
Allow Kerberos Au- thentication	checkbox	✓	Before setting this option, ensure <i>Kerberos Realms</i> (page 177) and <i>Kerberos Keytabs</i> (page 178) are configured and FreeNAS [®] can communicate with the Kerberos Domain Controller (KDC).
Allow TCP Port For- warding	checkbox		Set to allow users to bypass firewall restric- tions using the SSH port forwarding feature (https://www.symantec.com/connect/articles/ssh-port- forwarding).
Compress Connec- tions	checkbox		Set to attempt to reduce latency over slow networks.
SFTP Log Level	drop-down menu	\checkmark	Select the syslog(3) (https://www.freebsd.org/cgi/man.cgi?query=syslog) level of the SFTP server.
SFTP Log Facility	drop-down menu	✓	Select the syslog(3) (https://www.freebsd.org/cgi/man.cgi?query=syslog) facil- ity of the SFTP server.
Extra Options	string	V	Add any additional sshd_config(5) (https://www.freebsd.org/cgi/man.cgi?query=sshd_config) options not covered in this screen, one per line. These options are case-sensitive and misspellings can prevent the SSH service from starting.

Table 11.13: SSH Configuration Options

A few sshd_config(5) (https://www.freebsd.org/cgi/man.cgi?query=sshd_config) options that are useful to enter in the *Extra Options* field include:

- increase the *ClientAliveInterval* if SSH connections tend to drop
- *ClientMaxStartup* defaults to 10. Increase this value if more concurrent SSH connections are required.

11.15.1 SCP Only

When SSH is configured, authenticated users with a user account created using $Account \rightarrow Users \rightarrow Add User$ can use ssh to log into the FreeNAS[®] system over the network. The user home directory is the pool or dataset specified in the *Home Directory* field of the FreeNAS[®] account for that user. While the SSH login defaults to the user home directory, users are able to navigate outside their home directory, which can pose a security risk.

It is possible to allow users to use scp and sftp to transfer files between their local computer and their home directory on the FreeNAS[®] system, while restricting them from logging into the system using ssh. To configure this scenario, go to $Account \rightarrow Users \rightarrow View Users$, select the user, and click *Modify User*. Change the *Shell* to *scponly*. Repeat for each user that needs restricted SSH access.

Test the configuration from another system by running the sftp, ssh, and scp commands as the user. sftp and scp will work but ssh will fail.

Note: Some utilities like WinSCP and Filezilla can bypass the scponly shell. This section assumes that users are accessing the system using the command line versions of scp and sftp.

11.15.2 Troubleshooting SSH

Keywords listed in sshd_config(5) (https://www.freebsd.org/cgi/man.cgi?query=sshd_config) are case sensitive. This is important to remember when adding any *Extra options*. The configuration will not function as intended if the upper and lowercase letters of the keyword are not an exact match.

If clients are receiving "reverse DNS" or timeout errors, add an entry for the IP address of the FreeNAS[®] system in the Host name database field of Network \rightarrow Global Configuration.

When configuring SSH, always test the configuration as an SSH user account to ensure the user is limited by the configuration and they have permission to transfer files within the intended directories. If the user account is experiencing problems, the SSH error messages are specific in describing the problem. Type this command within *Shell* (page 289) to read these messages as they occur:

tail -f /var/log/messages

Additional messages regarding authentication errors are found in /var/log/auth.log.

11.16 TFTP

Trivial File Transfer Protocol (TFTP) is a light-weight version of FTP typically used to transfer configuration or boot files between machines, such as routers, in a local environment. TFTP provides an extremely limited set of commands and provides no authentication.

If the FreeNAS[®] system will be used to store images and configuration files for network devices, configure and start the TFTP service. Starting the TFTP service opens UDP port 69.

Figure 11.18 shows the TFTP configuration screen and Table 11.14 summarizes the available options.

Settings		88
Directory:	/tftproot	Browse
Allow New Files:		
Port:	69	ì
Username:	nobody	ì
File Permission:	Owner Group Other Read 🔽 🔽 💭 Write ☑ 🗐 🗐 Execute	
Extra options:		Ì
OK Cancel		

Fig. 11.18: TFTP Configuration

Setting	Value	Description
Directory	browse	Browse to an existing directory to be used for storage. Some devices re-
	button	quire a specific directory name. Refer to the device documentation for de-
		tails.
Allow New Files	checkbox	Enable if network devices need to send files to the system (for example, to
		back up their configuration).
Port	integer	Enter the UDP port to listen for TFTP requests. Default is 69.
Username	drop-down	Select the account to be used for TFTP requests. The account must have
	menu	permission to access the <i>Directory</i> .
File Permissions	checkboxes	Set permissions for newly created files. The default is everyone can read
		and only the owner can write. Some devices require less strict permis-
		sions.
Extra options	string	Add any additional tftpd(8) (https://www.freebsd.org/cgi/man.cgi?query=tftpd
		options not shown in this screen. Add one option on each line.

Table 11.14: TFTP (Configuration Options
---------------------	-----------------------

11.17 UPS

FreeNAS[®] uses NUT (http://networkupstools.org/) (Network UPS Tools) to provide UPS support. If the FreeNAS[®] system is connected to a UPS device, configure the UPS service then start it in *Services* \rightarrow *Control Services*.

Figure 11.19 shows the UPS configuration screen:

Settings			_	8
UPS Mode	:	Master		
Identifier:		ups	i	
Driver:			(i)	
Port:		•		
Auxiliary p	oarameters (ups.conf):			ì
Auxiliary p	oarameters (upsd.conf):			ì
Descriptio	in:			
Shutdown	mode:	UPS goes on battery		
Shutdown	timer:	30	i	
Shutdown	Command:	/sbin/shutdown -p now	ì	
No Comm	unication Warning Time:		١	
Monitor U	ser:	upsmon		
Monitor P	assword:	fixmepass		
Extra user	rs (upsd.users):			
Remote M	onitor:			
Send Ema	il Status Updates:			
To email:			ì	
Email Sub	ject:	UPS report generated by %h	1	
Power Off	UPS:			
OK	ncel			

Table 11.15 summarizes the options in the UPS Configuration screen.

Setting	Value	Description
UPS Mode	drop-down	Select <i>Master</i> if the UPS is plugged directly into the system serial port. The
	menu	UPS will remain the last item to shut down. Select <i>Slave</i> to have the system
	mena	shut down before <i>Master</i> .
Identifier	string	Describe the UPS device. Can contain alphanumeric, period, comma, hy-
lacitatie	Stille	phen, and underscore characters.
Driver / Remote Host	drop-down	For a list of supported devices, see the Network UPS Tools compatibility
Driver / Remote host	menu	list (https://networkupstools.org/stable-hcl.html).
	menu	Driver becomes Remote Host when UPS Mode is set to Slave. The IP ad-
		dress of the system configured as the UPS <i>Master</i> system. See this
		post (https://forums.freenas.org/index.php?resources/configuring-ups-
		support-for-single-or-multiple-freenas-servers.30/) for more details about
		configuring multiple systems with a single UPS.
Port / Hostname / Remote	drop-down	<i>Port</i> : Enter the serial or USB port the UPS is plugged into (see <i>NOTE</i>
Port Port	menu	(page 251)).
FOIL	menu	<i>Port</i> becomes <i>Remote Port</i> when the <i>UPS Mode</i> is set to <i>Slave</i> . The open
		network port number of the UPS <i>Master</i> system. The default port is 3493.
Auvilian / Daramatora	string	Enter any additional options from ups.conf(5)
Auxiliary Parameters	string	(https://www.freebsd.org/cgi/man.cgi?query=ups.conf).
(ups.conf)	string	
Auxiliary Parameters	string	Enter any additional options from upsd.conf(5)
(upsd.conf)		(https://www.freebsd.org/cgi/man.cgi?query=upsd.conf).
Description	string	Optional. Enter any notes about the UPS service.
Shutdown mode	drop-down	Choose when the UPS initiates shutdown. Choices are UPS goes on battery
	menu	and UPS reaches low battery.
Shutdown timer	integer	Select a value in seconds for the UPS to wait before initiating shutdown.
		Shutdown will not occur if the power is restored while the timer is count-
		ing down. The value only applies when <i>Shutdown Mode</i> is set to <i>UPS goes</i>
		on battery.
Shutdown Command	string	Enter the command to run to shut down the computer when battery
		power is low or shutdown timer runs out.
No Communication Warn-	string	Enter a value in seconds to wait before alerting that the service cannot
ing Time		reach any UPS. Warnings continue until the situation is fixed.
Monitor User	string	Enter a user to associate with this service. The recommended default user
		is upsmon.
Monitor Password	string	Default is the known value <i>fixmepass</i> . Change this to enhance system se-
		curity. Cannot contain a space or #.
Extra users	string	Enter the accounts with administrative access. See upsd.users(5)
		(http://networkupstools.org/docs/man/upsd.users.html) for examples.
Remote monitor	checkbox	Set for the default configuration to listen on all interfaces using the known
		values of user upsmon and password fixmepass.
Send Email Status Updates	checkbox	Set to enable the FreeNAS [®] system to send email updates to the config-
		ured <i>To email</i> address.
To email	email ad-	Enter the email address to receive status updates. Separate multiple email
	dress	addresses with a semicolon (;).
Email Subject	string	Enter a subject line to be used in email status updates.
Power Off UPS	checkbox	Set to power off the UPS after shutting down the FreeNAS system.

Table 11.15: UPS Configuration Options

Note: For USB devices, the easiest way to determine the correct device name is to enable the *Show console messages* option in *System* \rightarrow *Advanced*. Plug in the USB device and look for a */dev/ugen* or */dev/uhid* device name in the console messages.

Tip: Some UPS models might be unresponsive with the default polling frequency. This can show in FreeNAS[®] logs as a recurring error like: libusb_get_interrupt: Unknown error.

If this error occurs, increase the polling frequency by adding an entry to *Auxiliary Parameters (ups.conf)*: pollinterval = 10. The default polling frequency is two seconds.

upsc(8) (http://networkupstools.org/docs/man/upsc.html) can be used to get status variables from the UPS daemon such as the current charge and input voltage. It can be run from *Shell* (page 289) using this syntax:

upsc ups@localhost

The *upsc(8)* man page gives some other usage examples.

upscmd(8) (http://networkupstools.org/docs/man/upscmd.html) can be used to send commands directly to the UPS, assuming the hardware supports the command being sent. Only users with administrative rights can use this command. These users are created in the *Extra users* field.

11.17.1 Multiple Computers with One UPS

A UPS with adequate capacity can power multiple computers. One computer is connected to the UPS data port with a serial or USB cable. This *master* makes UPS status available on the network for other computers. These *slave* computers are powered by the UPS, but receive UPS status data from the master computer. See the NUT User Manual (http://networkupstools.org/docs/user-manual.chunked/index.html) and NUT User Manual Pages (http://networkupstools.org/docs/man/index.html#User_man).

11.18 WebDAV

The WebDAV service can be configured to provide a file browser over a web connection. Before starting this service, at least one WebDAV share must be created using *Sharing* \rightarrow *WebDAV Shares* \rightarrow *Add WebDAV Share*. Refer to *WebDAV Shares* (page 194) for instructions on how to create a share and connect to it when the service is configured and started.

Figure 11.20 shows the WebDAV configuration screen. Table 11.16 summarizes the available options.

W	/ebDAV	_	ж
	Protocol:	НТТР	
	HTTP Port:	8080	ì
	HTTP Authentication:	Digest Authentication \checkmark (1)	
	Webdav Password:		ì
	Confirm WebDAV Password:		
	ОК Cancel		

Fig. 11.20: WebDAV Configuration Screen

Setting	Value	Description	
Protocol	drop-down	HTTP keeps the connection always unencrypted. HTTPS always encrypts	
	menu	the connection. <i>HTTP</i> + <i>HTTPS</i> allows both types of connections.	
HTTP Port	string	Specify a port for unencrypted connections. Only appears if the selected	
		<i>Protocol</i> is <i>HTTP</i> or <i>HTTP</i> + <i>HTTPS</i> . The default of 8080 is recommended. Do	
		not reuse a port number.	
HTTPS Port	string	Specify a port for encrypted connections. Only appears if the selected <i>Pro</i> -	
		<i>tocol</i> is <i>HTTPS</i> or <i>HTTP+HTTPS</i> . The default of <i>8081</i> is recommended. Do	
		not reuse a port number.	
Webdav SSL Certificate	drop-down	Select the SSL certificate to use for encrypted connections. Only appears	
	menu	if the selected <i>Protocol</i> is <i>HTTPS</i> or <i>HTTP+HTTPS</i> . To create a certificate, use	
		System $ ightarrow$ Certificates.	
HTTP Authentication	drop-down	Choices are No Authentication, Basic Authentication (unencrypted), or Digest	
	menu	Authentication (encrypted).	
Webdav Password	string	Default is <i>davtest</i> . This is a known value and is recommended to be	
		changed.	

Table 11.16: WebDAV Configuration Options

PLUGINS

Warning: The legacy plugins infrastructure has been deprecated and is no longer supported. Plugins installation has been removed from the legacy UI but it can still be used to manage existing plugins. It is recommended to reinstall all legacy plugins using the new UI.

12.1 Installed Plugins

Entries for installed PBI will appear in these locations:

- the Installed tab of Plugins
- the *Plugins* section of the tree
- the Jails section of the tree

The entry in the *Installed* tab of Plugins displays the plugin name and version, the name of the PBI installed, the name of the jail, whether the application status is *ON* or *OFF*, and a button to delete the application and its associated jail.

Note: The *Service status* of a plugin must be turned to *ON* before the installed application is available. Before starting the service, check to see if it has a configuration menu by clicking its entry in the *Plugins* section of the tree. If the application is configurable, this will open a screen that contains the available configuration options. Plugins which are not configurable display a message with a hyperlink for accessing the software. However, that hyperlink does **not** work until the plugin is started.

Always review the configuration options of a plugin before attempting to start it. Some plugins have options that need to be set before their service will successfully start. If the application has not been configured before, check the website of the application to see what documentation is available.

If the application requires access to the data stored on the FreeNAS[®] system, click the entry for the associated jail in the *Jails* section of the tree and add a storage as described in *Add Storage* (page 261).

Access the shell of the jail containing the application by clicking the entry for the associated jail in the *Jails* section of the tree. You can then click its shell icon as described in *Managing Jails* (page 259).

Once the configuration is complete, click the red *OFF* button for the entry for the plugin. If the service starts successfully, it will change to a blue *ON*. If it fails to start, click the jail's *Shell* icon and type tail /var/log/messages to see if any errors were logged.

12.2 Deleting Plugins

Deleting a plugin deletes the associated jail as it is no longer required. **Before deleting a plugin**, make sure that there is no data or configuration options in the jail that need to be saved. Back up that data **before** deleting the plugin.

In the example shown in Figure 12.1, Sabnzbd is installed and the user has clicked the *Delete* button. A pop-up message displays. **This is the one and only warning.**

Plugins							
Available	Installed	Configurat	tion				
Plugin nan	ne Version	PBI		Jail	Service status	Actions	
Sabn:	_{zbd} 0.7.19	sabnzbd-	0.7.19-amd64	sabnzbd_1	OFF	Delete	
			Delete plug	gin		88	
				ire you want Cancel	to delete Plugir	ı sabnzbd?	

Fig. 12.1: Deleting an Installed Plugin

JAILS

This section describes how to use Jails, which allow users who are comfortable with the command line to have more control over software installation and management.

Warning: The jails infrastructure now uses uses the iocage backend and the warden backend has been deprecated and is no longer supported. Jail creation has been removed from the legacy UI but it can still be used to manage existing warden jails. It is recommended to recreate all legacy jails using the new UI, copy over any existing configurations, and delete the old jail datasets once the new jails are working as expected. To create new Jails, log into the new UI.

By default, a FreeBSD jail (https://en.wikipedia.org/wiki/Freebsd_jail) is created. This provides a very light-weight, operating system-level virtualization. Consider it as another independent instance of FreeBSD running on the same hardware, without all of the overhead usually associated with virtualization. The jail installs the FreeBSD software management utilities so FreeBSD ports can be compiled and FreeBSD packages can be installed from the command line of the jail.

It is important to understand that any users, groups, installed software, and configurations within a jail are isolated from both the FreeNAS[®] operating system and any other jails running on that system.

The rest of this section describes:

- Jails Configuration (page 256)
- Managing Jails (page 259)
- Starting Installed Software (page 264)

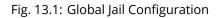
13.1 Jails Configuration

Jails are stored in a volume or dataset. **Using a separate dataset for the** *Jail Root* **is strongly recommended**. The volume or dataset to be used must already exist or can be created with *Volume Manager* (page 125).

Note: The Jail Root volume or dataset cannot be created on a Share (page 180).

Begin global jail configuration by choosing $Jails \rightarrow Configuration$ to open the screen shown in Figure 13.1. Jails are automatically installed into their own dataset under the specified path as they are created. For example, if the *Jail Root* is set to /mnt/volume1/dataset1 and a jail named *jail1* is created, it is installed into its own dataset named /mnt/volume1/ dataset1/jail1.

Jails	
Configuration	
Jail Root:	Browse
IPv4 DHCP:	(i)
IPv6 Autoconfigure:	
Save Advanced Mod	le



Warning: If any *Plugins* (page 254) are already installed, the *Jail Root, IPv4 Network, IPv4 Network Start Address*, and *IPv4 Network End Address* are automatically filled. Double-check that the pre-configured IP address values are appropriate for the jails and do not conflict with addresses used by other systems on the network.

Table 13.1 summarizes the fields in this configuration screen. Refer to the text below the table for more details on how to properly configure the *Jail Root* and network settings. Some settings are only available in *Advanced Mode*. To see these settings, either click the *Advanced Mode* button or configure the system to always display these settings by checking the box *Show advanced fields by default* in *System* \rightarrow *Advanced*.

Setting	Value	Advanced Mode	Description
Jail Root	browse but- ton		Mandatory. Jails cannot be added until this is set.
IPv4 DHCP	checkbox		Check this box if the network has a DHCP server.
IPv4 Network	string	✓	The format is IP address of <i>network/CIDR mask</i> .
IPv4 Network Start Address	string	\checkmark	Enter the first IP address in the reserved range in the format <i>host/CIDR mask</i> .
IPv4 Network End Address	string	\checkmark	Enter the last IP address in the reserved range in the format <i>host/CIDR mask</i> .
IPv6 Autoconfigure	checkbox		Check this box if the network has a DHCPv6 server and IPv6 will be used to access jails.
IPv6 Network	string	\checkmark	Enter the network address for a properly configured IPv6 net- work.
IPv6 Network Start Address	string	\checkmark	Enter the first IP address in the reserved range for a properly configured IPv6 network.
IPv6 Network End Address	string	\checkmark	Enter the last IP address in the reserved range for a properly configured IPv6 network.
Collection URL	string	\checkmark	Changing the default may break the ability to install jails.

Table 13.1:	Jail Configuration	n Options
	jun connguiudo	i options

When selecting the *Jail Root*, ensure that the size of the selected volume or dataset is sufficient to hold the number of jails to be installed as well as any software, log files, and data to be stored within each jail. At a bare minimum, budget at least 2 GiB per jail and do not select a dataset that is less than 2 GiB in size.

Note: When adding storage to a jail, be aware that the path size is limited to 88 characters. Make sure that the length of the volume name plus the dataset name plus the jail name does not exceed this limit.

If the network contains a DHCP server, it is recommended to check the box *IPv4 DHCP* (or *IPv6 Autoconfigure*, for a properly

configured IPv6 network). This prevents IP address conflicts on the network as the DHCP server automatically assigns the jail the next available lease and records the lease as in use.

If a static IP address is needed so that users always know the IP address of the jail, enter the start and end address for the IPv4 and/or IPv6 network. The range defined by the start and end addresses will be automatically assigned as jails are created. For example, when creating 5 jails on the 192.168.1.0 network, enter a IPv4 Network Start Address of 192.168.1.100 and a IPv4 Network End Address of 192.168.1.104.

When creating a start and end range on a network that contains a DHCP server, it is important to also reserve those addresses on the DHCP server. Otherwise, the DHCP server is not aware that those addresses are being used by jails. This lead to IP address conflicts and weird networking errors on the network.

FreeNAS® automatically detects and displays the IPv4 Network to which the administrative interface is connected. This setting is important. The IP addresses used by the jails must be pingable from the FreeNAS[®] system for the jails and any installed software to be accessible. If the network topology requires changing the default value, a default gateway and possibly a static route need to be added to the specified network. After changing this value, ensure that the subnet mask value is correct, as an incorrect mask can make the IP network unreachable. When in doubt, keep the default setting for IPv4 Network. With VMware, make sure that the vswitch is set to "promiscuous mode". With VirtualBox, make sure Network -> Advanced -> Promiscuous Mode is not set to "Deny".

After clicking the Save button to save the configuration, the system is ready to create and manage jails as described in the rest of this chapter.

Table 13.2 summarizes the available options. Most settings are only available in Advanced Mode and are not needed if the intent is to create a FreeBSD jail. To see these settings, either click the Advanced Mode button or configure the system to always display these settings by checking the box Show advanced fields by default in System \rightarrow Advanced.

Setting	Value	Advanced Mode	Description
Jail Name	string		Mandatory. Can only contain letters, numbers, dashes, or the underscore character.
Template	drop-down menu	\checkmark	Contains any created custom templates.
IPv4 DHCP	checkbox	\checkmark	If unchecked, make sure that the defined address does not conflict with the DHCP server's pool of available addresses.
IPv4 address	integer	√	This and the other IPv4 settings are grayed out if <i>IPv4 DHCP</i> is checked. Enter a unique IP address that is in the local network and not already used by anyother computer.
IPv4 netmask	drop-down menu	\checkmark	Select the subnet mask associated with <i>IPv4 address</i> .
IPv4 bridge address	integer	\checkmark	Grayed out unless <i>VIMAGE</i> is checked. See NOTE below.
IPv4 bridge netmask	drop-down menu	~	Select the subnet mask associated with <i>IPv4 bridge address</i> . Grayed out unless <i>VIMAGE</i> is checked.
IPv4 default gateway	string	\checkmark	Grayed out unless <i>VIMAGE</i> is checked.
IPv6 Autoconfigure	checkbox	\checkmark	If unchecked, make sure that the defined address does not conflict with the DHCP server's pool of available addresses.
IPv6 address	integer	√	This and other IPv6 settings are grayed out if <i>IPv6 Autoconfig-ure</i> is checked. Enter a unique IPv6 address that is in the local network and not already used by any other computer.
IPv6 prefix length	drop-down menu	\checkmark	Select the prefix length associated with <i>IPv6 address</i> .
IPv6 bridge address	integer	\checkmark	Grayed out unless <i>VIMAGE</i> is checked. See NOTE below.
IPv6 bridge prefix length	drop-down menu	\checkmark	Grayed out unless <i>VIMAGE</i> is checked. Select the prefix length associated with <i>IPv6 address</i> .
IPv6 default gateway	string	\checkmark	Grayed out unless <i>VIMAGE</i> is checked. Used to set the jail's default gateway IPv6 address.

Table 13.2: Jail Configuration Options

Continued on next page

Setting	Value	Advanced	Description
		Mode	
MAC	string	√	Grayed out unless <i>VIMAGE</i> is checked. Unique static MAC ad- dresses must be entered for every jail created if a static MAC address is entered.
NIC	drop-down	√	Grayed out if <i>VIMAGE</i> is checked. Can be used to specify the
	menu		interface to use for jail connections.
Sysctls	string	\checkmark	Comma-delimited list of sysctls to set inside jail (like <i>al-low.sysvipc=1,allow.raw_sockets=1</i>)
Autostart	checkbox	 ✓ 	Uncheck if the jail will be started manually.
VIMAGE	checkbox	\checkmark	Gives a jail its own virtualized network stack. Requires promis- cuous mode be enabled on the interface.
NAT	checkbox	\checkmark	Grayed out for Linux jails or if <i>VIMAGE</i> is unchecked. Enables Network Address Translation for the jail.

Table 13.2 – continued from previous page

Note: The IPv4 and IPv6 bridge interface is used to bridge the epair(4) (https://www.freebsd.org/cgi/man.cgi?query=epair) device, which is automatically created for each started jail, to a physical network device. The default network device is the one that is configured with a default gateway. So, if *em0* is the FreeBSD name of the physical interface and three jails are running, these virtual interfaces are automatically created: *bridge0, epair0a, epair1a,* and *epair2a.* The physical interface *em0* will be added to the bridge, as well as each epair device. The other half of the epair is placed inside the jail and is assigned the IP address specified for that jail. The bridge interface is assigned an alias of the default gateway for that jail or the bridge IP, if configured; either is correct.

13.2 Managing Jails

Click *Jails* to view and configure the added jails. In the example shown in Figure 13.2, the list entry for the jail named *xdm_1* has been clicked to enable that jail's configuration options. The entry indicates the jail name, IP address, whether it will start automatically at system boot, if it is currently running, and jail type: *standard* for a FreeBSD jail, or *pluginjail* if it was installed using *Plugins* (page 254).

Jails						
Jails	Storage Te	mplates Configur	ation			
Jails can no l	longer be created th	nrough legacy UI. P	lease use the new UI to create jails.			
Jail			IPv4 Address	Autostart	Status	Туре
freebsd1			DHCP	true	Running	standard
۴ 🖌		E				

Fig. 13.2: Viewing Jails

From left to right, these configuration icons are available:

Edit Jail: edit the jail settings which were described in Table 13.2.

After a jail has been created, the jail name and type cannot be changed. These fields are grayed out.

Note: To modify the IP address information for a jail, use the *Edit Jail* button instead of the associated networking commands from the command line of the jail.

Add Storage: configure the jail to access an area of storage as described in Add Storage (page 261).

Start/Stop: this icon changes appearance depending on the current *Status* of the jail. When the jail is not running, the icon is green and clicking it starts the jail. When the jail is already running, the icon is red and clicking it stops the jail. A stopped jail and its applications are inaccessible until it is restarted.

Restart: restart the jail.

Shell: access a *root* command prompt to configure the selected jail from the command line. When finished, type exit to close the shell.

Delete: delete the jail and any periodic snapshots of it. The contents of the jail are entirely removed.

Warning: Back up data and programs in the jail before deleting it. There is no way to recover the contents of a jail after deletion.

13.2.1 Accessing a Jail Using SSH

ssh can be used to access a jail instead of the jail's *Shell* icon. This requires starting the ssh service and creating a user account for ssh access. Start by clicking the *Shell* icon for the desired jail.

Find the sshd_enable= line in the jail's /etc/rc.conf and set it to "YES":

```
sshd_enable="YES"
```

Then start the SSH daemon:

service sshd start

The first time the service runs, the jail's RSA key pair is generated and the key fingerprint and random art image displayed.

Add a user account by typing adduser and following the prompts. If the user needs superuser privileges, they must be added to the *wheel* group. For those users, enter *wheel* at this prompt:

Login group is user1. Invite user1 into other groups? []: wheel

After creating the user, set the *root* password so that the new user will be able to use the su command to gain superuser privilege. To set the password, type passwd then enter and confirm the desired password.

Finally, test from another system that the user can successfully ssh in and become the superuser. In this example, a user named *user1* uses ssh to access the jail at 192.168.2.3. The first time the user logs in, they will be asked to verify the fingerprint of the host:

```
ssh user1@192.168.2.3
The authenticity of host '192.168.2.3 (192.168.2.3)' can't be established.
RSA key fingerprint is 6f:93:e5:36:4f:54:ed:4b:9c:c8:c2:71:89:c1:58:f0.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.2.3' (RSA) to the list of known hosts.
Password: type_password_here
```

Note: Each jail has its own user accounts and service configuration. These steps must be repeated for each jail that requires SSH access.

13.2.2 Add Storage

It is possible to give a FreeBSD jail access to an area of storage on the FreeNAS[®] system. This is useful for applications that store a large amount of data or if an application in a jail needs access to the data stored on the FreeNAS[®] system. One example is transmission, which stores torrents. The storage is added using the mount_nullfs(8) (https://www.freebsd.org/cgi/man.cgi?query=mount_nullfs) mechanism, which links data that resides outside of the jail as a storage area within the jail.

To add storage, click the *Add Storage* button for a highlighted jail entry to open the screen shown in Figure 13.3. This screen can also be accessed by expanding the jail name in the tree view and clicking *Storage* \rightarrow *Add Storage*.

expand all collapse all	Jails						
Account	Jails Storage Templates Configuration						
🗄 🎬 System	Jails can no longer be created through	h legacy UI. Please use th	the new UI to create jails.				
🗄 🙆 Tasks		- /	,				
E See Network							
📧 📂 Storage	Jail	IPv4 Addr	ress	Autostart		Status	Туре
Directory Service	freebsd1	DHCP		true		Running	standard
📧 👩 Sharing							
🗄 💕 Services							
Plugins							
🖃 🎹 Jails							
Add Jail Templates			Language				
III View Jails			Add Storage		28		
III View Jail Templates							
🔧 Configuration			Jail:	freebsd1 🔻			
🖃 🏂 freebsd1							
🔧 Edit			Source:		Browse		
🖃 🚰 Storage							
Add Storage			Destination:		Browse		
UMs VMs			Read-Only:				
Reporting							
Guide			Create directory:	🔽 🛈			
🖀 Wizard							
Display System Processes			OK Cancel				
Shell							

Fig. 13.3: Adding Storage to a Jail

Browse to the Source and Destination, where:

- Source: is the directory or dataset on the FreeNAS[®] system which will be accessed by the jail. This directory **must** reside outside of the volume or dataset being used by the jail. This is why it is recommended to create a separate dataset to store jails, so the dataset holding the jails is always separate from any datasets used for storage on the FreeNAS[®] system.
- *Destination:* select an **existing, empty** directory within the jail to link to the *Source* storage area. If that directory does not exist yet, enter the desired directory name and check the *Create directory* box.

Storage is typically added because the user and group account associated with an application installed inside of a jail needs to access data stored on the FreeNAS[®] system. Before selecting the *Source*, it is important to first ensure that the permissions of the selected directory or dataset grant permission to the user/group account inside of the jail. This is not the default, as the users and groups created inside of a jail are totally separate from the users and groups of the FreeNAS[®] system.

The workflow for adding storage usually goes like this:

1. Determine the name of the user and group account used by the application. For example, the installation of the transmission application automatically creates a user account named *transmission* and a group account also named *transmission*. When in doubt, check the files /etc/passwd (to find the user account) and /etc/group (to find the group account) inside the jail. Typically, the user and group names are similar to the application name. Also, the UID and GID are usually the same as the port number used by the service.

A *media* user and group (GID 8675309) are part of the base system. Having applications run as this group or user makes it possible to share storage between multiple applications in a single jail, between multiple jails, or even between the host and jails.

- 2. On the FreeNAS[®] system, create a user account and group account that match the user and group names used by the application in the jail.
- 3. Decide whether the jail should have access to existing data or if a new area of storage will be set aside for the jail to use.
- 4. If the jail will access existing data, edit the permissions of the volume or dataset so the user and group accounts have the desired read and write access. If multiple applications or jails are to have access to the same data, create a new group and add each needed user account to that group.
- 5. If an area of storage is being set aside for that jail or individual application, create a dataset. Edit the permissions of that dataset so the user and group account has the desired read and write access.
- 6. Use the *Add Storage* button of the jail and select the configured volume/dataset as the *Source*.

To prevent writes to the storage, check the box *Read-Only*.

By default, the *Create directory* box is checked. This means that the directory will automatically be created under the specified *Destination* path if the directory does not already exist.

After storage has been added or created, it appears in the tree under the specified jail. In the example shown in Figure 13.4, a dataset named tank/data has been chosen as the *Source* as it contains the files stored on the FreeNAS[®] system. When the storage was created, the user browsed to /usr/local/ in the *Destination* field, then entered *test* as the directory. Since this directory did not already exist, it was created, because the *Create directory* box was left checked. The resulting storage was added to the *freebsd1* entry in the tree as /usr/local/test. The user has clicked this /usr/local/test entry to access the *Edit* screen.

expand all collapse all	Jails						
📧 🏭 Account	Jails	Storage	Templates	Configura	ition		
🖅 🏫 System	Jails can r	o longer be crea	ated through le	gacy UI. Ple	ease use th	e new UI to cre	eate jails.
+ 👩 Tasks							
E Network					-		_
🛨 🚰 Storage	Jail	IPv4 Address	Autostart	Status	Туре		
Directory Service	freebsd1	DHCP	true	Running	standard		
🗄 👩 Sharing	/usr/loc	al/test				88	1
GP Services	Marnoc						
🛖 Plugins							
🖃 🎹 Jails	Jail:		freebsd1 🔻				
Add Jail Templates							
🔟 View Jails	Sour	ce:	/mnt/tank/data			Browse	
View Jail Templates	Deat	ination:	/usr/local/test			Browse	
🔧 Configuration	Dest	ination:	usmocantest			browse	
🖃 🦹 freebsd1	Read	I-Only:					
🔧 Edit							
🖃 🚰 Storage	Crea	te directory:] (i)				
🔧 /usr/local/test							
🔧 /usr/local/test	Mou	nted?	~				
🔧 /usr/local/test	ок	Cancel	Delete				
Add Storage							

Fig. 13.4: Example Storage

Storage is normally mounted as it is created. To unmount the storage, uncheck the *Mounted?* box.

Note: A mounted dataset does not automatically mount any of its child datasets. While the child datasets may appear to be browsable inside the jail, any changes are not visible. Since each dataset is considered to be its own filesystem, each child dataset must have its own mount point. Separate storage must be created for any child datasets which need to be mounted.

To delete the storage, click the *Delete* button.

Warning: It is important to realize that added storage is really just a pointer to the selected storage directory on the FreeNAS[®] system. It does **not** copy that data to the jail. **Files that are deleted from the** *Destination* **directory in the jail are really deleted from the** *Source* **directory on the** FreeNAS[®] **system.** However, removing the jail storage entry only removes the pointer, leaving the data intact but not accessible from the jail.

13.3 Starting Installed Software

After packages or ports are installed, they need to be configured and started. If you are familiar with the software, look for the configuration file in /usr/local/etc or a subdirectory of it. Many FreeBSD packages contain a sample configuration file as a reference. If you are unfamiliar with the software, you will need to spend some time at the software's website to learn which configuration options are available and which configuration files require editing.

Most FreeBSD packages that contain a startable service include a startup script which is automatically installed to /usr/ local/etc/rc.d/. After the configuration is complete, the starting of the service can be tested by running the script with the onestart option. As an example, if openvpn is installed into the jail, these commands run its startup script and verify that the service started:

```
/usr/local/etc/rc.d/openvpn onestart
Starting openvpn.
/usr/local/etc/rc.d/openvpn onestatus
openvpn is running as pid 45560.
sockstat -4
USER COMMAND
                      PTD
                              FD
                                      PROTO
                                               LOCAL ADDRESS
                                                               FOREIGN ADDRESS
root openvpn
                      48386
                              4
                                      udp4
                                               *:54789
                                                                *:*
```

If it produces an error:

/usr/local/etc/rc.d/openvpn onestart
Starting openvpn.
/usr/local/etc/rc.d/openvpn: WARNING: failed to start openvpn

Run tail /var/log/messages to see if any error messages hint at the problem. Most startup failures are related to a misconfiguration: either a typo or a missing option in a configuration file.

After verifying that the service starts and is working as intended, add a line to /etc/rc.conf to start the service automatically when the jail is started. The line to start a service always ends in _enable="YES" and typically starts with the name of the software. For example, this is the entry for the openvpn service:

openvpn_enable="YES"

When in doubt, the startup script shows the line to put in /etc/rc.conf. This is the description in /usr/local/etc/rc. d/openvpn:

```
# This script supports running multiple instances of openvpn.
# To run additional instances link this script to something like
# % ln -s openvpn openvpn_foo
# and define additional openvpn_foo_* variables in one of
# /etc/rc.conf, /etc/rc.conf.local or /etc/rc.conf.d /openvpn_foo
#
# Below NAME should be substituted with the name of this script. By default
# it is openvpn, so read as openvpn_enable. If you linked the script to
# openvpn_foo, then read as openvpn_foo_enable etc.
#
# The following variables are supported (defaults are shown).
```

```
# You can place them in any of
# /etc/rc.conf, /etc/rc.conf.local or /etc/rc.conf.d/NAME
#
# NAME_enable="NO"
# set to YES to enable openvpn
```

The startup script also indicates if any additional parameters are available:

```
# NAME_if=
# driver(s) to load, set to "tun", "tap" or "tun tap"
#
# it is OK to specify the if_ prefix.
#
# # optional:
# NAME_flags=
# additional command line arguments
# NAME_configfile="/usr/local/etc/openvpn/NAME.conf"
# --config file
# NAME_dir="/usr/local/etc/openvpn"
# --cd directory
```

CHAPTER FOURTEEN

VIRTUAL MACHINES

A Virtual Machine (*VM*) is an environment on a host computer that can be used as if it were a separate physical computer. VMs can be used to run multiple operating systems simultaneously on a single computer. Operating systems running inside a VM see emulated virtual hardware rather than the actual hardware of the host computer. This provides more isolation than *Jails* (page 256), although there is additional overhead. A portion of system RAM is assigned to each VM, and each VM uses a *zvol* (page 135) for storage. While a VM is running, these resources are not available to the host computer or other VMs.

FreeNAS[®] VMs use the bhyve(8) (https://www.freebsd.org/cgi/man.cgi?query=bhyve) virtual machine software. This type of virtualization requires an Intel processor with Extended Page Tables (EPT) or an AMD processor with Rapid Virtualization Indexing (RVI) or Nested Page Tables (NPT).

To verify that an Intel processor has the required features, use *Shell* (page 289) to run grep VT-x /var/run/dmesg.boot. If the *EPT* and *UG* features are shown, this processor can be used with *bhyve*.

To verify that an AMD processor has the required features, use *Shell* (page 289) to run grep POPCNT /var/run/dmesg. boot. If the output shows the POPCNT feature, this processor can be used with *bhyve*.

Note: By default, new VMs have the bhyve(8) (https://www.freebsd.org/cgi/man.cgi?query=bhyve) –H option set. This causes the virtual CPU thread to yield when a HLT instruction is detected, and prevents idle VMs from consuming all of the host's CPU.

Note: AMD K10 "Kuma" processors include POPCNT but do not support NRIPS, which is required for use with bhyve. Production of these processors ceased in 2012 or 2013.

14.1 Creating VMs

Select *VMs* \rightarrow *Add VM* for the *Add VM* dialog shown in Figure 14.1:

A	dd VM		8
	VM Type:	Virtual Machine	
	Name:		ì
	Description:		ì
	Virtual CPUs:	1	Ì
	Memory Size (MiB):		Ì
	Boot Method:	UEFI 🔻 (1)	
	Autostart:		
	OK Cancel		

Fig. 14.1: Add VM

VM configuration options are described in Table 14.1.

Table 14.1:	VM	Options
-------------	----	---------

		-
Setting	Value	Description
VM Type	drop-down	Select the VM type. Virtual Machine is a typical instance, and Docker VM is a
	menu	special VM to run Docker.
Name	string	Enter a name to identify the VM.
Description	string	Enter a short description of the VM or its purpose.
Virtual CPUs	integer	Select the number of virtual CPUs to allocate to the VM. The maximum is 16 unless the host CPU limits the maximum. The VM operating system might also have operational or licensing restrictions on the number of CPUs.
Memory Size (MiB)	integer	Allocate the amount of RAM in mebibytes (https://simple.wikipedia.org/wiki/Mebibyte) for the VM.
Boot Method	drop-down menu	Select <i>UEFI</i> for newer operating systems, or <i>UEFI-CSM</i> for (Compatibility Support Mode) older operating systems that only understand BIOS booting.
Autostart	checkbox	Set to start the VM automatically when the system boots.

14.2 Adding Devices to a VM

After creating the VM, click it to select it, then click *Devices* and *Add Device* to add virtual hardware to it:

Add device 🛛 🕅		
VM:	🔻	
Type:	💌	
ОК	Cancel	

Fig. 14.2: Add Devices to a VM

Select the name of the VM from the VM drop-down menu, then select the Type of device to add. These types are available:

- Network Interfaces (page 268)
- Disk Devices (page 269)
- Raw Files (page 269)
- CD-ROMs (page 270)
- VNC Interface (page 270)

Note: Docker VMs (page 273) are not compatible with VNC connections.

Figure 14.3 shows the fields that appear when Network Interface is the selected Type.

14.2.1 Network Interfaces

E	dit		Ж
	VM:	samplevm	
	Туре:	Network Interface 🔻	
	Adapter Type:	Intel e82545 (e1000) 🔻	
	NIC to attach:	em0 💌	
	MAC Address:		i
	OK Cancel	Delete	

Fig. 14.3: VM Network Interface Device

The default *Adapter Type* emulates an Intel e82545 (e1000) Ethernet card for compatibility with most operating systems. *VirtIO* can provide better performance when the operating system installed in the VM supports VirtIO paravirtualized network drivers.

If the system has multiple physical network interface cards, use the *Nic to attach* drop-down menu to specify which physical interface to associate with the VM.

By default, the VM receives an auto-generated random MAC address. To override the default with a custom value, enter the desired address into the *MAC Address* field.

14.2.2 Disk Devices

Zvols (page 135) are typically used as virtual hard drives. After *creating a zvol* (page 135), associate it with the VM by selecting *Add device*.

Add device	8
VM:	samplevm
Туре:	Disk
ZVol:	•
Mode:	AHCI
Disk sectorsize:	0
OK Cancel	

Fig. 14.4: VM Disk Device

Choose the VM, select a Type of Disk, select the created zvol, then set the Mode:

- AHCI emulates an AHCI hard disk for best software compatibility.
- *VirtIO* uses paravirtualized drivers and can provide better performance, but requires the operating system installed in the VM to support VirtIO disk devices.

If a specific sector size is required, enter the number of bytes into *Disk sector size*. The default of 0 uses an autotune script to determine the best sector size for the zvol.

14.2.3 Raw Files

Raw Files are similar to *Zvol* (page 135) disk devices, but the disk image comes from a file. These are typically used with existing read-only binary images of drives, like an installer disk image file meant to be copied onto a USB stick.

After obtaining and copying the image file to the FreeNAS[®] system, select *Add device*, choose the *VM*, select a *Type* of *Raw File*, browse to the image file, then set the *Mode*. *AHCI* emulates an AHCI hard disk for best software compatibility. *VirtIO* uses paravirtualized drivers and can provide better performance, but requires the operating system installed in the VM to support VirtIO disk devices.

If a specific sector size is required, enter the number of bytes into *Disk sectorsize*. The default of 0 uses an autotuner to find and set the best sector size for the file.

Add device		8
VM:	samplevm	
Туре:	Raw File	
Mode:	AHCI	
Raw File:	/mnt/volume1/FreeBSD-11.1.	Browse
Disk sectorsize:	0	i
OK Cancel		

Fig. 14.5: VM Raw File Disk Device

14.2.4 CD-ROM Devices

Edit		88
VM:	samplevm	
Туре:	CD-ROM 🔽	
CD-ROM (ISO):	/mnt/volume1/isos/FreeBSD-	Browse
OK Cancel	Delete	

Adding a CD-ROM device makes it possible to boot the VM from a CD-ROM image, typically an installation CD. The image must be present on an accessible portion of the FreeNAS[®] storage. In this example, a FreeBSD installation image is shown:



Note: VMs from other virtual machine systems can be recreated for use in FreeNAS[®]. Back up the original VM, then create a new FreeNAS[®] VM with virtual hardware as close as possible to the original VM. Binary-copy the disk image data into the *zvol* (page 135) created for the FreeNAS[®] VM with a tool that operates at the level of disk blocks, like dd(1) (https://www.freebsd.org/cgi/man.cgi?query=dd). For some VM systems, it is best to back up data, install the operating system from scratch in a new FreeNAS[®] VM, and restore the data into the new VM.

14.2.5 VNC Interface

VMs set to *UEFI* booting are also given a VNC (Virtual Network Computing) remote connection. A standard VNC (https://en.wikipedia.org/wiki/Virtual_Network_Computing) client can connect to the VM to provide screen output and keyboard and mouse input.

Note: Each VM can only have a single VNC device.

Note: *Docker VMs* (page 273) are not compatible with VNC connections and cannot have a VNC interface.

Note: Using a non-US keyboard via VNC is not yet supported. As a workaround, select the US keymap on the system running the VNC client, then configure the operating system running in the VM to use a keymap that matches the physical keyboard. This will enable passthrough of all keys regardless of the keyboard layout.

Figure 14.7 shows the fields that appear when *VNC* is the selected *Type*.

Edit		ж
VM:	samplevm	
Туре:		
Resolution:	1920×1080 🔻	
VNC port:	0	ì
Bind to:	0.0.0.0	
Wait to boot:		
Password:		ì
VNC Web:		
OK Cancel	Delete	

Fig. 14.7: VM VNC Device

The *Resolution* drop-down menu can be used to modify the default screen resolution used by the VNC session.

The *VNC port* can be set to 0, left empty for FreeNAS[®] to assign a port when the VM is started, or set to a fixed, preferred port number.

Select the IP address for VNC to listen on with the Bind to drop-down menu.

Set Wait to boot to indicate that the VNC client should wait until the VM has booted before attempting the connection.

To automatically pass the VNC password, enter it into the *Password* field. Note that the password is limited to 8 characters. To use the VNC web interface, set *VNC Web*.

Tip: If a RealVNC 5.X Client shows the error RFB protocol error: invalid message type, disable the Adapt to network speed option and move the slider to Best quality. On later versions of RealVNC, select File \rightarrow Preferences, click Expert,

ProtocolVersion, then select 4.1 from the drop-down menu.

14.2.6 Virtual Serial Ports

VMs automatically include a virtual serial port.

- /dev/nmdm1B is assigned to the first VM
- /dev/nmdm2B is assigned to the second VM

And so on. These virtual serial ports allow connecting to the VM console from the Shell (page 289).

Tip: The nmdm (https://www.freebsd.org/cgi/man.cgi?query=nmdm) device is dynamically created. The actual nmdm name can differ on each system.

To connect to the first VM:

```
cu -s 9600 -l /dev/nmdm1B
```

See cu(1) (https://www.freebsd.org/cgi/man.cgi?query=cu) for more information on operating cu.

14.3 Running VMs

Select *VMs* to see a list of configured VMs. Configuration and control buttons appear at the bottom of the screen when an individual VM is selected with a mouse click:

VMs VMs Add VM						
Name	Description	Info	Virtual CPUs	Memory Size (MiB)	Boot Method	Autostart
samplevm		State: STOPPED VNC Port: 5901	1	2048	UEFI	false

Fig. 14.8: VM Configuration and Control Buttons

The name, description, running state, VNC port (if present), and other configuration values are shown. Click on an individual VM for additional options.

Some standard buttons are shown for all VMs:

- Edit changes VM settings.
- Delete removes the VM (page 273).
- Devices is used to add and remove devices to this VM.

When a VM is not running, these buttons are available:

- *Start* starts the VM.
- Clone clones or copies the VM to a new VM. The new VM is given the same name as the original, with cloneN appended.

When a VM is already running, these buttons are available:

- Stop shuts down the VM.
- Power off immediately halts the VM, equivalent to disconnecting the power on a physical computer.
- *Restart* restarts the VM.
- *Vnc via Web* starts a web VNC connection to the VM. The VM must have a VNC device and *VNC Web* enabled in that device.

14.4 Deleting VMs

A VM is deleted by clicking the VM, then *Delete* at the bottom of the screen. A dialog shows any related devices that will also be deleted and asks for confirmation.

Tip: *Zvols* (page 135) used in *disk devices* (page 269) and image files used in *raw file* (page 269) devices are *not* removed when a VM is deleted. These resources can be removed manually after it is determined that the data in them has been backed up or is no longer needed.

14.5 Docker VM

Docker (https://www.docker.com/what-docker) is Open Source software for automating application deployment inside containers. A container provides a complete filesystem, runtime, system tools, and system libraries, so applications always see the same environment.

Rancher (https://rancher.com/) is a web interface tool for managing Docker containers.

FreeNAS[®] runs the Rancher web interface within the Docker VM.

14.5.1 Docker VM Requirements

The system BIOS **must** have virtualization support enabled for a Docker VM to run properly after installation. On Intel systems this is typically an option called *VT-x*. AMD systems generally have an *SVM* option.

20 GiB of storage space is required for the Rancher VM. For setup, the SSH (page 245) service must be enabled.

The Rancher VM requires 2 GiB of RAM while running.

14.5.2 Create the Docker VM

Figure 14.9 shows the window that appears after going to the VMs page, clicking Add VM, and selecting Docker VM as the VM Type.

A	Add VM		
	VM Туре:	Docker VM	
	Name:	RancherUI	i
	Description:	RancherUI VM	Ì
	Virtual CPUs:	1	Ì
	Memory Size (MiB):	2,048	Ì
	Autostart:	(i)	
	Root Password:	••••	
	Docker Disk File:	/mnt/pool1/rancherui.img	Browse
	Size of Docker Disk File (GiB):	20	
	OK Cancel		

Fig. 14.9: Docker VM Configuration

Setting	Value	Description
VM	drop-down menu	Choose to
Туре		create ei-
		ther a stan-
		dard <i>Virtual</i>
		Machine or
		a Docker
		VM.
Name	string	Enter a de-
		scriptive
		name for
		the Docker
		VM.
Descript	iostring	Describe
		this Docker
		VM.
	Continuos	on next name

Table 14.2: Docker VM Options

Continued on next page

n
al-
er
é
n
ess
n.
р-
м /S-
, 5
•
nal
C-
:he
of
ונ
he
r
er
۱
rt
er
•
3
k
th
er
-
-
in

Table 14.2 – continued from previous page

 a space.

 Continued on next page

Setting	Value	Description
Docker	string	Browse to
Disk	String	the location
File		
File		to store a
		new raw
		file. Add /,
		a unique
		name to
		the end of
		the path,
		and .img
		to create
		a new raw
		file with
		that name.
		Example:
		/mnt/
		pool1/
		rancherui.
		img
Size of	integer	Allocate
Docker		storage size
Disk		in GiB for
File		the new
(GiB)		raw file.
		20 is the
		minimum
		recommen- dation.

Table 14.2 – continued from prev	vious page
----------------------------------	------------

Recommendations for the Docker VM:

- Enter Rancher UI VM for the Description.
- Leave the number of Virtual CPUs at 1.
- Enter 2048 for the Memory Size.
- Leave 20 as the Size of Docker Disk File (GiB).

Click *OK* to create the virtual machine.

To make any changes to the raw file after creating the Docker VM, click on the *Devices* button for the VM to show the devices attached to that VM. Click on the *RAW* device to select it, then click *Edit*. Figure 14.10 shows the options for editing the Docker VM raw file options.

E¢	dit	_	8
	VM:	RancherUI 🔻	
	Туре:	Raw File	
	Mode:	AHCI 💌	
	Raw File:	k/vm-storage/rancherui.img	Close
		🖃 🗁 /	
		📃 🗁 mnt	
		📄 🗁 tank	
		귿 vm-storage	
	Disk boot:		
	Password:	•••••	ì
	Disk sectorsize:	0	i
	Disk size:	20G	ì
	OK Cancel	Delete	

Fig. 14.10: Docker VM Image Storage

The *raw file options* (page 269) section describes the options in this window.

14.5.3 Start the Docker VM

Click VMs, then click on the Docker VM line to select it. Click the Start button and Yes to start the VM.

14.5.4 SSH into the Docker VM

It is possible to SSH into a running Docker VM. Go to the VMs page and find the entry for the Docker VM. The *Info* column shows the *Com Port* for the Docker VM. In this example, /dev/nmdm12B is used.

Use an SSH client to connect to the FreeNAS[®] server. Remember this also requires the *SSH* (page 245) service to be running. Depending on the FreeNAS[®] system configuration, it might also require changes to the *SSH* service settings, like setting *Login as Root with Password*.

At the FreeNAS[®] console prompt, connect to the running Docker VM with cu (https://www.freebsd.org/cgi/man.cgi?query=cu), replacing /dev/nmdm12B with the value from the Docker VM *Com Port*:

cu -l /dev/nmdm12B -s 9600

If the terminal does not show a rancher login: prompt, press Enter. The Docker VM can take some time to start and display the login prompt.

14.5.5 Installing and Configuring the Rancher Server

Using the Docker VM to install and configure the Rancher Server is done from the command line. Open the *Shell* (page 289) and enter the command cu -1 /dev/nmdm12B -s 9600, where /dev/nmdm12B is the *Com Port* value in the *Info* column for the Docker VM.

If the terminal does not show a rancher login: prompt after a few moments, press Enter.

Enter *rancher* as the username, press Enter, then type the password that was entered when the raw file was created above and press Enter again. After logging in, a [rancher@rancher ~]\$ prompt is displayed.

Ensure Rancher has functional networking and can ping an outside website. Adjust the VM *Network Interface* (page 268) and reboot the VM if necessary.

Download and install the Rancher system with this command:

sudo docker run -d --restart=unless-stopped -p 8080:8080 rancher/server

Note: If the error Cannot connect to the Docker daemon is shown, run sudo dockerd. Then give the sudo docker run command above again.

Installation time varies with processor and network connection speed, but typically takes a few minutes. After the process finishes and a command prompt is shown, type this command:

ifconfig eth0 | grep 'inet addr'

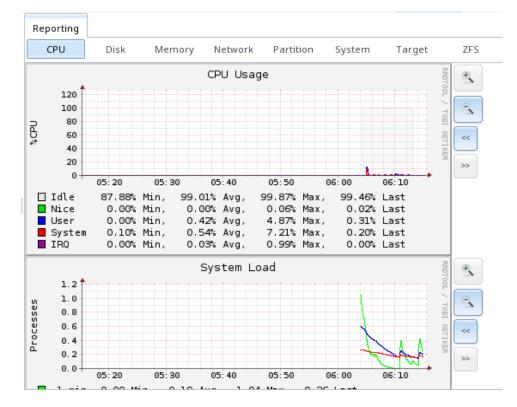
The first value is the IP address of the Rancher server. Enter the IP address and port 8080 as the URL in a web browser. For example, if the IP address was 10.231.3.208, enter 10.231.3.208:8080 as the URL in the web browser.

The Rancher server takes a few minutes to start. The web browser might show a connection error while the Rancher web interface is still starting. If the browser shows a connection has timed out or a similar error, wait one minute and try again.

In the Rancher web interface, click Add a host, ensure the radial This site's address button is set, and click Save. Follow the instructions that now display and run the sudo docker run --rm --privileged -v command in the Docker Host Serial shell. After the command runs a message displays Launched Rancher Agent:. Refresh or go to the Hosts page of the Rancher web interface to confirm the Docker Host displays in the web interface. Rancher is now configured and ready for use.

For more information on using RancherOS, see the RancherOS Documentation (https://rancher.com/docs/os/v1.x/en/).

REPORTING



Reporting displays several graphs, as seen in Figure 15.1. Click the tab for a device type to see those specific graphs.

Fig. 15.1: Reporting Graphs

FreeNAS[®] uses collectd (https://collectd.org/) to provide reporting statistics. The resulting graphs are grouped into several tabs on the Reporting page:

• CPU

- CPU (https://collectd.org/wiki/index.php/Plugin:CPU) shows the amount of time spent by the CPU in various states such as executing user code, executing system code, and being idle.
- Disk
 - Disk (https://collectd.org/wiki/index.php/Plugin:Disk) shows statistics on I/O, percent busy, latency, operations per second, pending I/O requests, and disk temperature.
- Memory
 - Memory (https://collectd.org/wiki/index.php/Plugin:Memory) displays memory usage.
 - Swap (https://collectd.org/wiki/index.php/Plugin:Swap) displays the amount of free and used swap space.
- Network

- Interface (https://collectd.org/wiki/index.php/Plugin:Interface) shows received and transmitted traffic in bits per second for each configured interface.
- Partition
 - Disk space (https://collectd.org/wiki/index.php/Plugin:DF) displays free and used space for each volume and dataset. However, the disk space used by an individual zvol is not displayed as it is a block device.
- System
 - Processes and Uptime (https://collectd.org/wiki/index.php/Plugin:Processes) displays the number of processes. It is grouped by state.
 - Uptime (https://collectd.org/wiki/index.php/Plugin:Uptime) keeps track of the system uptime, the average running time, and the maximum reached uptime.
- Target
 - Target shows bandwidth statistics for iSCSI ports.
- ZFS
 - ZFS (https://collectd.org/wiki/index.php/Plugin:ZFS_ARC) shows ARC size, hit ratio, and requests.

Reporting data is saved to permit viewing and monitoring usage trends over time. This data is preserved across system upgrades and restarts.

Data files are saved in /var/db/collectd/rrd/.

The reporting data file recording method is controlled by the *System* \rightarrow *System Dataset Reporting database* option. When deselected, data files are recorded in a temporary filesystem and copied hourly to on-disk files.

When System \rightarrow System Dataset Reporting database is enabled, data files are written directly to the System Dataset (page 76).

Warning: Reporting data is frequently written and should not be stored on the boot pool or boot device.

Use the magnifier buttons next to each graph to increase or decrease the displayed time increment from 10 minutes, hourly, daily, weekly, or monthly. The << and >> buttons can be used to scroll through the output.

Update on using Graphite with FreeNAS (http://cmhramblings.blogspot.com/2015/12/update-on-using-graphite-with-freenas.html) contains instructions for sending the collected information to a Graphite (http://graphiteapp.org/) server.

CHAPTER SIXTEEN

WIZARD

FreeNAS[®] provides a wizard which helps complete the steps needed to quickly configure FreeNAS[®] for serving data over a network. The wizard can be run at any time by clicking the *Wizard* icon.

Figure 16.1 shows the first wizard configuration screen.

W	lizard	X
	Language:	English
	Console Keyboard Map:	
	Timezone:	America/Los_Angeles 🔻
	Next	

Fig. 16.1: Configuration Wizard

Note: You can exit the wizard at any time by clicking the *Exit* button. However, exiting the wizard will not save any selections. The wizard can always be run again by clicking the *Wizard* icon. Alternately, the FreeNAS[®] GUI can be used to configure the system, as described in the rest of this Guide.

This screen can be used to change the default language, keyboard map, and timezone. After making your selections, click *Next*. The next screen depends on whether or not the storage disks have already been formatted into a ZFS pool.

Figure 16.2 shows the configuration screen that appears if the storage disks have not yet been formatted.

Wizard	
Volume Name:	
 Automatic (Reasonable defaults using the available drives) Virtualization (RAID 10: Moderate Redundancy, Maximum Performance, Minimum Capacity) Backups (RAID Z2: Moderate Redundancy, Moderate Performance, Moderate Capacity) Media (RAID Z1: Minimum Redundancy, Moderate Performance, Moderate Capacity) Logs (RAID 0: No Redundancy, Maximum Performance, Maximum Capacity) 	
Estimated Total Size: 0 Disks to be formatted: ada1, ada2, ada3 Next Exit	

Fig. 16.2: Volume Creation Wizard

Note: The wizard will not recognize an **encrypted** ZFS pool. If your ZFS pool is GELI-encrypted, cancel the wizard and use the instructions in *Importing an Encrypted Pool* (page 138) to import the encrypted volume. You can then rerun the wizard afterwards, if you wish to use it for post-configuration, and it will recognize that the volume has been imported and will not prompt to reformat the disks.

Enter a name for the ZFS pool that conforms to these naming conventions (https://docs.oracle.com/cd/E23824_01/html/821-1448/gbcpt.html). It is recommended to choose a name that will stick out in the logs (e.g. **not** data or freenas).

Decide if the pool should provide disk redundancy, and if so, which type. The *ZFS Primer* (page 320) discusses RAIDZ redundancy in more detail. If you prefer to make a more complex configuration, click the *Exit* button to close the wizard and instead use *Volume Manager* (page 125).

These redundancy types are available:

- **Automatic:** automatically creates a mirrored, RAIDZ1, or RAIDZ2 pool, depending upon the number of disks. If you prefer to control the type of redundancy, select one of the other options.
- RAID 10: creates a striped mirror and requires a minimum of 4 disks.
- **RAIDZ2:** requires a minimum of 4 disks. Up to 2 disks can fail without data loss.
- RAIDZ1: requires a minimum of 3 disks. Up to 1 disk can fail without data loss.
- Stripe: requires a minimum of 1 disk. Provides **no** redundancy, meaning if any of the disks in the stripe fails, all data in the stripe is lost.

Once you have made your selection, click Next to continue.

If the disks have already been formatted with ZFS and the disks have **not** been encrypted, the next screen will instead prompt to import the volume, as shown in Figure 16.3.



Fig. 16.3: Volume Import Screen

Select the existing volume from the drop-down menu and click *Next* to continue. The next screen in the wizard is shown in Figure 16.4.

_			
W	/izard		
	Directory Service:	Active Directory	
	-	· · · · · · · · · · · · · · · · · · ·	
	Domain Name (DNS/Realm-Name):		
	Domain Account Name:		
	Domain Account Password:		
	Previous Next Exit		

Fig. 16.4: Directory Service Selection

If the FreeNAS[®] system is on a network that does not contain an Active Directory, LDAP, or NIS server, click *Next* to skip to the next screen.

However, if the FreeNAS[®] system is on a network containing an Active Directory, LDAP, or NIS server and you wish to import the users and groups from that server, select the type of directory service in the *Directory Service* drop-down menu. The rest of the fields in this screen will vary, depending upon which directory service is selected. Available configuration options for each directory service are summarized in Tables 16.1 through 16.3.

Note: Additional configuration options are available for each directory service. The wizard can be used to set the initial values required to connect to that directory service. You can then review the other available options in *Directory Services* (page 169) to determine if additional configuration is required.

Table 16.1: Active Directory Options			
Setting	Value	Description	
Domain Name	string	Enter the name of Active Directory domain (e.g. <i>example.com</i>) or child do- main (e.g. <i>sales.example.com</i>).	
Domain Account Name	string	Enter the name of the Active Directory administrator account.	
Domain Account Password	string	Enter the password for the Active Directory administrator account.	

Table 16.2: LDAP Options

Setting	Value	Description
Hostname	string	Hostname or IP address of LDAP server.
Base DN	string	Top level of the LDAP directory tree to be used when searching for re-
		sources. Example: <i>dc=test,dc=org</i>

Continued on next page

Table 16.2 – continued from previous page			
Setting	Value	Description	
Bind DN	string	Name of the administrative account on the LDAP server. Example: <i>cn=Manager,dc=test,dc=org</i>)	
Base password	string	Password for the administrative account on the LDAP server.	

Tabla	16 3.	NIC	Options
rable	10.5.	INID	Options

		•
Setting	Value	Description
NIS domain	string	Name of the NIS domain.
NIS servers	string	Enter a comma-delimited list of hostnames or IP addresses.
Secure mode	checkbox	Set for ypbind(8) (https://www.freebsd.org/cgi/man.cgi?query=ypbind) to refuse to bind to any NIS server that is not running as root on a TCP port number over <i>1024</i> .
Manycast	checkbox	Set for <i>ypbind</i> to bind to the server that responds the fastest. This is useful when no local NIS server is available on the same subnet.

The next configuration screen, shown in Figure 16.5, is used to create network shares.

Wizard	36
Share name: Purpose Windows (SMB) Allow Guest Mac OS X (AFP) Time Machine Generic Unix (NFS) Block Storage (iSCSI) Size:	
Add Delete Update	
Name	
	-
	v
Previous Next Exit	

Fig. 16.5: Network Shares

FreeNAS[®] supports several types of shares for providing storage data to the clients in a network. The initial wizard can be used to quickly make shares using default permissions which should "just work" for common scenarios. For more complex scenarios, refer to the section on *Sharing* (page 180).

To create a share using the wizard, enter a name for the share, then select the *Purpose* of the share:

- Windows (SMB): this type of share can be accessed by any operating system using a SMB client. Check the box for *Allow Guest* to allow users to access the share without a password. SMB shares created with the wizard can be fine-tuned afterward with *Windows (SMB) Shares* (page 195).
- Mac OS X (AFP): this type of share can be accessed by Mac OS X users. Check the box for *Time Machine* if Mac users will be using the FreeNAS[®] system as a backup device. AFP shares created with the wizard can be fine-tuned afterward with *Apple (AFP) Shares* (page 181).
- Generic Unix (NFS): this type of share can be accessed by any operating system using a NFS client. NFS shares created using the wizard can be fine-tuned afterward with Unix (NFS) Shares (page 188).
- **Block Storage (iSCSI):** this type of share can be accessed by any operating system using iSCSI initiator software. Enter the size of the block storage to create in the format *20G* (for 20 GiB). iSCSI shares created with the wizard can be fine-tuned afterward with *iSCSI* (page 230).

W	izard		ж
	User:	root Create User	i
	Group:	wheel Create Group	i)
	Mode:	Owner Group Other Read 🔽 🔽 🔽 Write 💟 🔲 🔲 Execute 💟 ☑ 💟	
	Return	Cancel	

After selecting the *Purpose*, click the *Ownership* button to see the screen shown in Figure 16.6.

Fig. 16.6: Share Permissions

The default permissions for the share are displayed. To create a user or group, enter the desired name, then check the *Create User* box to create that user and the *Create Group* box to create the group. Check or uncheck the boxes in the *Mode* section to set the initial access permissions for the share. When finished, click the *Return* button to return to the share creation screen. Click the *Add* button to finish creating that share, which will then appear in the *Name* frame.

The *Delete* button can be used to remove the share highlighted in the *Name* frame. To edit a share, highlight it, make the change, then press the *Update* button.

When finished making shares, click the *Next* button to advance to the screen shown in Figure 16.7.

Wizard									
	Console messages:	i							
	Root E-mail:		i						
	From email:	root@freenas.local	i						
	Outgoing mail server:		i						
	Port to connect to:	25	i						
	TLS/SSL:	Plain	i						
	Use SMTP Authentication:								
	Username:		i						
	Password:								
	Password confirmation:		i						
	Previous Send Test Mail	Next Exit							

Fig. 16.7: Miscellaneous Settings

This screen can be used to configure these settings:

- **Console messages:** check this box if you would like to view system messages at the bottom of the graphical administrative interface. This can be handy when troubleshooting a service that will not start. When using the console message view, if you click the console messages area, it will pop-up as a window, allowing you to scroll through the output and to copy its contents.
- **Root E-mail:** FreeNAS[®] provides an "Alert" icon in the upper right corner to provide a visual indication of events that warrant administrative attention. The alert system automatically emails the *root* user account whenever an alert is issued. **It is important** to enter the email address of the person to receive these alerts and other administrative emails. The rest of the email settings in this screen should also be reviewed and edited as necessary. Before leaving this screen, click the "Send Test Mail" button to ensure that email notifications are working correctly.
- From email: the from email address to use when sending email notifications.
- Outgoing mail server: hostname or IP address of SMTP server.
- Port to connect to: port number used by the SMTP server.
- TLS/SSL: encryption type used by the SMTP server.
- Use SMTP Authentication: check this box if the SMTP server requires authentication.
- Username: enter the username if the SMTP server requires authentication.
- **Password:** enter the password if the SMTP server requires authentication.

When finished, click *Next*. A message will indicate that the wizard is ready to perform all of the saved actions. To make changes, click the *Return to Wizard* button to review your edits. If you click the *Exit without saving* button, none of your selections will be saved. To save your edits, click the *Confirm* button. A status bar will indicate when the wizard has completed applying the new settings.

In addition to the settings that you specify, the wizard will automatically enable *S.M.A.R.T. Tests* (page 110), create a boot environment, and add the new boot environment to the boot menu. If you also wish to save a backup of the configuration database to the system being used to access the administrative graphical interface, go to *System* \rightarrow *General*, click the *Save*

Config button, and browse to the directory where the configuration will be saved. **Always back up your configuration after making any configuration changes**.

The rest of this Guide describes the FreeNAS[®] graphical interface in more detail. The layout of this Guide follows the order of the menu items in the tree located in the left frame of the graphical interface.

Note: It is important to use the GUI (or the Console Setup menu) for all configuration changes. FreeNAS[®] uses a configuration database to store its settings. While it is possible to use the command line to modify your configuration, changes made at the command line **are not** written to the configuration database. This means that any changes made at the command line will not persist after a reboot and will be overwritten by the values in the configuration database during an upgrade.

CHAPTER SEVENTEEN

DISPLAY SYSTEM PROCESSES

Clicking *Display System Processes* opens a screen showing the output of top(1) (https://www.freebsd.org/cgi/man.cgi?query=top). An example is shown in Figure 17.1.

Running Processes 🛞												
last pid: 4533; 21 processes: 1					¥, 0.04	4, 0.00	u	0 0+01:	17:36	06:26:29		
Mem: 103M Active, 118M Inact, 224M Wired, 3220K Cache, 152M Buf, 7375M Free ARC: 2543K Total, 1052K MFU, 1126K MRU, 16K Anon, 90K Header, 258K Other												
Swap: 8192M Total, 8192M Free												
PID USERNAME	THR	PRI	NICE	SIZE	RES	STATE	С	TIME	WCPU	COMMAND		
2014 root	6	20	θ	382M	138M	usem	3	0:07	0.00%	python2.7		
2586 root	1	52	0			ttyin	3	0:01		python2.7		
3942 root	7	20				uwait		0:00		collectd		
1742 root	1	20		22216K	3852K	select	2	0:00	0.00%	ntpd		
3387 www	1	20		26028K	5540K	kgread	1	0:00	0.00%	nginx		
1557 root	1	20	0	12044K	1724K	select	2	0:00	0.00%	syslogd		
2200 root	1	52	0	14128K	1808K	nanslp	1	0:00	0.00%	cron		
2442 root	1	20		14128K	1852K	select		0:00	0.00%	rpcbind		
1290 root	1	20	0	10376K	4400K	select	3	0:00				
2088 root	1	20	0	26028K	5028K	pause	2	0:00	0.00%	nginx		
4533 root	1	20	0	16556K	2184K	CPU3	3	0:00	0.00%	top		
2591 root	1	52	0	12044K	1620K	ttyin	3	0:00		getty		
2446 root	1	23	0	12040K		select	0	0:00		mountd		
2587 root	1	52	0	12044K	1620K	ttyin	θ	0:00	0.00%	getty		
2589 root	1	52	0	12044K	1620K	ttvin	θ	0:00	0.00%	getty		
2593 root	1	52	0	12044K	1620K	ttyin	0	0:00		getty		
2588 root	1			12044K		ttyin	3	0:00		getty		
2592 root	1			12044K		ttvin	2	0:00		getty		

Fig. 17.1: System Processes Running on FreeNAS®

The display automatically refreshes itself. The display is read-only.

CHAPTER EIGHTEEN

SHELL

Beginning with version 8.2.0, the FreeNAS[®] GUI provides a web shell, making it convenient to run command line tools from the web browser as the *root* user. The link to Shell is the fourth entry from the bottom of the menu tree. In Figure 18.1, the link has been clicked and Shell is open.

Shell			
[root@fi	reenas	5 ∼]#	
Paste 80x	c25 🔻		
	x25		
	x30 x50		
	2x25		
	2x43		
132	2x50		

Fig. 18.1: Web Shell

The prompt indicates that the current user is *root*, the hostname is *freenas*, and the current working directory is ~ (*root*'s home directory).

Note: The default shell for a new install of FreeNAS[®] is zsh. FreeNAS[®] systems that are upgraded from an earlier version will continue to use csh as the default shell. The default shell can be changed by going to $Account \rightarrow Users$. Select the desired user and click *Modify User*. Choose the desired shell from the *Shell* drop-down.

To change the size of the shell, click the *80x25* drop-down menu and select a different size.

To copy text from shell, highlight the text, right-click, and select *Copy* from the right-click menu. To paste into the shell, click the *Paste* button, paste the text into the box that opens, and click the *OK* button to complete the paste operation.

Shell provides a history of commands used. Use the arrow keys to see previously entered commands and press Enter to repeat the command. The keys Home, End, and Delete are also supported in the shell. The shell also provides tab completion. Type a few letters and press tab to complete a command name or filename in the current directory. Type exit to leave the session.

Using the shell prevents access to any of the other GUI menus. To have access to a prompt while using the GUI menus, use *tmux* (page 313) instead as it supports multiple shell sessions and the detachment and reattachment of sessions.

Note: Not all of Shell's features render correctly in Chrome. Firefox is the recommended browser for using Shell.

Most FreeBSD command line utilities are available in the *Shell*. Additional troubleshooting utilities that are provided by FreeNAS[®] are described in *Command Line Utilities* (page 300).

CHAPTER NINETEEN

LOG OUT

Click the Log Out entry in the ${\rm FreeNAS}^{\circledast}$ GUI to log out.

The screen changes back to log in screen shown in Figure 19.1

Welcome to	FreeNAS®
Username:	
Password:	
Log In	🚺 systems
(New Web Interface

Fig. 19.1: Log in to FreeNAS®

CHAPTER TWENTY

REBOOT

Clicking the *Reboot* entry in the tree shows the warning message in Figure 20.1. The browser screen color changes to red to indicate that this option will negatively impact current users of the FreeNAS[®] system.

🔤 Display System Processes	
🚾 Shell	Reboot
🗶 Log Out	
Ne Reboot	Warning!
Shutdown	You are about to REBOOT the system, what would you like to do?

Fig. 20.1: Reboot Warning Message

An additional warning message appears when a restart is attempted on a system with a scrub or resilver in progress. In this case, it is recommended to *Cancel* the reboot request and to periodically run <code>zpool status</code> from Shell until it is verified that the scrub or resilver process is complete. Once complete, the reboot request can be reissued.

Click the *Cancel* button to cancel the reboot request. Otherwise, click the *Reboot* button to reboot the system. Rebooting the system disconnects all clients, including the web administration GUI. The URL in the web browser changes to add */sys-tem/reboot/* to the end of the IP address. Wait a few minutes for the system to boot, then use the back button in the browser to return to the IP address of the FreeNAS[®] system. The GUI login screen appears after a successful reboot. If the login screen does not appear, using a monitor and keyboard to physically access the FreeNAS[®] system is required to determine the problem that is preventing the system from resuming normal operation.

CHAPTER TWENTYONE

SHUTDOWN

Clicking the *Shutdown* entry in the tree opens the warning message shown in Figure 21.1. The browser window color changes to red to indicate that this command will negatively impact current users of the FreeNAS[®] system.

Display System Processes	
🚾 Shell	Shutdown 8
💥 Log Out	
兴 Reboot	Warning!
Shutdown	You are about to SHUTDOWN the system, what would you like to do?

Fig. 21.1: Shutdown Warning Message

If a scrub or resilver is running, a warning is shown. Clicking *Cancel* is recommended. zpool status can be run from the *Shell* (page 289) to watch for the scrub or resilver to complete. Then the system can be shut down normally.

Confirm the command and click *Shutdown* to shutdown the system. Shutting down the system disconnects all clients, including the web administration GUI. Physical access to the FreeNAS[®] system is required to turn it back on.

CHAPTER TWENTYTWO

SUPPORT ICON

The *Support* icon, the third icon from the left in the top menubar, provides a shortcut to *System* \rightarrow *Support*. This screen can be used to create a support ticket. Refer to *Support* (page 92) for detailed usage instructions.

CHAPTER TWENTYTHREE

GUIDE

The Documentation icon, the second icon from the left in the top menubar, links to the online documentation (this guide).

CHAPTER TWENTYFOUR

ALERT

The FreeNAS[®] alert system provides a visual warning of any conditions that require administrative attention. The *Alert* button in the far right corner flashes red when there is an outstanding alert. In the example alert shown in Figure 24.1, the system is warning that the S.M.A.R.T. service is not running.

ıg	Wizard	Support	Guide) Alert
Ale	ert System			36

🔹 📝 WARNING: April 18, 2016, 5:49 a.m. - smartd is not running.



Informational messages have a green *OK*, warning messages flash yellow, and messages requiring attention are listed as a red *CRITICAL*. CRITICAL messages are also emailed to the root user account. To remove the flashing alert for a message, deselect the option next to it.

Behind the scenes, an alert daemon checks for various alert conditions, such as volume and disk status, and writes the current conditions to /var/tmp/alert. The daemon retrieves the current alert status every minute and changes the solid green alert icon to flashing red when a new alert is detected.

Current alerts are viewed from the Shell option of the Console Setup Menu (Figure 3.1) or from the Web Shell (Figure 18.1) by running alertcli.py.

Some of the conditions that trigger an alert include:

- used space on a volume, dataset, or zvol goes over 80%; the alert goes red at 95%
- new OpenZFS feature flags are available for the pool; this alert can be unchecked if a pool upgrade is not desired at present
- a new update is available
- the system reboots itself
- · non-optimal multipath states are detected
- ZFS pool status changes from HEALTHY
- a S.M.A.R.T. error occurs
- the system dataset does not reside on the boot pool
- the system is unable to bind to the *WebGUI IPv4 Address* set in *System* ightarrow *General*
- the system can not find an IP address configured on an iSCSI portal
- the NTP server cannot be contacted
- a periodic snapshot or replication task fails
- a VMware login or a VMware-Snapshot (page 167) task fails

- deleting a VMware snapshot fails
- a Certificate Authority or certificate is invalid or malformed
- an update failed, or the system needs to reboot to complete a successful update
- a re-key operation fails on an encrypted pool
- LDAP failed to bind to the domain
- any member interfaces of a lagg interface are not active
- the status of an Avago MegaRAID SAS controller has changed; mfiutil(8) (https://www.freebsd.org/cgi/man.cgi?query=mfiutil) is included for managing these devices
- a scrub is paused

CHAPTER TWENTYFIVE

SUPPORT RESOURCES

FreeNAS[®] has a large installation base and an active user community. This means that many usage questions have already been answered and the details are available on the Internet. If an issue occurs while using FreeNAS[®], it can be helpful to spend a few moments searching the Internet for the word *FreeNAS* with some keywords that describe the error message or the function that is being implemented.

This section discusses resources available to FreeNAS[®] users:

- Website and Social Media (page 298)
- Forums (page 298)
- IRC (page 299)
- Videos (page 299)
- Professional Support (page 299)

25.1 Website and Social Media

The FreeNAS® website (http://www.freenas.org/) contains links to all of the available documentation, support, and social media resources. Major announcements are also posted to the main page.

Users are welcome to network on the FreeNAS[®] social media sites:

- LinkedIn (https://www.linkedin.com/groups/3903140/profile)
- Google+ (https://plus.google.com/110373675402281849911/posts)
- Facebook FreeNAS Community (https://www.facebook.com/freenascommunity)
- Facebook FreeNAS Consortium (please request to be added) (https://www.facebook.com/groups/1707686686200221)
- Twitter (https://twitter.com/freenas)

25.2 Forums

The FreeNAS Forums (https://forums.freenas.org/index.php) are an active online resource where people can ask questions, receive help, and share findings with other FreeNAS[®] users. New users are encouraged to post a brief message about themselves and how they use FreeNAS[®] in the Introductions (https://forums.freenas.org/index.php?forums/introductions.25/) forum.

The Resources (https://forums.freenas.org/index.php?resources/) section contains categorized, user-contributed guides on many aspects of building and using FreeNAS[®] systems.

Language-specific categories are available under International.

- Chinese (https://forums.freenas.org/index.php?forums/chinese-%E4%B8%AD%E6%96%87.60/)
- Dutch Nederlands (https://forums.freenas.org/index.php?forums/dutch-nederlands.35/)
- French Francais (https://forums.freenas.org/index.php?forums/french-francais.29/)

- German Deutsch (https://forums.freenas.org/index.php?forums/german-deutsch.31/)
- · Italian Italiano (https://forums.freenas.org/index.php?forums/italian-italiano.30/)
- Portuguese Português (https://forums.freenas.org/index.php?forums/portuguese-portugu%C3%AAs.44/)
- Romanian Română (https://forums.freenas.org/index.php?forums/romanian-rom%C3%A2n%C4%83.53/)
- Russian Русский (https://forums.freenas.org/index.php?forums/russian-%D0%A0%D1%83%D1%81%D0%BA%D0%B
- Spanish Español (https://forums.freenas.org/index.php?forums/spanish-espa%C3%B1ol.33/)
- Swedish Svenske (https://forums.freenas.org/index.php?forums/swedish-svenske.51/)
- Turkish Türkçe (https://forums.freenas.org/index.php?forums/turkish-t%C3%BCrk%C3%A7e.36/)

To join the forums, create an account with the Sign Up Now! link.

Before asking a question on the forums, check the Resources (https://forums.freenas.org/index.php?resources/) to see if the information is already there. See the Forum Rules (https://forums.freenas.org/index.php?threads/updated-forum-rules-4-11-17.45124/) for guidelines on posting your hardware information and how to ask questions that will get a response.

25.3 IRC

To ask a question in real time, use the *#freenas* channel on IRC Freenode (http://freenode.net/). Depending on the time of day and your time zone, FreeNAS[®] developers or other users may be available to provide assistance. If no one answers right away, remain on the channel, as other users tend to read the channel history to answer questions as time permits.

Typically, an IRC client (https://en.wikipedia.org/wiki/Comparison_of_Internet_Relay_Chat_clients) is used to access the *#freenas* IRC channel. Alternately, use webchat (http://webchat.freenode.net/?channels=freenas) from a web browser.

To get the most out of the IRC channel, keep these points in mind:

- Do not ask "Can anyone help me?". Just ask the question.
- Do not ask a question and then leave. Users who know the answer cannot help you if you disappear.
- If no one answers, the question may be difficult to answer or it has been asked before. Research other resources while waiting for the question to be answered.
- Do not post error messages in the channel. Instead, use a pasting service such as pastebin (https://pastebin.com/) and paste the resulting URL into the IRC discussion.

25.4 Videos

A series of instructional videos are available for FreeNAS[®]:

- Install Murmur (Mumble server) on FreeNAS/FreeBSD (https://www.youtube.com/watch?v=aAeZRNfarJc)
- FreeNAS® 9.10 Certificate Authority & SSL Certificates (https://www.youtube.com/watch?v=OT1Le5VQlc0)
- How to Update FreeNAS® 9.10 (https://www.youtube.com/watch?v=2nvb90AhgL8)
- FreeNAS® 9.10 LAGG & VLAN Overview (https://www.youtube.com/watch?v=wqSH_uQSArQ)
- FreeNAS 9.10 and Samba (SMB) Permissions (https://www.youtube.com/watch?v=RxggaE935PM)
- FreeNAS® 11 What's New (https://www.youtube.com/watch?v=-uJ_7eG88zk)
- FreeNAS® 11 How to Install (https://www.youtube.com/watch?v=R3f-Sr6y-c4)

25.5 Professional Support

In addition to free community resources, support might be available in your area through third-party consultants. Submit a support inquiry using the form at https://www.ixsystems.com/freenas-commercial-support/.

COMMAND LINE UTILITIES

Several command line utilities which are provided with FreeNAS[®] are demonstrated in this section. These utilities are used for benchmarking and performance testing:

- *lperf* (page 300): used for measuring maximum TCP and UDP bandwidth performance
- Netperf (page 303): a tool for measuring network performance
- *IOzone* (page 304): filesystem benchmark utility used to perform a broad filesystem analysis
- arcstat (page 306): used to gather ZFS ARC statistics

These utilities are specific to RAID controllers:

- *tw_cli* (page 311):_used to monitor and maintain 3ware RAID controllers
- MegaCli (page 312): used to configure and manage Broadcom MegaRAID SAS family of RAID controllers

This section also describes these utilities:

- *freenas-debug* (page 313): the backend used to dump FreeNAS[®] debugging information
- *tmux* (page 313): a terminal multiplexer similar to GNU screen
- Dmidecode (page 314): reports information about system hardware as described in the system's BIOS

26.1 Iperf

Iperf is a utility for measuring maximum TCP and UDP bandwidth performance. It can be used to chart network throughput over time. For example, it can be used to test the speed of different types of shares to determine which type best performs on the network.

FreeNAS[®] includes the Iperf server. To perform network testing, install an Iperf client on a desktop system that has network access to the FreeNAS[®] system. This section demonstrates how to use the xjperf GUI client (https://code.google.com/archive/p/xjperf/downloads) as it works on Windows, macOS, Linux, and BSD systems.

Since this client is Java-based, the appropriate JRE (http://www.oracle.com/technetwork/java/javase/downloads/index.html) must be installed on the client computer.

Linux and BSD users can install the Iperf package using the package management system for their operating system.

To start xjperf on Windows: unzip the downloaded file, start Command Prompt in Run as administrator mode, cd to the unzipped folder, and run jperf.bat.

To start xjperf on macOS, Linux, or BSD, unzip the downloaded file, cd to the unzipped directory, type chmod u+x jperf. sh, and run ./jperf.sh.

Once the client is ready, start the Iperf server on FreeNAS[®].

Note: Beginning with FreeNAS[®] version 11.1, both iperf2 (https://sourceforge.net/projects/iperf2/) and iperf3 (http://software.es.net/iperf/) are pre-installed. To use iperf2, use iperf1. To use iperf3, instead type iperf3. The examples below are for iperf2.

To see the available server options, open Shell and type:

iperf --help | more

or:

iperf3 --help | more

For example, to perform a TCP test and start the server in daemon mode (to get the prompt back), type:

```
iperf -sD
------
Server listening on TCP port 5001
TCP window size: 64.0 KByte (default)
------
Running Iperf Server as a daemon
The Iperf daemon process ID: 4842
```

Note: The daemon process stops when *Shell* (page 289) closes. Set up the environment, for example, shares configured and started, **before** starting the Iperf process.

From the desktop, open the client. Enter the IP of address of the FreeNAS[®] system, specify the running time for the test under *Application layer options* \rightarrow *Transmit* (the default test time is 10 seconds), and click the *Run Iperf*! button. Figure 26.1 shows an example of the client running on a Windows system while an SFTP transfer is occurring on the network.

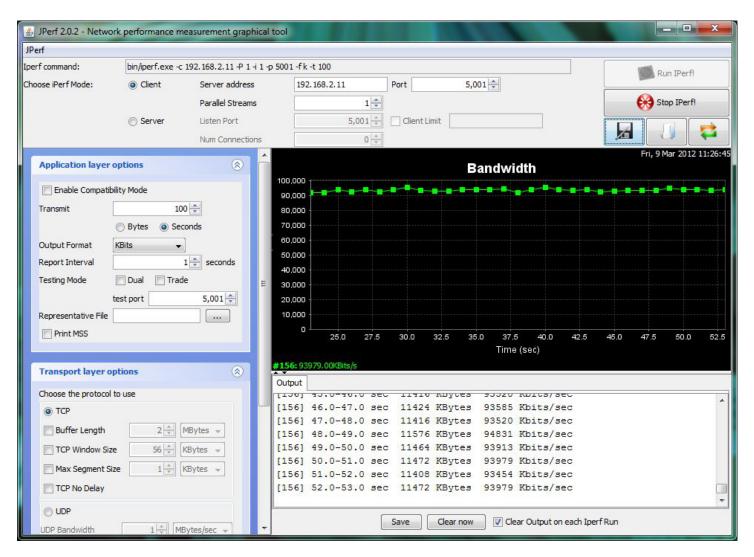


Fig. 26.1: Viewing Bandwidth Statistics Using xjperf

Depending upon the traffic being tested, for example, the type of share running on the network, UDP may need to be tested instead of TCP. To start the Iperf server in UDP mode, use iperf -sDu as the **u** specifies UDP; the startup message should indicate that the server is listening for UDP datagrams. If unsure whether the traffic to be tested is UDP or TCP, run this command to determine which services are running on the FreeNAS[®] system:

sockstat	-4 mor	re				
USER	COMMAND	PID	FD	PROTO	LOCAL ADDRESS	FOREIGN ADDRESS
root	iperf	4870	6	udp4	*:5001	*:*
root	iperf	4842	6	tcp4	*:5001	*:*
www	nginx	4827	3	tcp4	127.0.0.1:15956	127.0.0.1:9042
WWW	nginx	4827	5	tcp4	192.168.2.11:80	192.168.2.26:56964
WWW	nginx	4827	7	tcp4	*:80	*:*
root	sshd	3852	5	tcp4	*:22	*:*
root	python	2503	5	udp4	* • *	*:*
root	mountd	2363	7	udp4	*:812	*:*
root	mountd	2363	8	tcp4	*:812	*:*
root	rpcbind	2359	9	udp4	*:111	*:*
root	rpcbind	2359	10	udp4	*:886	*:*
root	rpcbind	2359	11	tcp4	*:111	*:*
root	nginx	2044	7	tcp4	*:80	*:*
root	python	2029	3	udp4	* • *	*:*
root	python	2029	4	tcp4	127.0.0.1:9042	*:*

root	python	2029	7	tcp4	127.0.0.1:9042 127.0.0.1:15956
root	ntpd	1548	20	udp4	*:123 *:*
root	ntpd	1548	22	udp4	192.168.2.11:123*:*
root	ntpd	1548	25	udp4	127.0.0.1:123 *:*
root	syslogd	1089	6	udp4	127.0.0.1:514 *:*

When testing is finished, either type killall iperf or close Shell to terminate the lperf server process.

26.2 Netperf

Netperf is a benchmarking utility that can be used to measure the performance of unidirectional throughput and end-to-end latency.

Before using the netperf command, start its server process with this command:

```
netserver
Starting netserver with host 'IN(6)ADDR_ANY' port '12865' and family AF_UNSPEC
```

The following command displays the available options for performing tests with the netperf command. The Netperf Manual (https://hewlettpackard.github.io/netperf/) describes each option in more detail and explains how to perform many types of tests. It is the best reference for understanding how each test works and how to interpret the results. When you are finished with the tests, type killall netserver to stop the server process.

```
netperf -h |more
Usage: netperf [global options] -- [test options]
Global options:
    -a send, recv
                      Set the local send, recv buffer alignment
    -A send, recv
                     Set the remote send, recv buffer alignment
    -B brandstr
                      Specify a string to be emitted with brief output
                     Report local CPU usage
    -c [cpu_rate]
    -C [cpu_rate]
                      Report remote CPU usage
    -d
                       Increase debugging output
    -D [secs, units] * Display interim results at least every secs seconds
                      using units as the initial guess for units per second
    -f G|M|K|q|m|k
                      Set the output units
    -F fill_file
                      Pre-fill buffers with data from fill_file
    -h
                      Display this text
    -H name|ip,fam * Specify the target machine and/or local ip and family
    -i max,min
                      Specify the max and min number of iterations (15,1)
    -I lvl[,intvl]
                      Specify confidence level (95 or 99) (99)
                      and confidence interval in percentage (10)
    -i
                      Keep additional timing statistics
    -l testlen
                      Specify test duration (>0 secs) (<0 bytes|trans)
    -L name|ip,fam * Specify the local ip|name and address family
    -o send, recv
                      Set the local send, recv buffer offsets
    -0 send, recv
                      Set the remote send, recv buffer offset
    -n numcpu
                      Set the number of processors for CPU util
    -N
                      Establish no control connection, do 'send' side only
    -p port,lport*
                      Specify netserver port number and/or local port
    -P 0|1
                      Don't/Do display test headers
    -r
                      Allow confidence to be hit on result only
                      Wait seconds between test setup and test start
    -s seconds
    -S
                      Set SO_KEEPALIVE on the data connection
                      Specify test to perform
    -t testname
    -T lcpu,rcpu
                      Request netperf/netserver be bound to local/remote cpu
    -v verbosity
                      Specify the verbosity level
    -W send, recv
                      Set the number of send, recv buffers
    -v level
                      Set the verbosity level (default 1, min 0)
    -V
                      Display the netperf version and exit
```

For those options taking two parms, at least one must be specified; specifying one value without a comma will set both parms to that value, specifying a value with a leading comma will set just the second parm, a value with a trailing comma will set just the first. To set each parm to unique values, specify both and separate them with a comma.

For these options taking two parms, specifying one value with no comma will only set the first parms and will leave the second at the default value. To set the second value it must be preceded with a comma or be a comma-separated pair. This is to retain previous netperf behavior.

26.3 IOzone

IOzone is a disk and filesystem benchmarking tool. It can be used to test file I/O performance for the following operations: read, write, re-read, re-write, read backwards, read strided, fread, fwrite, random read, pread, mmap, aio_read, and aio_write.

FreeNAS[®] ships with IOzone, meaning that it can be run from Shell. When using IOzone on FreeNAS[®], cd to a directory in a volume that you have permission to write to, otherwise an error about being unable to write the temporary file will occur.

Before using IOzone, read through the IOzone documentation PDF (http://www.iozone.org/docs/IOzone_msword_98.pdf) as it describes the tests, the many command line switches, and how to interpret the results.

These resources provide good starting points on which tests to run, when to run them, and how to interpret the results:

- How To Measure Linux Filesystem I/O Performance With iozone (https://www.cyberciti.biz/tips/linux-filesystembenchmarking-with-iozone.html)
- Analyzing NFS Client Performance with IOzone (http://www.iozone.org/docs/NFSClientPerf_revised.pdf)
- 10 iozone Examples for Disk I/O Performance Measurement on Linux (https://www.thegeekstuff.com/2011/05/iozoneexamples)

Type the following command to receive a summary of the available switches. As you can see from the number of options, IOzone is comprehensive so it can take some time to learn how to use the tests effectively.

Starting with version 9.2.1, FreeNAS[®] enables compression on newly created ZFS pools by default. Since IOzone creates test data that is compressible, this can skew test results. To configure IOzone to generate incompressible test data, include the options **-+w 1 -+y 1 -+C 1**.

Alternatively, consider temporarily disabling compression on the ZFS pool or dataset when running IOzone benchmarks.

Note: If a visual representation of the collected data is preferred, scripts are available to render IOzone's output in Gnuplot (http://www.gnuplot.info/).

```
iozone -h | more
iozone: help mode
Usage: iozone[-s filesize_Kb] [-r record_size_Kb] [-f [path]filename] [-h]
             [-i test] [-E] [-p] [-a] [-A] [-z] [-Z] [-m] [-M] [-t children]
             [-1 min_number_procs] [-u max_number_procs] [-v] [-R] [-x] [-o]
             [-d microseconds] [-F path1 path2...] [-V pattern] [-j stride]
             [-T] [-C] [-B] [-D] [-G] [-I] [-H depth] [-k depth] [-U mount_point]
             [-S cache_size] [-O] [-L cacheline_size] [-K] [-g maxfilesize_Kb]
             [-n minfilesize_Kb] [-N] [-Q] [-P start_cpu] [-e] [-c] [-b Excel.xls]
             [-J milliseconds] [-X write_telemetry_filename] [-W] [-W]
             [-Y read_telemetry_filename] [-y minrecsize_Kb] [-q maxrecsize_Kb]
             [-+u] [-+m cluster_filename] [-+d] [-+x multiplier] [-+p # ]
             [-+r] [-+t] [-+X] [-+Z] [-+w percent dedupable] [-+y percent_interior_dedup]
             [-+C percent_dedup_within]
         -a Auto mode
         -A Auto2 mode
         -b Filename Create Excel worksheet file
         -B Use mmap() files
         -c Include close in the timing calculations
```

```
-C Show bytes transferred by each child in throughput testing
-d # Microsecond delay out of barrier
-D Use msync(MS_ASYNC) on mmap files
-e Include flush (fsync, fflush) in the timing calculations
-E Run extension tests
-f filename to use
   filenames for each process/thread in throughput test
-F
-g # Set maximum file size (in Kbytes) for auto mode (or #m or #g)
-G Use msync(MS_SYNC) on mmap files
-h help
-H # Use POSIX async I/O with # async operations
-i # Test to run (0=write/rewrite, 1=read/re-read, 2=random-read/write
     3=Read-backwards, 4=Re-write-record, 5=stride-read, 6=fwrite/re-fwrite
     7=fread/Re-fread, 8=random_mix, 9=pwrite/Re-pwrite, 10=pread/Re-pread
     11=pwritev/Re-pwritev, 12=preadv/Re-preadv)
-I Use VxFS VX_DIRECT, O_DIRECT, or O_DIRECTIO for all file operations
-j # Set stride of file accesses to (# * record size)
-J # milliseconds of compute cycle before each I/O operation
-k # Use POSIX async I/O (no bcopy) with # async operations
-K Create jitter in the access pattern for readers
-1 # Lower limit on number of processes to run
-L # Set processor cache line size to value (in bytes)
-m Use multiple buffers
-M Report uname -a output
-n # Set minimum file size (in Kbytes) for auto mode (or #m or #g)
-N Report results in microseconds per operation
-o Writes are synch (O_SYNC)
-O Give results in ops/sec.
-p Purge on
-P # Bind processes/threads to processors, starting with this cpu
-q # Set maximum record size (in Kbytes) for auto mode (or #m or #g)
-Q Create offset/latency files
-r # record size in Kb
  or -r #k .. size in Kb
  or -r #m .. size in Mb
  or -r #g .. size in Gb
-R Generate Excel report
-s # file size in Kb
  or -s #k .. size in Kb
  or -s #m .. size in Mb
  or -s #g .. size in Gb
-S # Set processor cache size to value (in Kbytes)
-t # Number of threads or processes to use in throughput test
-T Use POSIX pthreads for throughput tests
-u # Upper limit on number of processes to run
-U Mount point to remount between tests
-v version information
-V # Verify data pattern write/read
-w Do not unlink temporary file
-W Lock file when reading or writing
-x Turn off stone-walling
-X filename Write telemetry file. Contains lines with (offset reclen compute_time) in ascii
-y # Set minimum record size (in Kbytes) for auto mode (or #m or #g)
-Y filename Read telemetry file. Contains lines with (offset reclen compute_time) in ascii
-z Used in conjunction with -a to test all possible record sizes
-Z Enable mixing of mmap I/O and file I/O
-+E Use existing non-Iozone file for read-only testing
-+K Sony special. Manual control of test 8.
-+m Cluster_filename Enable Cluster testing
-+d File I/O diagnostic mode. (To troubleshoot a broken file I/O subsystem)
-+u Enable CPU utilization output (Experimental)
```

```
-+x # Multiplier to use for incrementing file and record sizes
-+p # Percentage of mix to be reads
-+r Enable O_RSYNC|O_SYNC for all testing.
-+t Enable network performance test. Requires -+m
-+n No retests selected.
-+k Use constant aggregate data set size.
-+q Delay in seconds between tests.
-+1 Enable record locking mode.
-+L Enable record locking mode, with shared file.
-+B Sequential mixed workload.
-+A # Enable madvise. 0 = normal, 1=random, 2=sequential 3=dontneed, 4=willneed
-+N Do not truncate existing files on sequential writes.
-+S # Dedup-able data is limited to sharing within each numerically identified file set
-+V Enable shared file. No locking.
-+X Enable short circuit mode for filesystem testing ONLY
   ALL Results are NOT valid in this mode.
-+Z Enable old data set compatibility mode. WARNING.. Published
   hacks may invalidate these results and generate bogus, high values for results.
-+w ## Percent of dedup-able data in buffers.
-+y ## Percent of dedup-able within & across files in buffers.
-+C ## Percent of dedup-able within & not across files in buffers.
-+H Hostname Hostname of the PIT server.
-+P Service Service of the PIT server.
-+z Enable latency histogram logging.
```

26.4 arcstat

Arcstat is a script that prints out ZFS ARC (https://en.wikipedia.org/wiki/Adaptive_replacement_cache) statistics. Originally it was a perl script created by Sun. That perl script was ported to FreeBSD and was then ported as a Python script for use on FreeNAS[®].

Watching ARC hits/misses and percentages will provide an indication of how well the ZFS pool is fetching from the ARC rather than using disk I/O. Ideally, there will be as many things fetching from cache as possible. Keep the load in mind while reviewing the stats. For random reads, expect a miss and having to go to disk to fetch the data. For cached reads, expect it to pull out of the cache and have a hit.

Like all cache systems, the ARC takes time to fill with data. This means that it will have a lot of misses until the pool has been in use for a while. If there continues to be lots of misses and high disk I/O on cached reads, there is cause to investigate further and tune the system.

The FreeBSD ZFS Tuning Guide (https://wiki.freebsd.org/ZFSTuningGuide) provides some suggestions for commonly tuned sysctl values. It should be noted that performance tuning is more of an art than a science and that any changes made will probably require several iterations of tune and test. Be aware that what needs to be tuned will vary depending upon the type of workload and that what works for one person's network may not benefit yours.

In particular, the value of pre-fetching depends upon the amount of memory and the type of workload, as seen in this example:

• Understanding ZFS: Prefetch (http://cuddletech.com/?p=204)

FreeNAS[®] provides two command line scripts which can be manually run from *Shell* (page 289):

- arc_summary.py: provides a summary of the statistics
- arcstat.py: used to watch the statistics in real time

The advantage of these scripts is that they can be used to provide real time (right now) information, whereas the current GUI reporting mechanism is designed to only provide graphs charted over time.

This forum post (https://forums.freenas.org/index.php?threads/benchmarking-zfs.7928/) demonstrates some examples of using these scripts with hints on how to interpret the results.

To view the help for arcstat.py:

```
arcstat.py -h
[-havxp] [-f fields] [-o file] [-s string] [interval [count]]

        -h : Print this help message
        -a : Print all possible stats
        -v : List all possible field headers and definitions
        -x : Print extended stats
        -f : Specify specific fields to print (see -v)
        -o : Redirect output to the specified file
        -s : Override default field separator with custom character or string
        -p : Disable auto-scaling of numerical fields

Examples:
    arcstat -o /tmp/a.log 2 10
    arcstat -v
    arcstat -f time,hit%,dh%,ph%,mh% 1
```

To view ARC statistics in real time, specify an interval and a count. This command will display every 1 second for a count of five.

arcstat.p	oy 1 5										
time	read	miss	miss%	dmis	dm%	pmis	pm%	mmis	mm%	arcsz	С
06:19:03	7	0	0	0	0	0	0	0	0	153M	6.6G
06:19:04	257	0	0	0	0	0	0	0	0	153M	6.6G
06:19:05	193	0	0	0	0	0	0	0	0	153M	6.6G
06:19:06	193	0	0	0	0	0	0	0	0	153M	6.6G
06:19:07	255	0	0	0	0	0	0	0	0	153M	6.6G

Table 26.1 briefly describes the columns in the output.

Column	Description
read	total ARC accesses/second
miss	ARC misses/second
miss%	ARC miss percentage
dmis	demand data misses/second
dm%	demand data miss percentage
pmis	prefetch misses per second
pm%	prefetch miss percentage
mmis	metadata misses/second
mm%	metadata miss percentage
arcsz	arc size
С	arc target size

To receive a summary of statistics, use:

arcsulli	mary.py						
System	Memory:						
	2.36%	93.40	MiB	Active,	8.95%	353.43	MiB Inact
	8.38%	330.89	MiB	Wired,	0.15%	5.90	MiB Cache
	80.16%	3.09	GiB	Free,	0.00%	0	Bytes Gap
	Real In	stalled:				4.00	GiB
	Real Av	ailable:			99.31%	3.97	GiB
	Real Ma	naged:			97.10%	3.86	GiB
	Logical	Total:				4.00	GiB
	Logical	Used:			13.93%	570.77	MiB
	Logical	Free:			86.07%	3.44	GiB
Kernel	Memory:					87.62	MiB

Da	ta:	69.91%	61.25	MiB
Те	xt:	30.09%	26.37	MiB
Kernel Me	mory Map:		3.86	GiB
	ze:		201.70	
	ee:	94.89%	3.66	GiB
	ry: (HEALTHY)			
	orage pool Version:		5000	
	lesystem Version:		5	
	mory Throttle Count:		0	
ARC Misc:	leted:		8	
	tex Misses:		0	
	ict Skips:		0	
ARC Size:	-	5.83%	170.45	MiB
		100.00%		GiB
			365.69	
	x Size (High Water):	8:1	2.86	GiB
	Breakdown:			
Re	cently Used Cache Size:	50.00%	1.43	GiB
	equently Used Cache Size:	50.00%		GiB
	Breakdown:			
El	ements Max:		5.90k	
	ements Current:	100.00%	5.90k	
	llisions:		72	
	ain Max:		1	
	ains:		23	
	accesses:		0.4.6 0.5.	954.06k
	che Hit Ratio:		946.25k	
		0.82%		
			943.00k	
	ta Demand Efficiency: CHE HITS BY CACHE LIST:	<i>э</i> ≫.∠∪る	458.77k	
	Anonymously Used:	0.34%	3.25k	
	Most Recently Used:	3.73%	35.33k	
			907.67k	
		0.00%	0	
	Most Frequently Used Ghost:	0.00%	0	
	CHE HITS BY DATA TYPE:			
	Demand Data:	48.10%	455.10k	
	Prefetch Data:	0.00%	0	
	Demand Metadata:	51.56%	487.90k	
	Prefetch Metadata:	0.34%	3.25k	
CA	CHE MISSES BY DATA TYPE:			
	Demand Data:	46.93%	3.66k	
	Prefetch Data:	0.00%	0	
	Demand Metadata:	49.76%		
	Prefetch Metadata:	3.30%	258	
	le (sysctl):			
	rn.maxusers		590	
	.kmem_size		41413754	188
	.kmem_size_scale		1	
	.kmem_size_min		0	
	.kmem_size_max		1319413	950874
	s.zfs.vol.unmap_enabled		1	
	s.zfs.vol.mode		2	
	s.zfs.sync_pass_rewrite		2	
	s.zfs.sync_pass_dont_compress		5	
	s.zfs.sync_pass_deferred_free		2 0	
	s.zfs.zio.exclude_metadata s.zfs.zio.use_uma		1	
	s.zfs.cache_flush_disable		1	
VI	0.210.0a0nc_1103n_013abie		5	

vfs.zfs.zil_replay_disable	0
vfs.zfs.version.zpl	5
vfs.zfs.version.spa	5000
vfs.zfs.version.acl	1
vfs.zfs.version.ioctl	5
vfs.zfs.debug	0
	0
vfs.zfs.super_owner	
vfs.zfs.min_auto_ashift	9
vfs.zfs.max_auto_ashift	13
vfs.zfs.vdev.write_gap_limit	4096
vfs.zfs.vdev.read_gap_limit	32768
vfs.zfs.vdev.aggregation_limit	131072
vfs.zfs.vdev.trim_max_active	64
vfs.zfs.vdev.trim_min_active	1
vfs.zfs.vdev.scrub_max_active	2
vfs.zfs.vdev.scrub_min_active	1
vfs.zfs.vdev.async_write_max_active	10
vfs.zfs.vdev.async_write_min_active	1
vfs.zfs.vdev.async_read_max_active	3
vfs.zfs.vdev.async_read_min_active	1
vfs.zfs.vdev.sync_write_max_active	10
vfs.zfs.vdev.sync_write_min_active	10
vfs.zfs.vdev.sync_read_max_active	10
vfs.zfs.vdev.sync_read_min_active	10
vfs.zfs.vdev.max_active	1000
vfs.zfs.vdev.async_write_active_max_o	dirty_percent60
vfs.zfs.vdev.async_write_active_min_o	dirty_percent30
vfs.zfs.vdev.mirror.non_rotating_see	k_inc1
vfs.zfs.vdev.mirror.non_rotating_inc	0
vfs.zfs.vdev.mirror.rotating_seek_of:	fset1048576
vfs.zfs.vdev.mirror.rotating_seek_ind	c 5
vfs.zfs.vdev.mirror.rotating_inc	0
vfs.zfs.vdev.trim_on_init	1
vfs.zfs.vdev.larger_ashift_minimal	0
vfs.zfs.vdev.bio_delete_disable	0
vfs.zfs.vdev.bio_flush_disable	0
vfs.zfs.vdev.cache.bshift	16
vfs.zfs.vdev.cache.size	0
vfs.zfs.vdev.cache.max	16384
vfs.zfs.vdev.metaslabs_per_vdev	200
vfs.zfs.vdev.trim_max_pending	10000
vfs.zfs.txg.timeout	5
vfs.zfs.trim.enabled	1
vfs.zfs.trim.max_interval	1
vfs.zfs.trim.timeout	30
vfs.zfs.trim.txg_delay	32
vfs.zfs.space_map_blksz	4096
vfs.zfs.spa_slop_shift	5
vfs.zfs.spa_asize_inflation	24
vfs.zfs.deadman_enabled	1
vfs.zfs.deadman_checktime_ms	5000
vfs.zfs.deadman_synctime_ms	1000000
vfs.zfs.recover	0
vfs.zfs.spa_load_verify_data	1
vfs.zfs.spa_load_verify_metadata	1
vfs.zfs.spa_load_verify_metadata vfs.zfs.spa_load_verify_maxinflight	10000
vfs.zfs.check_hostid	1
vfs.zfs.mg_fragmentation_threshold	1 85
vis.zis.mg_noalloc_threshold	0
vfs.zfs.condense_pct	200
vfs.zfs.metaslab.bias_enabled	1
v15.215.mcca5tab.bta5_enabted	±

vfs.zfs.metaslab.lba_weighting_enabled	1
vfs.zfs.metaslab.fragmentation_factor_e	nabled1
vfs.zfs.metaslab.preload_enabled	1
vfs.zfs.metaslab.preload_limit	3
vfs.zfs.metaslab.unload_delay	8
vfs.zfs.metaslab.load_pct	50
vfs.zfs.metaslab.min_alloc_size	33554432
vfs.zfs.metaslab.df_free_pct	4
vfs.zfs.metaslab.df_alloc_threshold	131072
vfs.zfs.metaslab.debug_unload	0
vfs.zfs.metaslab.debug_load	0
vfs.zfs.metaslab.fragmentation_threshol	d70
vfs.zfs.metaslab.gang_bang	16777217
vfs.zfs.free_bpobj_enabled	1
vfs.zfs.free_max_blocks	18446744073709551615
vfs.zfs.no_scrub_prefetch	0
vfs.zfs.no_scrub_io	0
vfs.zfs.resilver_min_time_ms	3000
vfs.zfs.free_min_time_ms	1000
vfs.zfs.scan_min_time_ms	1000
vfs.zfs.scan_idle	50
vfs.zfs.scrub_delay	4
vfs.zfs.resilver_delay	2
vfs.zfs.top_maxinflight	32
vfs.zfs.delay_scale	500000
vfs.zfs.delay_min_dirty_percent	60
vfs.zfs.dirty_data_sync	67108864
vfs.zfs.dirty_data_max_percent	10
vfs.zfs.dirty_data_max_max	4294967296
vfs.zfs.dirty_data_max	426512793
vfs.zfs.max_recordsize	1048576
vfs.zfs.zfetch.array_rd_sz	1048576
vfs.zfs.zfetch.max_distance	8388608
vfs.zfs.zfetch.min_sec_reap	2
vfs.zfs.zfetch.max_streams	8
vfs.zfs.prefetch_disable	1
vfs.zfs.mdcomp_disable	0
vfs.zfs.nopwrite_enabled	1
vfs.zfs.dedup.prefetch	1
vfs.zfs.l2c_only_size	0
vfs.zfs.mfu_ghost_data_lsize	0
vfs.zfs.mfu_ghost_metadata_lsize	0
vfs.zfs.mfu_ghost_size	0
vfs.zfs.mfu_data_lsize	26300416
vfs.zfs.mfu_metadata_lsize	1780736
vfs.zfs.mfu_size	29428736
vfs.zfs.mru_ghost_data_lsize	0
vfs.zfs.mru_ghost_metadata_lsize	0
vfs.zfs.mru_ghost_size	0
vfs.zfs.mru_data_lsize	122090496
vfs.zfs.mru_metadata_lsize	2235904
vfs.zfs.mru_size	139389440
vfs.zfs.anon_data_lsize	0
vfs.zfs.anon_metadata_lsize	0
vfs.zfs.anon_size	163840
vfs.zfs.l2arc_norw	1
vfs.zfs.l2arc_feed_again	1
vfs.zfs.l2arc_noprefetch	1
vfs.zfs.l2arc_feed_min_ms	200
vfs.zfs.l2arc_feed_secs	1
vfs.zfs.l2arc_headroom	2

vfs.zfs.l2arc_write_boost	8388608
vfs.zfs.l2arc_write_max	8388608
vfs.zfs.arc_meta_limit	766908416
vfs.zfs.arc_free_target	7062
vfs.zfs.arc_shrink_shift	7
vfs.zfs.arc_average_blocksize	8192
vfs.zfs.arc_min	383454208
vfs.zfs.arc_max	3067633664

When reading the tunable values, 0 means no, 1 typically means yes, and any other number represents a value. To receive a brief description of a "sysctl" value, use sysctl -d. For example:

sysctl -d vfs.zfs.zio.use_uma
vfs.zfs.zio.use_uma: Use uma(9) for ZIO allocations

The ZFS tunables require a fair understanding of how ZFS works, meaning that reading man pages and searching for the meaning of acronyms is required. **Do not change a tunable's value without researching it first.** If the tunable takes a numeric value (rather than 0 for no or 1 for yes), do not make one up. Instead, research examples of beneficial values that match your workload.

If any of the ZFS tunables are changed, continue to monitor the system to determine the effect of the change. Using sysct1 at the command line to test the changes first is recommended. For example, to disable pre-fetch (i.e. change disable to 1 or yes):

```
sysctl vfs.zfs.prefetch_disable=1
vfs.zfs.prefetch_disable: 0 -> 1
```

The output will indicate the old value followed by the new value. If the change is not beneficial, change it back to the original value. If the change turns out to be beneficial, it can be made permanent by creating a *sysctl* using the instructions in *Tunables* (page 77).

26.5 tw_cli

FreeNAS[®] includes the tw_cli command line utility for providing controller, logical unit, and drive management for AMCC/3ware ATA RAID Controllers. The supported models are listed in the man pages for the twe(4) (https://www.freebsd.org/cgi/man.cgi?query=twe) and twa(4) (https://www.freebsd.org/cgi/man.cgi?query=twa) drivers.

Before using this command, read its man page (https://www.cyberciti.biz/files/tw_cli.8.html) as it describes the terminology and provides some usage examples.

When tw_cli is entered in Shell, the prompt will change, indicating interactive mode is enabled where all sorts of maintenance commands on the controller and its arrays can be run.

Alternately, you can specify one command to run. For example, to view the disks in the array:

tw_cli Unit	/c0 sho UnitTyp		Status	%RCmpl	%V/I/M	Stripe	Size(GB)	Cache	AVrfy
u0	RAID-6		OK	_	_	256K	5587.88		RiW	ON
u1	SPARE		OK	_	-	_	931.505		-	OFF
u2	RAID-10		OK	_	-	256K	1862.62		RiW	ON
VPort	Status	Unit	Size		Туре	Phy Enc	l-Slot	Model		
p8	OK	u0	931.51	GB SAS	_	/c0/e0/	slt0	SEAGATE	ST31000	640SS
p9	OK	u0	931.51	GB SAS	-	/c0/e0/	slt1	SEAGATE	ST31000	640SS
p10	OK	u0	931.51	GB SAS	-	/c0/e0/	slt2	SEAGATE	ST31000	640SS
p11	OK	u0	931.51	GB SAS	-	/c0/e0/	slt3	SEAGATE	ST31000	640SS
p12	OK	u0	931.51	GB SAS	-	/c0/e0/	slt4	SEAGATE	ST31000	640SS
p13	OK	u0	931.51	GB SAS	-	/c0/e0/	slt5	SEAGATE	ST31000	640SS

p14	OK	u0	931.51 GB SAS	-	/c0/e0/	/slt6	SEAGATE	ST31000640SS
p15	OK	u0	931.51 GB SAS	-	/c0/e0/	/slt7	SEAGATE	ST31000640SS
p16	OK	u1	931.51 GB SAS	-	/c0/e0/	/slt8	SEAGATE	ST31000640SS
p17	OK	u2	931.51 GB SATA	-	/c0/e0/	/slt9	ST31000	340NS
p18	OK	u2	931.51 GB SATA	-	/c0/e0/	/slt10	ST31000	340NS
p19	OK	u2	931.51 GB SATA	-	/c0/e0/	/slt11	ST31000	340NS
p20	OK	u2	931.51 GB SATA	-	/c0/e0/	/slt15	ST31000	340NS
Name	Online	eState	BBUReady	Status	Volt	Temp	Hours	LastCapTest
bbu	On		Yes	OK	OK	OK	212	03-Jan-2012

Or, to review the event log:

tw_cli	/c0 show events		
Ctl	Date	Severity	AEN Message
c0	[Thu Feb 23 2012 14:01:15]	INFO	Battery charging started
c0	[Thu Feb 23 2012 14:03:02]	INFO	Battery charging completed
c0	[Sat Feb 25 2012 00:02:18]	INFO	Verify started: unit=0
c0	[Sat Feb 25 2012 00:02:18]	INFO	<pre>Verify started: unit=2,subunit=0</pre>
c0	[Sat Feb 25 2012 00:02:18]	INFO	<pre>Verify started: unit=2, subunit=1</pre>
c0	[Sat Feb 25 2012 03:49:35]	INFO	<pre>Verify completed: unit=2,subunit=0</pre>
c0	[Sat Feb 25 2012 03:51:39]	INFO	<pre>Verify completed: unit=2,subunit=1</pre>
c0	[Sat Feb 25 2012 21:55:59]	INFO	Verify completed: unit=0
c0	[Thu Mar 01 2012 13:51:09]	INFO	Battery health check started
c0	[Thu Mar 01 2012 13:51:09]	INFO	Battery health check completed
c0	[Thu Mar 01 2012 13:51:09]	INFO	Battery charging started
c0	[Thu Mar 01 2012 13:53:03]	INFO	Battery charging completed
c0	[Sat Mar 03 2012 00:01:24]	INFO	Verify started: unit=0
c0	[Sat Mar 03 2012 00:01:24]	INFO	<pre>Verify started: unit=2, subunit=0</pre>
c0	[Sat Mar 03 2012 00:01:24]	INFO	<pre>Verify started: unit=2, subunit=1</pre>
c0	[Sat Mar 03 2012 04:04:27]	INFO	<pre>Verify completed: unit=2,subunit=0</pre>
c0	[Sat Mar 03 2012 04:06:25]	INFO	<pre>Verify completed: unit=2,subunit=1</pre>
c0	[Sat Mar 03 2012 16:22:05]	INFO	Verify completed: unit=0
c0	[Thu Mar 08 2012 13:41:39]	INFO	Battery charging started
c0	[Thu Mar 08 2012 13:43:42]	INFO	Battery charging completed
с0	[Sat Mar 10 2012 00:01:30]	INFO	Verify started: unit=0
c0	[Sat Mar 10 2012 00:01:30]	INFO	<pre>Verify started: unit=2, subunit=0</pre>
c0	[Sat Mar 10 2012 00:01:30]	INFO	<pre>Verify started: unit=2, subunit=1</pre>
с0	[Sat Mar 10 2012 05:06:38]	INFO	<pre>Verify completed: unit=2,subunit=0</pre>
с0	[Sat Mar 10 2012 05:08:57]	INFO	Verify completed: unit=2,subunit=1
с0	[Sat Mar 10 2012 15:58:15]	INFO	Verify completed: unit=0

If the disks added to the array do not appear in the GUI, try running this command:

tw_cli /c0 rescan

Use the drives to create units and export them to the operating system. When finished, run camcontrol rescan all and they should now be available in the FreeNAS[®] GUI.

This forum post (https://forums.freenas.org/index.php?threads/3ware-drive-monitoring.13835/) contains a handy wrapper script that will notify you of errors.

26.6 MegaCli

MegaCli is the command line interface for the Broadcom :MegaRAID SAS family of RAID controllers. FreeNAS[®] also includes the mfiutil(8) (https://www.freebsd.org/cgi/man.cgi?query=mfiutil) utility which can be used to configure and manage connected storage devices.

The MegaCli command is quite complex with several dozen options. The commands demonstrated in the Emergency Cheat Sheet (http://tools.rapidsoft.de/perc/perc-cheat-sheet.html) can get you started.

26.7 freenas-debug

The FreeNAS[®] GUI provides an option to save debugging information to a text file using System \rightarrow Advanced \rightarrow Save Debug. This debugging information is created by the freenas-debug command line utility and a copy of the information is saved to /var/tmp/fndebug.

This command can be run manually from *Shell* (page 289) to gather specific debugging information. To see a usage explanation listing all options, run the command without any options:

```
freenas-debug
Usage: /usr/local/bin/freenas-debug <options>
Where options are:
    -A Dump all debug information
    -B Dump System Configuration Database
    -C Dump SMB Configuration
    -D Dump Domain Controller Configuration
    -I Dump IPMI Configuration
    -M Dump SATA DOMs Information
    -N Dump NFS Configuration
    -S Dump SMART Information
    -T Loader Configuration Information
    -Z Remove old debug information
    -a Dump Active Directory Configuration
    -c Dump (AD|LDAP) Cache
    -e Email debug log to this comma-delimited list of email addresses
    -f Dump AFP Configuration
    -g Dump GEOM Configuration
    -h Dump Hardware Configuration
    -i Dump iSCSI Configuration
    -j Dump Jail Information
    -1 Dump LDAP Configuration
    -n Dump Network Configuration
    -s Dump SSL Configuration
    -t Dump System Information
    -v Dump Boot System File Verification Status and Inconsistencies
    -y Dump Sysctl Configuration
    -z Dump ZFS Configuration
```

Individual tests can be run alone. For example, when troubleshooting an Active Directory configuration, use:

freenas-debug -a

To collect the output of every module, use -A:

freenas-debug -A

26.8 tmux

tmux is a terminal multiplexer which enables a number of :terminals to be created, accessed, and controlled from a single :screen. tmux is an alternative to GNU screen. Similar to screen, tmux can be detached from a screen and continue running in the background, then later reattached. Unlike *Shell* (page 289), tmux allows you to have access to a command prompt while still providing access to the graphical administration screens.

To start a session, simply type tmux. As seen in Figure 26.2, a new session with a single window opens with a status line at the bottom of the screen. This line shows information on the current session and is used to enter interactive commands.

Shell	X
freenas#	
[θ] θ:bash*	"freenas.local" 09:11 03-Nov-14
Paste 80x25 V	

Fig. 26.2: tmux Session

To create a second window, press Ctrl+b then ". To close a window, type exit within the window.

tmux(1) (http://man.openbsd.org/cgi-bin/man.cgi/OpenBSD-current/man1/tmux.1?query=tmux) lists all of the key bindings and commands for interacting with tmux windows and sessions.

If Shell (page 289) is closed while tmux is running, it will detach its session. The next time Shell is open, run tmux attach to return to the previous session. To leave the tmux session entirely, type exit. If multiple windows are running, exit out of each first.

These resources provide more information about using tmux:

- A tmux Crash Course (https://robots.thoughtbot.com/a-tmux-crash-course)
- TMUX The Terminal Multiplexer (http://blog.hawkhost.com/2010/06/28/tmux-the-terminal-multiplexer/)

26.9 Dmidecode

Dmidecode reports hardware information as reported by the system BIOS. Dmidecode does not scan the hardware, it only reports what the BIOS told it to. A sample output can be seen here (http://www.nongnu.org/dmidecode/sample/dmidecode.txt).

To view the BIOS report, type the command with no arguments:

dmidecode | more

dmidecode(8) (https://linux.die.net/man/8/dmidecode) describes the supported strings and types.

26.10 Midnight Commander

Midnight Commander is a program used to manage files from the shell. Open the application by running the command mc. The arrow keys are used to navigate and select files. The function keys are used to perform operations such as renaming, editing and copying files. These resources provide more information about using mc:

- Midnight Commander wikipedia page (https://en.wikipedia.org/wiki/Midnight_Commander)
- Midnight Commander website (https://midnight-commander.org/)
- mc(1) (https://linux.die.net/man/1/mc)
- Basic Tutorial (http://linuxcommand.org/lc3_adv_mc.php)

CHAPTER TWENTYSEVEN

CONTRIBUTING TO FREENAS®

FreeNAS[®] is an open source community, relying on the input and expertise of its users to help grow and improve FreeNAS[®]. When you take time to assist the community, your contributions benefit everyone who uses FreeNAS[®].

This section describes some areas of participation to get you started. It is by no means an exhaustive list. If you have an idea that you think would benefit the FreeNAS[®] community, bring it up on one of the resources mentioned in *Support Resources* (page 298).

This section demonstrates how you can:

• Help with Translation (page 316)

27.1 Translation

Not everyone speaks English, and having a complete translation of the user interface into native languages can make FreeNAS[®] much more useful to communities around the world.

FreeNAS[®] uses Weblate (https://weblate.org/en/) to manage the translation of text shown in the FreeNAS[®] graphical administrative interface. Weblate provides an easy-to-use web-based editor and commenting system, making it possible for individuals to assist with translation or comment on existing translations.

To see the status of translations, open https://weblate.trueos.org/projects/freenas/, as shown in Figure 27.1.

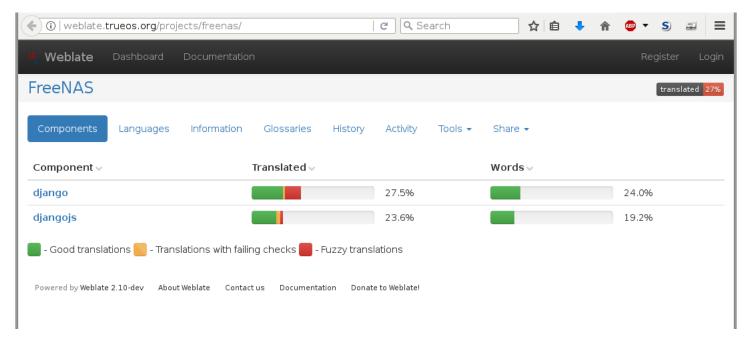


Fig. 27.1: FreeNAS[®] Translation System

To assist with translating FreeNAS[®], create an account by clicking the *Register* button. Enter the information requested, then a confirmation email will be sent. Follow the link in the email to set a password and complete the account creation. The Dashboard screen is shown after logging in:

() weblate.trueos.org	C Search	☆ 自	↓ ⋒	ABP 🔻	S	-	≡
Weblate Dashboard Documentation							€
Dashboard							
Watched translations - Search History Activity Tools -							
Choose your languages in preferences and you will get an overview here of Manage your languages Manage watched projects	f available translations for those I	anguages	in your wa	atched pr	ojects.		
Powered by Weblate 2.10-dev About Weblate Contact us Documentation Don	ate to Weblate!						
-							

Fig. 27.2: Weblate Dashboard

Click *Manage your languages* to choose languages for translation. Select languages, then click *Save*. Click the *Dashboard* link at the top of the screen to go back to the dashboard, then choose *Your languages* from the drop-down menu:

(i) weblate.trueos.	org/#your-languag	es		C Q Sea	arch	☆ 自	🕈 🏠 🕻	• • S	-	≡
Weblate Dashb	poard Document	ation								€
Dashboard										
Your languages 👻	Search History	Activity Tools								
Project 🗸	Tran	slated√		Wordsv	Review ~	Checks v	Suggestion	s√	Ð	
FreeNAS/django (Spa	anish)		64.7%	57.3%	20.0%	226	38	ø	' Translate	e
FreeNAS/django (Spa FreeNAS/djangojs (Sp			64.7% 60.1%	57.3% 36.8%	20.0% 0.7%	226 1	38 2		• Translate • Translate	_
	panish)	failing checks 🛑 - Fu	60.1%	36.8%						



Projects are a collection of text to be translated. In this example, the Django and DjangoJS projects have both been partially translated into Spanish. Click one of the entries under *Project* to help translate that project.

The Overview screen shows the current translation status along with categories of translatable strings:

FreeNAS[®] 11.2-RELEASE User Guide, Release 11.2

▶ ③ weblate.trueos.org/projects/freenas/django/es/	☆ 自 🔸 斋 🐵 ▼ S 👄
Veblate Dashboard Documentation	
eNAS / django / Spanish	translated
erview Search History Activity Statistics Files - Tools - Share -	
Translation status	
	_
Strings (2508)	64.7%
Words	57.3%
📕 - Good translations 📉 - Translations with failing checks 🛑 - Fuzzy translations	
Strings to check	Ø
All strings	2508 (14299 words)
Translated strings	1623 8203 words
Strings needing action	885 6096 words
Not translated strings	383 4119 words
Strings marked for review	502 1977 words
Strings with suggestions	38
Strings with any failing checks	226
Source and translation do not both end with a space	14
This message has more than one translation in this project	9
Source and translated strings are same	37
Source and translation do not both end with a question mark or it is not correctly spaced	0
Source and translation do not both end with a newline	0
Source and translation do not both end with an exclamation mark or it is not correctly spaced	0
Source and translation do not both end with a colon or colon is not correctly spaced	0
XML tags in translation do not match source	0
Python format string does not match source	B
Source and translation do not both start with same number of spaces	2
Source and translation do not both end with a full stop	160

Fig. 27.4: Translation Overview

Click on a category of string, like *Strings needing action*, to see the translation screen:

) 🛈 weblate.trueos.org/translate/freenas/django/es/?type=todo 🛛 😋 🔍 Search 🛣 🛧	i 🖡	ae 🔹 S	-
Weblate Dashboard Documentation			
eeNAS / django / Spanish / translate			
			X Zer
Translate 🔤	Things to ch	neck	
Source	Suggestions		
Translation Copy → + « » " "	Glossary		Ø
<u>Dirección</u> Final de red IPv6	Source	Translat	ion
	No related str glossary	ings were fo	und in
Save Suggest Commit message: Additional text to include in t	Source	Translati	on
Nearby messages Suggestions Comments Machine translation History	Source info	rmation	3
Nearby messages 11 Suggestions 1 Comments Machine translation History	Source strin	g location	

Fig. 27.5: Translate Strings

Enter translations here, clicking *Save* to save the work. The controls at the top of the screen can be used to skip forward and back in the list of strings to be translated. Click *Dashboard* at the top of the screen to return to the Dashboard.

All assistance with translations helps to benefit the FreeNAS[®] community. Thank you!

ZFS PRIMER

ZFS is an advanced, modern filesystem that was specifically designed to provide features not available in traditional UNIX filesystems. It was originally developed at Sun with the intent to open source the filesystem so that it could be ported to other operating systems. After the Oracle acquisition of Sun, some of the original ZFS engineers founded OpenZFS (http://open-zfs.org/wiki/Main_Page) to provide continued, collaborative development of the open source version. To differentiate itself from Oracle ZFS version numbers, OpenZFS uses feature flags. Feature flags are used to tag features with unique names in order to provide portability between OpenZFS implementations running on different platforms, as long as all of the feature flags enabled on the ZFS pool are supported by both platforms. FreeNAS[®] uses OpenZFS and each new version of FreeNAS[®] keeps up-to-date with the latest feature flags and OpenZFS bug fixes.

Here is an overview of the features provided by ZFS:

ZFS is a transactional, Copy-On-Write (COW) (https://en.wikipedia.org/wiki/ZFS#Copy-on-write_transactional_model) filesystem. For each write request, a copy is made of the associated disk blocks and all changes are made to the copy rather than to the original blocks. When the write is complete, all block pointers are changed to point to the new copy. This means that ZFS always writes to free space, most writes are sequential, and old versions of files are not unlinked until a complete new version has been written successfully. ZFS has direct access to disks and bundles multiple read and write requests into transactions. Most filesystems cannot do this, as they only have access to disk blocks. A transaction either completes or fails, meaning there will never be a write-hole (https://blogs.oracle.com/bonwick/raid-z) and a filesystem checker utility is not necessary. Because of the transactional design, as additional storage capacity is added, it becomes immediately available for writes. To rebalance the data, one can copy it to re-write the existing data across all available disks. As a 128-bit filesystem, the maximum filesystem or file size is 16 exabytes.

ZFS was designed to be a self-healing filesystem. As ZFS writes data, it creates a checksum for each disk block it writes. As ZFS reads data, it validates the checksum for each disk block it reads. Media errors or "bit rot" can cause data to change, and the checksum no longer matches. When ZFS identifies a disk block checksum error on a pool that is mirrored or uses RAIDZ, it replaces the corrupted data with the correct data. Since some disk blocks are rarely read, regular scrubs should be scheduled so that ZFS can read all of the data blocks to validate their checksums and correct any corrupted blocks. While multiple disks are required in order to provide redundancy and data correction, ZFS will still provide data corruption detection to a system with one disk. FreeNAS[®] automatically schedules a monthly scrub for each ZFS pool and the results of the scrub are displayed by selecting the *Volume* (page 125) and clicking *Volume Status*. Checking scrub results provides an early indication of potential disk problems.

Unlike traditional UNIX filesystems, **it is not necessary to define partition sizes when filesystems are created**. Instead, a group of disks, known as a *vdev*, are built into a ZFS *pool*. Filesystems are created from the pool as needed. As more capacity is needed, identical vdevs can be striped into the pool. In FreeNAS[®], *Volume Manager* (page 125) is used to create or extend ZFS pools. After a pool is created, it can be divided into dynamically-sized datasets or fixed-size zvols as needed. Datasets can be used to optimize storage for the type of data being stored as permissions and properties such as quotas and compression can be set on a per-dataset level. A zvol is essentially a raw, virtual block device which can be used for applications that need raw-device semantics such as iSCSI device extents.

ZFS supports real-time data compression. Compression happens when a block is written to disk, but only if the written data will benefit from compression. When a compressed block is accessed, it is automatically decompressed. Since compression happens at the block level, not the file level, it is transparent to any applications accessing the compressed data. ZFS pools created on FreeNAS[®] version 9.2.1 or later use the recommended LZ4 compression algorithm.

ZFS provides low-cost, instantaneous snapshots of the specified pool, dataset, or zvol. Due to COW, snapshots initially take no additional space. The size of a snapshot increases over time as changes to the files in the snapshot are written to disk. Snapshots can be used to provide a copy of data at the point in time the snapshot was created. When a file is deleted,

its disk blocks are added to the free list; however, the blocks for that file in any existing snapshots are not added to the free list until all referencing snapshots are removed. This makes snapshots a clever way to keep a history of files, useful for recovering an older copy of a file or a deleted file. For this reason, many administrators take snapshots often, store them for a period of time, and store them on another system. Such a strategy allows the administrator to roll the system back to a specific time. If there is a catastrophic loss, an off-site snapshot can restore the system up to the last snapshot interval, within 15 minutes of the data loss, for example. Snapshots are stored locally but can also be replicated to a remote ZFS pool. During replication, ZFS does not do a byte-for-byte copy but instead converts a snapshot into a stream of data. This design means that the ZFS pool on the receiving end does not need to be identical and can use a different RAIDZ level, volume size, or compression settings.

ZFS boot environments provide a method for recovering from a failed upgrade. In FreeNAS[®], a snapshot of the dataset the operating system resides on is automatically taken before an upgrade or a system update. This saved boot environment is automatically added to the GRUB boot loader. Should the upgrade or configuration change fail, simply reboot and select the previous boot environment from the boot menu. Users can also create their own boot environments in *System* \rightarrow *Boot* as needed, for example before making configuration changes. This way, the system can be rebooted into a snapshot of the system that did not include the new configuration changes.

ZFS provides a write cache in RAM as well as a ZFS Intent Log (ZIL (http://www.freenas.org/blog/zfs-zil-and-slogdemystified/)). The ZIL is a storage area that temporarily holds *synchronous* writes until they are written to the ZFS pool (https://pthree.org/2013/04/19/zfs-administration-appendix-a-visualizing-the-zfs-intent-log/). Adding a fast (low-latency), power-protected SSD as a SLOG (*Separate Log*) device permits much higher performance. This is a necessity for NFS over ESXi, and highly recommended for database servers or other applications that depend on synchronous writes. More detail on SLOG benefits and usage is available in these blog and forum posts:

- The ZFS ZIL and SLOG Demystified (http://www.freenas.org/blog/zfs-zil-and-slog-demystified/)
- Some insights into SLOG/ZIL with ZFS on FreeNAS® (https://forums.freenas.org/index.php?threads/some-insightsinto-slog-zil-with-zfs-on-freenas.13633/)
- ZFS Intent Log (http://nex7.blogspot.com/2013/04/zfs-intent-log.html)

Synchronous writes are relatively rare with SMB, AFP, and iSCSI, and adding a SLOG to improve performance of these protocols only makes sense in special cases. The *zilstat* utility can be run from *Shell* (page 289) to determine if the system will benefit from a SLOG. See this website (http://www.richardelling.com/Home/scripts-and-programs-1/zilstat) for usage information.

ZFS currently uses 16 GiB of space for SLOG. Larger SSDs can be installed, but the extra space will not be used. SLOG devices cannot be shared between pools. Each pool requires a separate SLOG device. Bandwidth and throughput limitations require that a SLOG device must only be used for this single purpose. Do not attempt to add other caching functions on the same SSD, or performance will suffer.

In mission-critical systems, a mirrored SLOG device is highly recommended. Mirrored SLOG devices are *required* for ZFS pools at ZFS version 19 or earlier. The ZFS pool version is checked from the *Shell* (page 289) with zpool get version poolname. A version value of - means the ZFS pool is version 5000 (also known as *Feature Flags*) or later.

ZFS provides a read cache in RAM, known as the ARC, which reduces read latency. FreeNAS[®] adds ARC stats to top(1) (https://www.freebsd.org/cgi/man.cgi?query=top) and includes the arc_summary.py and arcstat.py tools for monitoring the efficiency of the ARC. If an SSD is dedicated as a cache device, it is known as an L2ARC (http://www.brendangregg.com/blog/2008-07-22/zfs-l2arc.html). Additional read data is cached here, which can increase random read performance. L2ARC does *not* reduce the need for sufficient RAM. In fact, L2ARC needs RAM to function. If there is not enough RAM for a adequately-sized ARC, adding an L2ARC will not increase performance. Performance actually decreases in most cases, potentially causing system instability. RAM is always faster than disks, so always add as much RAM as possible before considering whether the system can benefit from an L2ARC device.

When applications perform large amounts of *random* reads on a dataset small enough to fit into L2ARC, read performance can be increased by adding a dedicated cache device. SSD cache devices only help if the active data is larger than system RAM but small enough that a significant percentage fits on the SSD. As a general rule, L2ARC should not be added to a system with less than 32 GiB of RAM, and the size of an L2ARC should not exceed ten times the amount of RAM. In some cases, it may be more efficient to have two separate pools: one on SSDs for active data, and another on hard drives for rarely used content. After adding an L2ARC device, monitor its effectiveness using tools such as arcstat. To increase the size of an existing L2ARC, stripe another cache device with it. The GUI will always stripe L2ARC, not mirror it, as the contents of L2ARC are recreated at boot. Failure of an individual SSD from an L2ARC pool will not affect the integrity of the pool, but may have

an impact on read performance, depending on the workload and the ratio of dataset size to cache size. Note that dedicated L2ARC devices cannot be shared between ZFS pools.

ZFS was designed to provide redundancy while addressing some of the inherent limitations of hardware RAID such as the write-hole and corrupt data written over time before the hardware controller provides an alert. ZFS provides three levels of redundancy, known as *RAIDZ*, where the number after the *RAIDZ* indicates how many disks per vdev can be lost without losing data. ZFS also supports mirrors, with no restrictions on the number of disks in the mirror. ZFS was designed for commodity disks so no RAID controller is needed. While ZFS can also be used with a RAID controller, it is recommended that the controller be put into JBOD mode so that ZFS has full control of the disks.

When determining the type of ZFS redundancy to use, consider whether the goal is to maximize disk space or performance:

- RAIDZ1 maximizes disk space and generally performs well when data is written and read in large chunks (128K or more).
- RAIDZ2 offers better data availability and significantly better mean time to data loss (MTTDL) than RAIDZ1.
- A mirror consumes more disk space but generally performs better with small random reads. For better performance, a mirror is strongly favored over any RAIDZ, particularly for large, uncacheable, random read loads.
- Using more than 12 disks per vdev is not recommended. The recommended number of disks per vdev is between 3 and 9. With more disks, use multiple vdevs.
- Some older ZFS documentation recommends that a certain number of disks is needed for each type of RAIDZ in order to achieve optimal performance. On systems using LZ4 compression, which is the default for FreeNAS[®] 9.2.1 and higher, this is no longer true.

These resources can also help determine the RAID configuration best suited to the specific storage requirements:

- Getting the Most out of ZFS Pools (https://forums.freenas.org/index.php?threads/getting-the-most-out-of-zfs-pools.16/)
- A Closer Look at ZFS, Vdevs and Performance (https://constantin.glez.de/2010/06/04/a-closer-look-zfs-vdevs-and-performance/)

Warning: RAID AND DISK REDUNDANCY ARE NOT A SUBSTITUTE FOR A RELIABLE BACKUP STRATEGY. BAD THINGS HAPPEN AND A GOOD BACKUP STRATEGY IS STILL REQUIRED TO PROTECT VALUABLE DATA. See *Periodic Snapshot Tasks* (page 150) and *Replication Tasks* (page 152) to use replicated ZFS snapshots as part of a backup strategy.

ZFS manages devices. When an individual drive in a mirror or RAIDZ fails and is replaced by the user, ZFS adds the replacement device to the vdev and copies redundant data to it in a process called *resilvering*. Hardware RAID controllers usually have no way of knowing which blocks were in use and must copy every block to the new device. ZFS only copies blocks that are in use, reducing the time it takes to rebuild the vdev. Resilvering is also interruptable. After an interruption, resilvering resumes where it left off rather than starting from the beginning.

While ZFS provides many benefits, there are some caveats:

- At 90% capacity, ZFS switches from performance- to space-based optimization, which has massive performance implications. For maximum write performance and to prevent problems with drive replacement, add more capacity before a pool reaches 80%. If using iSCSI, it is recommended to not let the pool go over 50% capacity to prevent fragmentation issues.
- When considering the number of disks to use per vdev, consider the size of the disks and the amount of time required for resilvering, which is the process of rebuilding the vdev. The larger the size of the vdev, the longer the resilvering time. When replacing a disk in a RAIDZ, it is possible that another disk will fail before the resilvering process completes. If the number of failed disks exceeds the number allowed per vdev for the type of RAIDZ, the data in the pool will be lost. For this reason, RAIDZ1 is not recommended for drives over 1 TiB in size.
- Using drives of equal sizes is recommended when creating a vdev. While ZFS can create a vdev using disks of differing sizes, its capacity will be limited by the size of the smallest disk.

For those new to ZFS, the Wikipedia entry on ZFS (https://en.wikipedia.org/wiki/Zfs) provides an excellent starting point to learn more about its features. These resources are also useful for reference:

• FreeBSD ZFS Tuning Guide (https://wiki.freebsd.org/ZFSTuningGuide)

- ZFS Administration Guide (https://docs.oracle.com/cd/E19253-01/819-5461/index.html)
- Becoming a ZFS Ninja (video) (https://www.youtube.com/watch?v=6_K55Ira1Cs)
- Slideshow explaining VDev, zpool, ZIL and L2ARC and other newbie mistakes! (https://forums.freenas.org/index.php?threads/slideshow-explaining-vdev-zpool-zil-and-l2arc-for-noobs.7775/)
- A Crash Course on ZFS (http://www.bsdnow.tv/tutorials/zfs)
- ZFS: The Last Word in File Systems Part 1 (video) (https://www.youtube.com/watch?v=uT2i2ryhCio)
- The Zettabyte Filesystem (https://www.youtube.com/watch?v=ptY6-K78McY)

VAAI

VMware's vStorage APIs for Array Integration, or VAAI, allows storage tasks such as large data moves to be offloaded from the virtualization hardware to the storage array. These operations are performed locally on the NAS without transferring bulk data over the network.

29.1 VAAI for iSCSI

VAAI for iSCSI supports these operations:

- Atomic Test and Set (ATS) allows multiple initiators to synchronize LUN access in a fine-grained manner rather than locking the whole LUN and preventing other hosts from accessing the same LUN simultaneously.
- *Clone Blocks (XCOPY)* copies disk blocks on the NAS. Copies occur locally rather than over the network. The operation is similar to Microsoft ODX (https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-R2-and-2012/hh831628(v=ws.11)).
- LUN Reporting allows a hypervisor to query the NAS to determine whether a LUN is using thin provisioning.
- *Stun* pauses running virtual machines when a volume runs out of space. The space issue can then be fixed and the virtual machines can continue rather than reporting write errors.
- *Threshold Warning* the system reports a warning when a configurable capacity is reached. In FreeNAS[®], this threshold can be configured at the pool level when using zvols (see Table 10.6) or at the extent level (see Table 10.11) for both file- and device-based extents. Typically, the warning is set at the pool level, unless file extents are used, in which case it must be set at the extent level.
- *Unmap* informs FreeNAS[®] that the space occupied by deleted files should be freed. Without unmap, the NAS is unaware of freed space created when the initiator deletes files. For this feature to work, the initiator must support the unmap command.
- *Zero Blocks* or *Write Same* zeros out disk regions. When allocating virtual machines with thick provisioning, the zero write is done locally, rather than over the network. This makes virtual machine creation and any other zeroing of disk regions much quicker.

CHAPTER

USING THE API

A REST (https://en.wikipedia.org/wiki/Representational_state_transfer) API is provided to be used as an alternate mechanism for remotely controlling a FreeNAS[®] system.

REST provides an easy-to-read, HTTP implementation of functions, known as resources, which are available beneath a specified base URL. Each resource is manipulated using the HTTP methods defined in **RFC 2616** (https://tools.ietf.org/html/rfc2616.html), such as GET, PUT, POST, or DELETE.

As shown in Figure 30.1, an online version of the API is available at api.freenas.org (http://api.freenas.org).

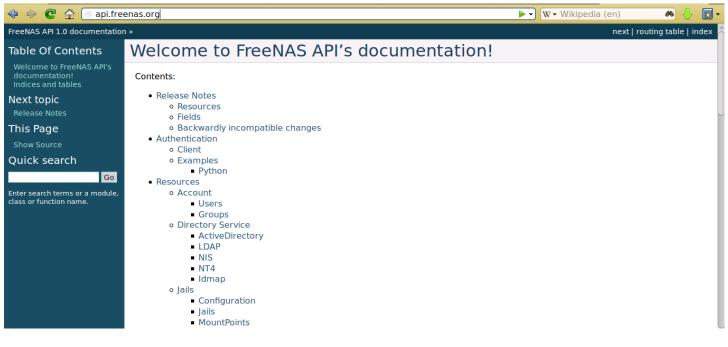


Fig. 30.1: API Documentation

The rest of this section shows code examples to illustrate the use of the API.

Note: A new API was released with FreeNAS[®] 11.1. The previous API is still present and in use because it is feature-complete. Documentation for the new API is available on the FreeNAS[®] system at the */api/docs/* URL. For example, if the FreeNAS[®] system is at IP address 192.168.1.119, enter *http://192.168.1.119/api/docs/* in a browser to see the API documentation. Work is under way to make the new API feature-complete. The new APIv2 uses WebSockets (https://developer.mozilla.org/en-US/docs/Web/API/WebSockets_API). This advanced technology makes it possible to open interactive communication sessions between web browsers and servers, allowing event-driven responses without the need to poll the server for a reply. When APIv2 is feature complete, the FreeNAS[®] documentation will include relevant examples that make use of the new API.

30.1 A Simple API Example

The api directory of the FreeNAS® github repository (https://github.com/freenas/freenas/tree/master/examples/api) contains some API usage examples. This section provides a walk-through of the newuser.py script, shown below, as it provides a simple example that creates a user.

A FreeNAS[®] system running at least version 9.2.0 is required when creating a customized script based on this example. To test the scripts directly on the FreeNAS[®] system, create a user account and select an existing volume or dataset for the user's *Home Directory*. After creating the user, start the SSH service using *Services* \rightarrow *Control Services*. That user will now be able to ssh to the IP address of the FreeNAS[®] system to create and run scripts. Alternately, scripts can be tested on any system with the required software installed as shown in the previous section.

To customize this script, copy the contents of this example into a filename that ends in .py. The text that is highlighted in red below can be modified in the new version to match the needs of the user being created. The text in black should not be changed. After saving changes, run the script by typing python scriptname.py. If all goes well, the new user account will appear in Account \rightarrow Users \rightarrow View Users in the FreeNAS[®] GUI.

Here is the example script with an explanation of the line numbers below it.

```
import json
   import requests
2
   r = requests.post(
3
     'https://freenas.mydomain/api/v1.0/account/users/',
     auth=('root', 'freenas'),
5
     headers={'Content-Type': 'application/json'},
     verify=False,
     data=json.dumps({
8
           'bsdusr_uid': '1100',
9
           'bsdusr_username': 'myuser',
10
           'bsdusr_mode': '755',
11
           'bsdusr_creategroup': 'True',
12
           'bsdusr_password': '12345',
13
           'bsdusr_shell': '/usr/local/bin/bash',
14
           'bsdusr_full_name': 'Full Name',
15
           'bsdusr_email': 'name@provider.com',
16
      })
17
    )
18
    print r.text
19
```

Where:

Lines 1-2: import the Python modules used to make HTTP requests and handle data in JSON format.

Line 4: replace *freenas.mydomain* with the *Hostname* value in *System* \rightarrow *System Information*. Note that the script will fail if the machine running it is not able to resolve that hostname. Change *https* to *http* to use HTTP rather than HTTPS to access the FreeNAS[®] system.

Line 5: replace *freenas* with the password used to access the FreeNAS[®] system.

Line 7: if you are using HTTPS and want to force validation of the SSL certificate, change False to True.

Lines 8-16: set the values for the user being created. The Users resource (http://api.freenas.org/resources/account.html#users) describes this in more detail. Allowed parameters are listed in the JSON Parameters section of that resource. Since this resource creates a FreeBSD user, the values entered must be valid for a FreeBSD user account.

Table 30.1 summarizes acceptable values. This resource uses JSON, so the boolean values are *True* or *False*.

JSON Parameter	Туре	Description
bsdusr_username	string	Enter a maximum of 32 characters. A maximum of 8 is recommended for interoperability. The username can include numerals but cannot include a space.

Table 30.1: JSON Parameters for Users Create Resource

Continued on next page

JSON Parameter	Туре	Description
bsdusr_full_name	string	This field can contain spaces and uppercase characters.
bsdusr_password	string	The password can include a mix of upper and lowercase letters, characters, and numbers.
bsdusr_uid	integer	By convention, user accounts have an ID greater than 1000 with a maxi- mum allowable value of 65,535.
bsdusr_group	integer	Specify the numeric ID of the group to create if <i>bsdusr_creategroup</i> is set to <i>False</i> .
bsdusr_creategroup	boolean	Set to <i>True</i> to create a primary group with the same numeric ID as <i>bs</i> - <i>dusr_uid</i> .
bsdusr_mode	string	Sets default numeric UNIX permissions for the home directory of the user.
bsdusr_shell	string	Specify the full path to a UNIX shell that is installed on the system.
bsdusr_password_disabl eo olean		The user is not allowed to log in when set to <i>True</i> .
bsdusr_locked	boolean	The user is not allowed to log in when set to <i>True</i> .
bsdusr_sudo	boolean	sudo is enabled for the user when set to True.
bsdusr_sshpubkey	string	Enter the contents of the SSH authorized keys file.

Table 30.1 – continued from previous page

Note: When using boolean values, JSON returns raw lowercase values but Python uses uppercase values. So use *True* or *False* in Python scripts even though the example JSON responses in the API documentation are displayed as *true* or *false*.

30.2 A More Complex Example

This section provides a walk-through of a more complex example found in the startup.py script. Use the searchbar within the API documentation to quickly locate the JSON parameters used here. This example defines a class and several methods to create a ZFS volume, create a ZFS dataset, share the dataset over CIFS, and enable the CIFS service. Responses from some methods are used as parameters in other methods. In addition to the import lines seen in the previous example, two additional Python modules are imported to provide parsing functions for command line arguments:

import argparse
import sys

2

3

Δ

5

6

7

8

9

10

11

12

13

14

15

16

17 18

19

20

It then creates a *Startup* class which is started with the hostname, username, and password provided by the user via the command line:

```
class Startup(object):
 def __init__(self, hostname, user, secret):
      self._hostname = hostname
       self._user = user
       self._secret = secret
       self._ep = 'http://%s/api/v1.0' % hostname
 def request(self, resource, method='GET', data=None):
       if data is None:
           data = ''
       r = requests.request(
           method.
           '%s/%s/' % (self._ep, resource),
           data=json.dumps(data),
           headers={'Content-Type': "application/json"},
           auth=(self._user, self._secret),
       )
       if r.ok:
           try:
               return r.json()
           except:
```

21 22

```
return r.text
raise ValueError(r)
```

A *get_disks* method is defined to get all the disks in the system as a *disk_name* response. The *create_pool* method uses this information to create a ZFS pool named *tank* which is created as a stripe. The *volume_name* and *layout* JSON parameters are described in the "Storage Volume" resource of the API documentation.:

```
def
       _get_disks(self):
           disks = self.request('storage/disk')
2
           return [disk['disk_name'] for disk in disks]
3
4
   def create_pool(self):
5
           disks = self._get_disks()
6
           self.request('storage/volume', method='POST', data={
7
               'volume_name': 'tank',
8
               'layout': [
9
                   {'vdevtype': 'stripe', 'disks': disks},
10
11
               ],
   })
12
```

The *create_dataset* method is defined which creates a dataset named MyShare:

The create_cifs_share method is used to share /mnt/tank/MyShare with guest-only access enabled. The cifs_name, cifs_path, cifs_guestonly JSON parameters, as well as the other allowable parameters, are described in the "Sharing CIFS" resource of the API documentation.:

```
1 def create_cifs_share(self):
2     self.request('sharing/cifs', method='POST', data={
3         'cifs_name': 'My Test Share',
4          'cifs_path': '/mnt/tank/MyShare',
5          'cifs_guestonly': True
6 })
```

Finally, the *service_start* method enables the CIFS service. The *srv_enable* JSON parameter is described in the Services resource.

INDEX

Symbols

802.1Q, 123

A

Add Group, 59 Add User, 62 Adding Devices to a VM, 267 AFP, 181, 221 Alert, 295 Alert Services, 85 Alerts, 85 API, 324 Apple Filing Protocol, 181, 221 arcstat, 306 Autotune, 73

В

Boot Environments, 69 Burn ISO, 18

С

CA, 87 Certificate Authority, 87 Certificates, 89 Checksum, 18 CIFS, 195, 238 Cloud Credentials, 82 Cloud Sync, 95 Compression, 135 Create Dataset, 132 Create Group, 59 Create the Docker VM, 273 Create User, 62 Creating VMs, 266 Cron Jobs, 100

D

DC, 222 DDNS, 223 Deduplication, 134 Delete Group, 60 Delete User, 64 Deleting VMs, 272 Dell PERC H330, 15 Dell PERC H730, 15 Dmidecode, 314 Docker VM, 273 Docker VM Requirements, 273 Domain Controller, 222 Download, 18 Dynamic DNS, 223

E

Email, 75 Encryption, 127 EtherChannel, 118

F

File Transfer Protocol, 225 Forums, 298 freenas-debug, 313 FTP, 225

G

Getting FreeNAS\ :sup:'®', 18 Groups, 58 Guide, 294

Η

Hardware Recommendations, 13 Highpoint RAID, 15 Hot Spares, 150

I

Install, 20 Internet Small Computer System Interface, 205 IOzone, 304 Iperf, 300 IRC, 299 iSCSI, 205 ISO, 18

J

Jails, 255

L

LACP, 118 LAGG, 118 Link Aggregation, 118 Link Layer Discovery Protocol, 230 LLDP, 230 Localize, 316 Log Out, 290

Μ

MegaCli, 312 Midnight Commander, 314 Minio, 236 Mirroring the Boot Device, 71 Multiple Boot Environments, 69

Ν

Netdata, 231 Netperf, 303 Network File System, 187, 232 Network Settings, 112 New Group, 59 New User, 62 NFS, 187, 232

Ρ

Path and Name Lengths, 12 Periodic Snapshot, 150 Plugin, 253 Professional Support, 299

R

Reboot, 291 Remove Group, 60 Remove User, 64 Replace Failed Drive, 147 Replication, 152 Reporting, 278 Resilver Priority, 161 RFC RFC 2616, 325 RFC 3721, 207 Route, 123 Rsync, 234 Rsync Tasks, 103 Running VMs, 272

S

S.M.A.R.T., 237 S.M.A.R.T. Tests, 110 S3, 236 Samba, 195, 238 SCP, 246 Scrub, 162 Secure Copy, 246 Secure Shell, 245 Self-Encrypting Drives, 74 Services, 218 Shadow Copies, 203 Shell, 288 Shutdown, 292 Simple Network Management Protocol, 243 SMB, 195, 238 Snapshot, 150 Snapshots, 165 SNMP, 243 Spares, 150 SSH, 245 Start Service, 219 Static Route, 123 Stop Service, 219 Support, 92, 293 System Dataset, 76

Т

Tasks, 94 TFTP, 247 Time Machine, 184 tmux, 313 Translate, 316 Translation, 316 Trivial File Transfer Protocol, 247 Trunking, 123 Tunables, 77 tw_cli, 311

U

Uninterruptible Power Supply, 248 Upgrade, 28 Upgrade ZFS Pool, 35 UPS, 248 USB Stick, 18 Users, 60

V

VAAI, 323 VAAI for iSCSI, 324 Virtualization, 36 VLAN, 123 VM, 36 VMs, 265 VMware Snapshot, 167 Volumes, 125

W

WebDAV, 194, 252 Windows File Share, 238 Windows Shares, 195 Wizard, 280

Ζ

ZVOL, 135