



# **Singularity Container Documentation**

*Release 3.2*

**User Docs**

**May 23, 2019**



# CONTENTS

<b>1</b>	<b>Quick Start</b>	<b>1</b>
1.1	Quick Installation Steps	1
1.1.1	Install system dependencies	1
1.1.2	Install Go	1
1.1.3	Clone the Singularity repository	2
1.1.4	Compile the Singularity binary	2
1.2	Overview of the Singularity Interface	2
1.3	Download pre-built images	4
1.4	Interact with images	5
1.4.1	Shell	6
1.4.2	Executing Commands	6
1.4.3	Running a container	7
1.4.4	Working with Files	7
1.5	Build images from scratch	8
1.5.1	Sandbox Directories	8
1.5.2	Converting images from one format to another	8
1.5.3	Singularity Definition Files	9
<b>2</b>	<b>Contributing</b>	<b>11</b>
2.1	Join the community	11
2.1.1	Singularity Google Group	11
2.1.2	Singularity on Slack	11
2.2	Raise an Issue	11
2.3	Write Documentation	11
2.4	Contribute to the code	12
2.4.1	Step 1. Fork the repo	12
2.4.2	Step 2. Checkout a new branch	12
2.4.3	Step 3. Make your changes	12
2.4.4	Step 4. Push your branch to your fork	13
2.4.5	Step 5. Submit a Pull Request	13
2.4.6	Step 6. Keep your branch in sync	13
<b>3</b>	<b>Installation</b>	<b>15</b>
3.1	Overview	15
3.2	Install on Linux	15
3.2.1	Before you begin	15
3.2.1.1	Install Dependencies	15
3.2.1.2	Install Go	16
3.2.2	Download Singularity from a release	16
3.2.3	Download from source	16

3.2.3.1	Compile Singularity . . . . .	17
3.2.3.2	Source bash completion file . . . . .	17
3.2.4	Build and install an RPM . . . . .	17
3.2.5	Remove an old version . . . . .	18
3.2.6	Distribution packages of Singularity . . . . .	18
3.2.6.1	Install the Debian/Ubuntu package using apt . . . . .	19
3.2.6.2	Install the CentOS/RHEL package using yum . . . . .	20
3.3	Install on Windows or Mac . . . . .	20
3.3.1	Setup . . . . .	20
3.3.1.1	Windows . . . . .	20
3.3.1.2	Mac . . . . .	20
3.3.2	Singularity Vagrant Box . . . . .	21
3.4	Singularity on a shared resource . . . . .	21
<b>4</b>	<b>Command Line Interface</b> . . . . .	<b>23</b>
4.1	singularity . . . . .	23
4.1.1	Synopsis . . . . .	23
4.1.2	Examples . . . . .	23
4.1.3	Options . . . . .	23
4.1.4	SEE ALSO . . . . .	23
4.2	singularity apps . . . . .	24
4.2.1	Synopsis . . . . .	24
4.2.2	Examples . . . . .	25
4.2.3	Options . . . . .	25
4.2.4	SEE ALSO . . . . .	25
4.3	singularity build . . . . .	25
4.3.1	Synopsis . . . . .	25
4.3.2	Examples . . . . .	26
4.3.3	Options . . . . .	27
4.3.4	SEE ALSO . . . . .	28
4.4	singularity cache . . . . .	28
4.4.1	Synopsis . . . . .	28
4.4.2	Examples . . . . .	28
4.4.3	Options . . . . .	28
4.4.4	SEE ALSO . . . . .	29
4.5	singularity cache clean . . . . .	29
4.5.1	Synopsis . . . . .	29
4.5.2	Examples . . . . .	29
4.5.3	Options . . . . .	29
4.5.4	SEE ALSO . . . . .	29
4.6	singularity cache list . . . . .	30
4.6.1	Synopsis . . . . .	30
4.6.2	Examples . . . . .	30
4.6.3	Options . . . . .	30
4.6.4	SEE ALSO . . . . .	30
4.7	singularity capability . . . . .	30
4.7.1	Synopsis . . . . .	30
4.7.2	Examples . . . . .	31
4.7.3	Options . . . . .	31
4.7.4	SEE ALSO . . . . .	31
4.8	singularity capability add . . . . .	31
4.8.1	Synopsis . . . . .	31
4.8.2	Examples . . . . .	31
4.8.3	Options . . . . .	32

4.8.4	SEE ALSO	32
4.9	singularity capability avail	32
4.9.1	Synopsis	32
4.9.2	Examples	32
4.9.3	Options	32
4.9.4	SEE ALSO	32
4.10	singularity capability drop	33
4.10.1	Synopsis	33
4.10.2	Examples	33
4.10.3	Options	33
4.10.4	SEE ALSO	33
4.11	singularity capability list	33
4.11.1	Synopsis	33
4.11.2	Examples	34
4.11.3	Options	34
4.11.4	SEE ALSO	34
4.12	singularity exec	34
4.12.1	Synopsis	34
4.12.2	Examples	34
4.12.3	Options	35
4.12.4	SEE ALSO	36
4.13	singularity inspect	36
4.13.1	Synopsis	36
4.13.2	Examples	36
4.13.3	Options	36
4.13.4	SEE ALSO	37
4.14	singularity instance	37
4.14.1	Synopsis	37
4.14.2	Examples	37
4.14.3	Options	37
4.14.4	SEE ALSO	37
4.15	singularity instance list	38
4.15.1	Synopsis	38
4.15.2	Examples	38
4.15.3	Options	38
4.15.4	SEE ALSO	38
4.16	singularity instance start	38
4.16.1	Synopsis	38
4.16.2	Examples	39
4.16.3	Options	39
4.16.4	SEE ALSO	40
4.17	singularity instance stop	40
4.17.1	Synopsis	41
4.17.2	Examples	41
4.17.3	Options	41
4.17.4	SEE ALSO	41
4.18	singularity key	41
4.18.1	Synopsis	41
4.18.2	Examples	42
4.18.3	Options	42
4.18.4	SEE ALSO	42
4.19	singularity key export	42
4.19.1	Synopsis	42
4.19.2	Examples	42

4.19.3	Options	43
4.19.4	SEE ALSO	43
4.20	singularity key import	43
4.20.1	Synopsis	43
4.20.2	Examples	43
4.20.3	Options	43
4.20.4	SEE ALSO	43
4.21	singularity key list	43
4.21.1	Synopsis	44
4.21.2	Examples	44
4.21.3	Options	44
4.21.4	SEE ALSO	44
4.22	singularity key newpair	44
4.22.1	Synopsis	44
4.22.2	Examples	44
4.22.3	Options	44
4.22.4	SEE ALSO	45
4.23	singularity key pull	45
4.23.1	Synopsis	45
4.23.2	Examples	45
4.23.3	Options	45
4.23.4	SEE ALSO	45
4.24	singularity key push	45
4.24.1	Synopsis	45
4.24.2	Examples	46
4.24.3	Options	46
4.24.4	SEE ALSO	46
4.25	singularity key remove	46
4.25.1	Synopsis	46
4.25.2	Examples	46
4.25.3	Options	46
4.25.4	SEE ALSO	46
4.26	singularity key search	46
4.26.1	Synopsis	47
4.26.2	Examples	47
4.26.3	Options	47
4.26.4	SEE ALSO	47
4.27	singularity oci	47
4.27.1	Synopsis	47
4.27.2	Examples	47
4.27.3	Options	48
4.27.4	SEE ALSO	48
4.28	singularity oci attach	48
4.28.1	Synopsis	48
4.28.2	Examples	48
4.28.3	Options	48
4.28.4	SEE ALSO	48
4.29	singularity oci create	49
4.29.1	Synopsis	49
4.29.2	Examples	49
4.29.3	Options	49
4.29.4	SEE ALSO	49
4.30	singularity oci delete	49
4.30.1	Synopsis	49

4.30.2	Examples	50
4.30.3	Options	50
4.30.4	SEE ALSO	50
4.31	singularity oci exec	50
4.31.1	Synopsis	50
4.31.2	Examples	50
4.31.3	Options	50
4.31.4	SEE ALSO	50
4.32	singularity oci kill	50
4.32.1	Synopsis	51
4.32.2	Examples	51
4.32.3	Options	51
4.32.4	SEE ALSO	51
4.33	singularity oci mount	51
4.33.1	Synopsis	51
4.33.2	Examples	51
4.33.3	Options	51
4.33.4	SEE ALSO	52
4.34	singularity oci pause	52
4.34.1	Synopsis	52
4.34.2	Examples	52
4.34.3	Options	52
4.34.4	SEE ALSO	52
4.35	singularity oci resume	52
4.35.1	Synopsis	52
4.35.2	Examples	52
4.35.3	Options	53
4.35.4	SEE ALSO	53
4.36	singularity oci run	53
4.36.1	Synopsis	53
4.36.2	Examples	53
4.36.3	Options	53
4.36.4	SEE ALSO	53
4.37	singularity oci start	54
4.37.1	Synopsis	54
4.37.2	Examples	54
4.37.3	Options	54
4.37.4	SEE ALSO	54
4.38	singularity oci state	54
4.38.1	Synopsis	54
4.38.2	Examples	54
4.38.3	Options	54
4.38.4	SEE ALSO	55
4.39	singularity oci umount	55
4.39.1	Synopsis	55
4.39.2	Examples	55
4.39.3	Options	55
4.39.4	SEE ALSO	55
4.40	singularity oci update	55
4.40.1	Synopsis	55
4.40.2	Examples	55
4.40.3	Options	56
4.40.4	SEE ALSO	56
4.41	singularity plugin	56

4.41.1	Synopsis . . . . .	56
4.41.2	Examples . . . . .	56
4.41.3	Options . . . . .	56
4.41.4	SEE ALSO . . . . .	56
4.42	singularity plugin compile . . . . .	57
4.42.1	Synopsis . . . . .	57
4.42.2	Examples . . . . .	57
4.42.3	Options . . . . .	57
4.42.4	SEE ALSO . . . . .	57
4.43	singularity plugin disable . . . . .	57
4.43.1	Synopsis . . . . .	57
4.43.2	Examples . . . . .	58
4.43.3	Options . . . . .	58
4.43.4	SEE ALSO . . . . .	58
4.44	singularity plugin enable . . . . .	58
4.44.1	Synopsis . . . . .	58
4.44.2	Examples . . . . .	58
4.44.3	Options . . . . .	58
4.44.4	SEE ALSO . . . . .	58
4.45	singularity plugin inspect . . . . .	58
4.45.1	Synopsis . . . . .	59
4.45.2	Examples . . . . .	59
4.45.3	Options . . . . .	59
4.45.4	SEE ALSO . . . . .	59
4.46	singularity plugin install . . . . .	59
4.46.1	Synopsis . . . . .	59
4.46.2	Examples . . . . .	59
4.46.3	Options . . . . .	59
4.46.4	SEE ALSO . . . . .	60
4.47	singularity plugin list . . . . .	60
4.47.1	Synopsis . . . . .	60
4.47.2	Examples . . . . .	60
4.47.3	Options . . . . .	60
4.47.4	SEE ALSO . . . . .	60
4.48	singularity plugin uninstall . . . . .	60
4.48.1	Synopsis . . . . .	60
4.48.2	Examples . . . . .	60
4.48.3	Options . . . . .	61
4.48.4	SEE ALSO . . . . .	61
4.49	singularity pull . . . . .	61
4.49.1	Synopsis . . . . .	61
4.49.2	Examples . . . . .	61
4.49.3	Options . . . . .	61
4.49.4	SEE ALSO . . . . .	62
4.50	singularity push . . . . .	62
4.50.1	Synopsis . . . . .	62
4.50.2	Examples . . . . .	62
4.50.3	Options . . . . .	62
4.50.4	SEE ALSO . . . . .	62
4.51	singularity remote . . . . .	62
4.51.1	Synopsis . . . . .	63
4.51.2	Examples . . . . .	63
4.51.3	Options . . . . .	63
4.51.4	SEE ALSO . . . . .	63



4.52	singularity remote add	63
4.52.1	Synopsis	63
4.52.2	Examples	63
4.52.3	Options	64
4.52.4	SEE ALSO	64
4.53	singularity remote list	64
4.53.1	Synopsis	64
4.53.2	Examples	64
4.53.3	Options	64
4.53.4	SEE ALSO	64
4.54	singularity remote login	64
4.54.1	Synopsis	65
4.54.2	Examples	65
4.54.3	Options	65
4.54.4	SEE ALSO	65
4.55	singularity remote remove	65
4.55.1	Synopsis	65
4.55.2	Examples	65
4.55.3	Options	65
4.55.4	SEE ALSO	66
4.56	singularity remote status	66
4.56.1	Synopsis	66
4.56.2	Examples	66
4.56.3	Options	66
4.56.4	SEE ALSO	66
4.57	singularity remote use	66
4.57.1	Synopsis	66
4.57.2	Examples	67
4.57.3	Options	67
4.57.4	SEE ALSO	67
4.58	singularity run	67
4.58.1	Synopsis	67
4.58.2	Examples	67
4.58.3	Options	68
4.58.4	SEE ALSO	69
4.59	singularity run-help	69
4.59.1	Synopsis	69
4.59.2	Examples	69
4.59.3	Options	70
4.59.4	SEE ALSO	70
4.60	singularity search	70
4.60.1	Synopsis	70
4.60.2	Examples	70
4.60.3	Options	71
4.60.4	SEE ALSO	71
4.61	singularity shell	71
4.61.1	Synopsis	71
4.61.2	Examples	71
4.61.3	Options	72
4.61.4	SEE ALSO	73
4.62	singularity sign	73
4.62.1	Synopsis	73
4.62.2	Examples	74
4.62.3	Options	74

4.62.4	SEE ALSO	74
4.63	singularity test	74
4.63.1	Synopsis	74
4.63.2	Examples	74
4.63.3	Options	75
4.63.4	SEE ALSO	76
4.64	singularity verify	76
4.64.1	Synopsis	76
4.64.2	Examples	76
4.64.3	Options	76
4.64.4	SEE ALSO	77
4.65	singularity version	77
4.65.1	Synopsis	77
4.65.2	Options	77
4.65.3	SEE ALSO	77
<b>5</b>	<b>Build a Container</b>	<b>79</b>
5.1	Overview	79
5.2	Downloading an existing container from the Container Library	79
5.3	Downloading an existing container from Docker Hub	80
5.4	Creating writable <code>--sandbox</code> directories	80
5.5	Converting containers from one format to another	80
5.6	Building containers from Singularity definition files	80
5.7	Build options	81
5.7.1	<code>--builder</code>	81
5.7.2	<code>--detached</code>	81
5.7.3	<code>--force</code>	81
5.7.4	<code>--json</code>	81
5.7.5	<code>--library</code>	81
5.7.6	<code>--notest</code>	82
5.7.7	<code>--remote</code>	82
5.7.8	<code>--sandbox</code>	82
5.7.9	<code>--section</code>	82
5.7.10	<code>--update</code>	82
5.8	More Build topics	82
<b>6</b>	<b>Definition Files</b>	<b>83</b>
6.1	Overview	83
6.2	Header	83
6.3	Sections	84
6.3.1	<code>%setup</code>	85
6.3.2	<code>%files</code>	85
6.3.3	<code>%environment</code>	86
6.3.4	<code>%post</code>	87
6.3.5	<code>%runscript</code>	87
6.3.6	<code>%startscript</code>	88
6.3.7	<code>%test</code>	88
6.3.8	<code>%labels</code>	88
6.3.9	<code>%help</code>	89
6.4	Multi-Stage Builds	89
6.5	Apps	90
6.6	Best Practices for Build Recipes	92
<b>7</b>	<b>Build Environment</b>	<b>93</b>

7.1	Overview . . . . .	93
7.2	Cache Folders . . . . .	93
7.3	Cache commands . . . . .	93
	7.3.1 Listing Cache . . . . .	94
	7.3.2 Cleaning the Cache . . . . .	94
7.4	Temporary Folders . . . . .	95
7.5	Pull Folder . . . . .	95
7.6	Environment Variables . . . . .	95
	7.6.1 Defaults . . . . .	96
	7.6.1.1 Docker . . . . .	96
	7.6.1.2 Library . . . . .	96
<b>8</b>	<b>Cloud Library</b> . . . . .	<b>97</b>
8.1	Overview . . . . .	97
8.2	Make an Account . . . . .	97
8.3	Creating a Access token . . . . .	97
8.4	Pushing a Container . . . . .	98
8.5	Pulling a container . . . . .	98
	8.5.1 Pulling your own container . . . . .	98
8.6	Verify/Sign your Container . . . . .	99
8.7	Searching the Library for Containers . . . . .	99
	8.7.1 Using the CLI Search . . . . .	99
8.8	Remote Builder . . . . .	100
	8.8.1 Building from a definition file: . . . . .	100
<b>9</b>	<b>Support for Docker and OCI</b> . . . . .	<b>101</b>
9.1	Overview . . . . .	101
9.2	Running action commands on public images from Docker Hub . . . . .	101
9.3	Making use of public images from Docker Hub . . . . .	104
9.4	Making use of private images from Docker Hub . . . . .	105
	9.4.1 Authentication via Interactive Login . . . . .	106
	9.4.2 Authentication via Environment Variables . . . . .	106
9.5	Making use of private images from Private Registries . . . . .	107
9.6	Building images for Singularity from Docker Registries . . . . .	108
	9.6.1 Working from the Singularity Command Line . . . . .	108
	9.6.1.1 Remotely Hosted Images . . . . .	108
	9.6.1.2 Building Containers Remotely . . . . .	110
	9.6.1.3 Locally Available Images: Cached by Docker . . . . .	111
	9.6.1.4 Locally Available Images: Stored Archives . . . . .	112
	9.6.1.5 Pushing Locally Available Images to a Library . . . . .	114
	9.6.2 Working with Definition Files . . . . .	116
	9.6.2.1 Mandatory Header Keywords: Remotely Boostrapped . . . . .	116
	9.6.2.2 Remotely Bootstrapped and Built Containers . . . . .	117
	9.6.2.3 Mandatory Header Keywords: Locally Boostrapped . . . . .	119
	9.6.2.4 Optional Header Keywords . . . . .	120
	9.6.2.5 Private Images and Registries . . . . .	121
	9.6.2.6 Directing Execution . . . . .	121
	9.6.2.7 Container Metadata . . . . .	123
9.7	OCI Image Support . . . . .	124
	9.7.1 Overview . . . . .	124
	9.7.1.1 Image Pulls Revisited . . . . .	124
	9.7.1.2 Image Caching in Singularity . . . . .	125
	9.7.1.3 Compliance with the OCI Image Layout Specification . . . . .	125
	9.7.1.4 OCI Compliance and the Singularity Cache . . . . .	126

9.7.2	Building Containers for Singularity from OCI Images . . . . .	129
9.7.2.1	Working Locally from the Singularity Command Line: <code>oci</code> Bootstrap Agent . . . . .	129
9.7.2.2	Working Locally from the Singularity Command Line: <code>oci-archive</code> Bootstrap Agent . . . . .	130
9.7.2.3	Working from the Singularity Command Line with Remotely Hosted Images . . . . .	132
9.7.2.4	Working with Definition Files: Mandatory Header Keywords . . . . .	133
9.7.2.5	Working with Definition Files: Additional Considerations . . . . .	134
9.8	Best Practices . . . . .	135
9.9	Troubleshooting . . . . .	137
<b>10</b>	<b>Bind Paths and Mounts</b> . . . . .	<b>139</b>
10.1	Overview . . . . .	139
10.2	System-defined bind paths . . . . .	139
10.3	User-defined bind paths . . . . .	139
10.3.1	Specifying bind paths . . . . .	139
10.3.2	A note on using <code>--bind</code> with the <code>--writable</code> flag . . . . .	140
<b>11</b>	<b>Persistent Overlays</b> . . . . .	<b>141</b>
11.1	Overview . . . . .	141
11.2	Usage . . . . .	141
<b>12</b>	<b>Running Services</b> . . . . .	<b>143</b>
12.1	Overview . . . . .	143
12.2	Container Instances in Singularity . . . . .	143
12.3	Nginx “Hello-world” in Singularity . . . . .	145
12.4	Putting all together . . . . .	146
12.4.1	Building the image . . . . .	146
12.4.2	Running the Service . . . . .	147
12.4.3	Making it Fancy . . . . .	150
<b>13</b>	<b>Environment and Metadata</b> . . . . .	<b>153</b>
13.1	Overview . . . . .	153
13.2	Environment . . . . .	154
13.3	Labels . . . . .	155
13.4	The <code>inspect</code> command . . . . .	156
13.4.1	<code>--labels</code> . . . . .	156
13.4.2	<code>--deffile</code> . . . . .	156
13.4.3	<code>--runscript</code> . . . . .	157
13.4.4	<code>--test</code> . . . . .	158
13.4.5	<code>--environment</code> . . . . .	158
13.4.6	<code>--helpfile</code> . . . . .	159
13.4.7	<code>--json</code> . . . . .	159
13.5	Container Metadata . . . . .	159
<b>14</b>	<b>OCI Runtime Support</b> . . . . .	<b>161</b>
14.1	Overview . . . . .	161
14.2	Mounted OCI Filesystem Bundles . . . . .	161
14.2.1	Mounting an OCI Filesystem Bundle . . . . .	161
14.2.2	Content of an OCI Compliant Filesystem Bundle . . . . .	162
14.3	Creating OCI Compliant Container Instances . . . . .	174
14.4	Unmounting OCI Filesystem Bundles . . . . .	174
14.5	Kubernetes Integration . . . . .	175
<b>15</b>	<b>Key commands</b> . . . . .	<b>177</b>
15.1	Key <code>import</code> command . . . . .	177

15.2	Key export command . . . . .	178
15.3	Key remove command . . . . .	179
<b>16</b>	<b>Signing and Verifying Containers</b>	<b>181</b>
16.1	Verifying containers from the Container Library . . . . .	181
16.2	Signing your own containers . . . . .	181
16.2.1	Generating and managing PGP keys . . . . .	181
16.2.2	Searching for keys . . . . .	182
16.2.3	Signing and validating your own containers . . . . .	183
<b>17</b>	<b>Plugins</b>	<b>185</b>
17.1	Overview . . . . .	185
17.2	Using Plugins . . . . .	185
17.3	Writing a Plugin . . . . .	186
<b>18</b>	<b>Security Options</b>	<b>187</b>
18.1	Linux Capabilities . . . . .	187
18.2	Security related action options . . . . .	188
18.2.1	--add-caps . . . . .	188
18.2.2	--allow-setuid . . . . .	188
18.2.3	--keep-privs . . . . .	188
18.2.4	--drop-caps . . . . .	188
18.2.5	--security . . . . .	189
<b>19</b>	<b>Network virtualization</b>	<b>191</b>
19.1	--dns . . . . .	191
19.2	--hostname . . . . .	191
19.3	--net . . . . .	191
19.4	--network . . . . .	192
19.5	--network-args . . . . .	192
<b>20</b>	<b>Limiting container resources with cgroups</b>	<b>195</b>
20.1	Overview . . . . .	195
20.2	Examples . . . . .	195
20.2.1	Limiting memory . . . . .	195
20.2.2	Limiting CPU . . . . .	196
20.2.2.1	shares . . . . .	196
20.2.2.2	quota/period . . . . .	196
20.2.2.3	cpus/mems . . . . .	196
20.2.3	Limiting IO . . . . .	197
20.2.3.1	Limiting device access . . . . .	198
<b>21</b>	<b>Appendix</b>	<b>199</b>
21.1	Singularity's environment variables . . . . .	199
21.1.1	A . . . . .	199
21.1.2	B . . . . .	199
21.1.3	C . . . . .	199
21.1.4	D . . . . .	200
21.1.5	E . . . . .	200
21.1.6	F . . . . .	200
21.1.7	G . . . . .	200
21.1.8	H . . . . .	201
21.1.9	I . . . . .	201
21.1.10	J . . . . .	201
21.1.11	K . . . . .	201

21.1.12	L	201
21.1.13	N	201
21.1.14	O	202
21.1.15	P	202
21.1.16	R	202
21.1.17	S	202
21.1.18	T	202
21.1.19	U	202
21.1.20	W	203
21.2	Build Modules	203
21.2.1	library bootstrap agent	203
21.2.1.1	Overview	203
21.2.1.2	Keywords	203
21.2.2	docker bootstrap agent	204
21.2.2.1	Overview	204
21.2.2.2	Keywords	204
21.2.2.3	Notes	204
21.2.3	shub bootstrap agent	205
21.2.3.1	Overview	205
21.2.3.2	Keywords	205
21.2.3.3	Notes	205
21.2.4	localimage bootstrap agent	205
21.2.4.1	Overview	205
21.2.4.2	Keywords	205
21.2.4.3	Notes	206
21.2.5	yum bootstrap agent	206
21.2.5.1	Overview	206
21.2.5.2	Keywords	206
21.2.5.3	Notes	206
21.2.6	debootstrap build agent	207
21.2.6.1	Overview	207
21.2.6.2	Keywords	207
21.2.6.3	Notes	207
21.2.7	arch bootstrap agent	208
21.2.7.1	Overview	208
21.2.7.2	Keywords	208
21.2.7.3	Notes	208
21.2.8	busybox bootstrap agent	208
21.2.8.1	Overview	208
21.2.8.2	Keywords	208
21.2.8.3	Notes	208
21.2.9	zypper bootstrap agent	209
21.2.9.1	Overview	209
21.2.9.2	Keywords	209

## QUICK START

This guide is intended for running Singularity on a computer where you have root (administrative) privileges.

If you need to request an installation on your shared resource, see the requesting an installation help page for information to send to your system administrator.

For any additional help or support contact the Sylabs team: <https://www.sylabs.io/contact/>

### 1.1 Quick Installation Steps

You will need a Linux system to run Singularity.

See the [installation page](#) for information about installing older versions of Singularity.

#### 1.1.1 Install system dependencies

You must first install development libraries to your host. Assuming Ubuntu (apply similar to RHEL derivatives):

```
$ sudo apt-get update && sudo apt-get install -y \  
  build-essential \  
  libssl-dev \  
  uuid-dev \  
  libpgm-dev \  
  squashfs-tools \  
  wget \  
  git
```

---

**Note:** Note that `squashfs-tools` is an image build dependency only and is not required for Singularity build and run commands.

---

#### 1.1.2 Install Go

Singularity 3.0 is written primarily in Go, and you will need Go installed to compile it from source.

This is one of several ways to [install and configure Go](#).

First, visit the [Go download page](#) and pick the appropriate Go archive ( $\geq 1.11.1$ ). Copy the link address and download with `wget` like so:

```
$ export VERSION=1.11 OS=linux ARCH=amd64
$ cd /tmp
$ wget https://dl.google.com/go/go$VERSION.$OS-$ARCH.tar.gz
```

Then extract the archive to /usr/local

```
$ sudo tar -C /usr/local -xzf go$VERSION.$OS-$ARCH.tar.gz
```

Finally, set up your environment for Go

```
$ echo 'export GOPATH=${HOME}/go' >> ~/.bashrc
$ echo 'export PATH=/usr/local/go/bin:${PATH}:${GOPATH}/bin' >> ~/.bashrc
$ source ~/.bashrc
```

### 1.1.3 Clone the Singularity repository

Go is a bit finicky about where things are placed. Here is the correct way to build Singularity from source.

```
$ mkdir -p $GOPATH/src/github.com/sylabs
$ cd $GOPATH/src/github.com/sylabs
$ git clone https://github.com/sylabs/singularity.git
$ cd singularity
```

### 1.1.4 Compile the Singularity binary

Now you are ready to build Singularity. Dependencies will be automatically downloaded. You can build Singularity using the following commands:

```
$ cd $GOPATH/src/github.com/sylabs/singularity
$ ./mconfig
$ make -C builddir
$ sudo make -C builddir install
```

Singularity must be installed as root to function properly.

## 1.2 Overview of the Singularity Interface

Singularity's command line interface allows you to build and interact with containers transparently. You can run programs inside a container as if they were running on your host system. You can easily redirect IO, use pipes, pass arguments, and access files, sockets, and ports on the host system from within a container.

The `help` command gives an overview of Singularity options and subcommands as follows:

```
$ singularity help

Linux container platform optimized for High Performance Computing (HPC) and
Enterprise Performance Computing (EPC)

Usage:
  singularity [global options...]

Description:
```

(continues on next page)



(continued from previous page)

Singularity containers provide an application virtualization layer enabling mobility of compute via both application and environment portability. With Singularity one is capable of building a root file system that runs on any other Linux system where Singularity is installed.

## Options:

```
-d, --debug          print debugging information (highest verbosity)
-h, --help          help for singularity
-q, --quiet         suppress normal output
-s, --silent        only print errors
-t, --tokenfile string path to the file holding your sylabs
                    authentication token (default
                    "/home/jacob/.singularity/sylabs-token")
-v, --verbose       print additional information
```

## Available Commands:

```
apps          List available apps within a container
build         Build a Singularity image
cache         Manage the local cache
capability    Manage Linux capabilities for users and groups
exec          Run a command within a container
help          Help about any command
inspect       Show metadata for an image
instance      Manage containers running as services
key           Manage OpenPGP keys
oci           Manage OCI containers
pull          Pull an image from a URI
push          Push a container to a Library URI
run           Run the user-defined default command within a container
run-help      Show the user-defined help for an image
search        Search a Library for images
shell         Run a shell within a container
sign          Attach a cryptographic signature to an image
test          Run the user-defined tests within a container
verify        Verify cryptographic signatures attached to an image
version       Show the version for Singularity
```

## Examples:

```
$ singularity help <command>
    Additional help for any Singularity subcommand can be seen by appending
    the subcommand name to the above command.
```

For additional help or support, please visit <https://www.sylabs.io/docs/>

Information about subcommand can also be viewed with the help command.

```
$ singularity help verify
Verify cryptographic signatures on container
```

## Usage:

```
singularity verify [verify options...] <image path>
```

## Description:

The verify command allows a user to verify cryptographic signatures on SIF container files. There may be multiple signatures for data objects and multiple data objects signed. By default the command searches for the primary

(continues on next page)

(continued from previous page)

```
partition signature. If found, a list of all verification blocks applied on
the primary partition is gathered so that data integrity (hashing) and
signature verification is done for all those blocks.
```

Options:

```
-g, --groupid uint32  group ID to be verified
-h, --help           help for verify
-i, --id uint32     descriptor ID to be verified
-u, --url string     key server URL (default "https://keys.sylabs.io")
```

Examples:

```
$ singularity verify container.sif
```

For additional help or support, please visit <https://www.sylabs.io/docs/>

Singularity uses positional syntax (i.e. the order of commands and options matters).

Global options affecting the behavior of all commands follow the main `singularity` command. Then sub commands are passed followed by their options and arguments.

For example, to pass the `--debug` option to the main `singularity` command and run Singularity with debugging messages on:

```
$ singularity --debug run library://sylabsed/examples/lolcow
```

To pass the `--containall` option to the `run` command and run a Singularity image in an isolated manner:

```
$ singularity run --containall library://sylabsed/examples/lolcow
```

Singularity 2.4 introduced the concept of command groups. For instance, to list Linux capabilities for a particular user, you would use the `list` command in the `capabilities` command group like so:

```
$ singularity capability list --user dave
```

Container authors might also write help docs specific to a container or for an internal module called an `app`. If those help docs exist for a particular container, you can view them like so.

```
$ singularity help container.sif # See the container's help, if provided
$ singularity help --app foo container.sif # See the help for foo, if provided
```

### 1.3 Download pre-built images

You can use the `search` command to locate groups, collections, and containers of interest on the [Container Library](#).

```
$ singularity search alp
No users found for 'alp'

Found 1 collections for 'alp'
  library://jchavez/alpine

Found 5 containers for 'alp'
```

(continues on next page)

(continued from previous page)

```
library://jialipassion/official/alpine
  Tags: latest
library://dtrudg/linux/alpine
  Tags: 3.2 3.3 3.4 3.5 3.6 3.7 3.8 edge latest
library://sylabse/linux/alpine
  Tags: 3.6 3.7 latest
library://library/default/alpine
  Tags: 3.1 3.2 3.3 3.4 3.5 3.6 3.7 3.8 latest
library://sylabse/examples/alpine
  Tags: latest
```

You can use the `pull` and `build` commands to download pre-built images from an external resource like the [Container Library](#) or [Docker Hub](#).

When called on a native Singularity image like those provided on the Container Library, `pull` simply downloads the image file to your system.

```
$ singularity pull library://sylabse/linux/alpine
```

You can also use `pull` with the `docker:// uri` to reference Docker images served from a registry. In this case `pull` does not just download an image file. Docker images are stored in layers, so `pull` must also combine those layers into a usable Singularity file.

```
$ singularity pull docker://godlovedc/lolcow
```

Pulling Docker images reduces reproducibility. If you were to pull a Docker image today and then wait six months and pull again, you are not guaranteed to get the same image. If any of the source layers has changed the image will be altered. If reproducibility is a priority for you, try building your images from the Container Library.

You can also use the `build` command to download pre-built images from an external resource. When using `build` you must specify a name for your container like so:

```
$ singularity build ubuntu.sif library://ubuntu
$ singularity build lolcow.sif docker://godlovedc/lolcow
```

Unlike `pull`, `build` will convert your image to the latest Singularity image format after downloading it.

`build` is like a “Swiss Army knife” for container creation. In addition to downloading images, you can use `build` to create images from other images or from scratch using a definition file. You can also use `build` to convert an image between the container formats supported by Singularity.

## 1.4 Interact with images

You can interact with images in several ways. It is not actually necessary to `pull` or `build` an image to interact with it. The commands listed here will work with image URIs in addition to accepting a local path to an image.

For these examples we will use a `lolcow_latest.sif` image that can be pulled from the Container Library like so.

```
$ singularity pull library://sylabse/examples/lolcow
```

### 1.4.1 Shell

The shell command allows you to spawn a new shell within your container and interact with it as though it were a small virtual machine.

```
$ singularity shell lolcow_latest.sif

Singularity lolcow_latest.sif:~>
```

The change in prompt indicates that you have entered the container (though you should not rely on that to determine whether you are in container or not).

Once inside of a Singularity container, you are the same user as you are on the host system.

```
Singularity lolcow_latest.sif:~> whoami
david

Singularity lolcow_latest.sif:~> id
uid=1000(david) gid=1000(david) groups=1000(david),4(adm),24(cdrom),27(sudo),30(dip),
↪46(plugdev),116(lpadmin),126(sambashare)
```

shell also works with the `library://`, `docker://`, and `shub://` URIs. This creates an ephemeral container that disappears when the shell is exited.

```
$ singularity shell library://sylabsed/examples/lolcow
```

### 1.4.2 Executing Commands

The `exec` command allows you to execute a custom command within a container by specifying the image file. For instance, to execute the `cowsay` program within the `lolcow_latest.sif` container:

```
$ singularity exec lolcow_latest.sif cowsay moo

_____
< moo >
-----
      \   ^__^
       \  (oo)\_______
          (__)\       )\/\
             ||----w |
              ||     ||
```

`exec` also works with the `library://`, `docker://`, and `shub://` URIs. This creates an ephemeral container that executes a command and disappears.

```
$ singularity exec library://sylabsed/examples/lolcow cowsay "Fresh from the library!"

_____
< Fresh from the library! >
-----
      \   ^__^
       \  (oo)\_______
          (__)\       )\/\
             ||----w |
              ||     ||
```

### 1.4.3 Running a container

Singularity containers contain *runscripts*. These are user defined scripts that define the actions a container should perform when someone runs it. The runscript can be triggered with the run command, or simply by calling the container as though it were an executable.

```
$ singularity run lolcow_latest.sif
-----
/ You have been selected for a secret \
\ mission.                             /
-----
      ^__^
      (oo)\_______
         (__)\       )\/\
            ||----w |
            ||     ||

$ ./lolcow_latest.sif
-----
/ Q: What is orange and goes "click, \
\ click?" A: A ball point carrot.     /
-----
      ^__^
      (oo)\_______
         (__)\       )\/\
            ||----w |
            ||     ||
```

run also works with the `library://`, `docker://`, and `shub://` URIs. This creates an ephemeral container that runs and then disappears.

```
$ singularity run library://sylabsed/examples/lolcow
-----
/ Is that really YOU that is reading \
\ this?                               /
-----
      ^__^
      (oo)\_______
         (__)\       )\/\
            ||----w |
            ||     ||
```

### 1.4.4 Working with Files

Files on the host are reachable from within the container.

```
$ echo "Hello from inside the container" > $HOME/hostfile.txt
$ singularity exec lolcow_latest.sif cat $HOME/hostfile.txt
Hello from inside the container
```

This example works because `hostfile.txt` exists in the user's home directory. By default Singularity bind mounts `/home/$USER`, `/tmp`, and `$PWD` into your container at runtime.

You can specify additional directories to bind mount into your container with the `--bind` option. In this example, the `data` directory on the host system is bind mounted to the `/mnt` directory inside the container.

```
$ echo "Drink milk (and never eat hamburgers)." > /data/cow_advice.txt
$ singularity exec --bind /data:/mnt lolcow_latest.sif cat /mnt/cow_advice.txt
Drink milk (and never eat hamburgers).
```

Pipes and redirects also work with Singularity commands just like they do with normal Linux commands.

```
$ cat /data/cow_advice.txt | singularity exec lolcow_latest.sif cowsay
-----
< Drink milk (and never eat hamburgers). >
-----
      \   ^__^
       \  (oo)\_______
          (__)\       )\/\
             ||----w |
              ||     ||
```

## 1.5 Build images from scratch

Singularity v3.0 produces immutable images in the Singularity Image File (SIF) format. This ensures reproducible and verifiable images and allows for many extra benefits such as the ability to sign and verify your containers.

However, during testing and debugging you may want an image format that is writable. This way you can `shell` into the image and install software and dependencies until you are satisfied that your container will fulfill your needs. For these scenarios, Singularity also supports the `sandbox` format (which is really just a directory).

For more details about the different build options and best practices, read about the Singularity flow.

### 1.5.1 Sandbox Directories

To build into a `sandbox` (container in a directory) use the `build --sandbox` command and option:

```
$ sudo singularity build --sandbox ubuntu/ library://ubuntu
```

This command creates a directory called `ubuntu/` with an entire Ubuntu Operating System and some Singularity metadata in your current working directory.

You can use commands like `shell`, `exec`, and `run` with this directory just as you would with a Singularity image. If you pass the `--writable` option when you use your container you can also write files within the `sandbox` directory (provided you have the permissions to do so).

```
$ sudo singularity exec --writable ubuntu touch /foo
$ singularity exec ubuntu/ ls /foo
/foo
```

### 1.5.2 Converting images from one format to another

The `build` command allows you to build a container from an existing container. This means that you can use it to convert a container from one format to another. For instance, if you have already created a `sandbox` (directory) and want to convert it to the default immutable image format (`squashfs`) you can do so:

```
$ singularity build new-sif sandbox
```

Doing so may break reproducibility if you have altered your sandbox outside of the context of a definition file, so you are advised to exercise care.

### 1.5.3 Singularity Definition Files

For a reproducible, production-quality container you should build a SIF file using a Singularity definition file. This also makes it easy to add files, environment variables, and install custom software, and still start from your base of choice (e.g., the Container Library).

A definition file has a header and a body. The header determines the base container to begin with, and the body is further divided into sections that do things like install software, setup the environment, and copy files into the container from the host system.

Here is an example of a definition file:

```
BootStrap: library
From: ubuntu:16.04

%post
    apt-get -y update
    apt-get -y install fortune cowsay lolcat

%environment
    export LC_ALL=C
    export PATH=/usr/games:$PATH

%runscript
    fortune | cowsay | lolcat

%labels
    Author GodloveD
```

To build a container from this definition file (assuming it is a file named lolcow.def), you would call build like so:

```
$ sudo singularity build lolcow.sif lolcow.def
```

In this example, the header tells Singularity to use a base Ubuntu 16.04 image from the Container Library.

The `%post` section executes within the container at build time after the base OS has been installed. The `%post` section is therefore the place to perform installations of new applications.

The `%environment` section defines some environment variables that will be available to the container at runtime.

The `%runscript` section defines actions for the container to take when it is executed.

And finally, the `%labels` section allows for custom metadata to be added to the container.

This is a very small example of the things that you can do with a definition file. In addition to building a container from the Container Library, you can start with base images from Docker Hub and use images directly from official repositories such as Ubuntu, Debian, CentOS, Arch, and BusyBox. You can also use an existing container on your host system as a base.

If you want to build Singularity images but you don't have administrative (root) access on your build system, you can build images using the [Remote Builder](#).

This quickstart document just scratches the surface of all of the things you can do with Singularity!

If you need additional help or support, contact the Sylabs team: <https://www.sylabs.io/contact/>





## CONTRIBUTING

Singularity is an open source project, meaning we have the challenge of limited resources. We are grateful for any support that you can offer. Helping other users, raising issues, helping write documentation, or contributing code are all ways to help!

### 2.1 Join the community

This is a huge endeavor, and your help would be greatly appreciated! Post to online communities about Singularity, and request that your distribution vendor, service provider, and system administrators include Singularity for you!

#### 2.1.1 Singularity Google Group

If you have been using Singularity and having good luck with it, join our [Google Group](#) and help out other users!

#### 2.1.2 Singularity on Slack

Many of our users come to Slack for quick help with an issue. You can find us at [singularity-container](#).

### 2.2 Raise an Issue

For general bugs/issues, you can open an issue [at the GitHub repo](#). However, if you find a security related issue/problem, please email Sylabs directly at [security@sylabs.io](mailto:security@sylabs.io). More information about the Sylabs security policies and procedures can be found [here](#)

### 2.3 Write Documentation

We (like almost all open source software providers) have a documentation dilemma. . . We tend to focus on the code features and functionality before working on documentation. And there is very good reason for this: we want to share the love so nobody feels left out!

You can contribute to the documentation by [raising an issue to suggest an improvement](#) or by sending a [pull request on our repository for documentation](#).

The current documentation is generated with:

- [reStructured Text \(RST\)](#) and [ReadTheDocs](#).

Other dependencies include:

- Python 2.7.
- Sphinx.

More information about contributing to the documentation, instructions on how to install the dependencies, and how to generate the files can be obtained [here](#).

For more information on using Git and GitHub to create a pull request suggesting additions and edits to the docs, see the [section on contributing to the code](#). The procedure is identical for contributions to the documentation and the code base.

## 2.4 Contribute to the code

We use the traditional [GitHub Flow](#) to develop. This means that you fork the main repo, create a new branch to make changes, and submit a pull request (PR) to the master branch.

Check out our official [CONTRIBUTING.md](#) document, which also includes a [code of conduct](#).

### 2.4.1 Step 1. Fork the repo

To contribute to Singularity, you should obtain a GitHub account and fork the [Singularity](#) repository. Once forked, clone your fork of the repo to your computer. (Obviously, you should replace `your-username` with your GitHub username.)

```
$ git clone https://github.com/your-username/singularity.git && \  
  cd singularity/
```

### 2.4.2 Step 2. Checkout a new branch

[Branches](#) are a way of isolating your features from the main branch. Given that we've just cloned the repo, we will probably want to make a new branch from master in which to work on our new feature. Lets call that branch `new-feature`:

```
$ git checkout master && \  
  git checkout -b new-feature
```

---

**Note:** You can always check which branch you are in by running `git branch`.

---

### 2.4.3 Step 3. Make your changes

On your new branch, go nuts! Make changes, test them, and when you are happy commit the changes to the branch:

```
$ git add file-changed1 file-changed2...  
$ git commit -m "what changed?"
```

This commit message is important - it should describe exactly the changes that you have made. Good commit messages read like so:

```
$ git commit -m "changed function getConfig in functions.go to output csv to fix #2"
$ git commit -m "updated docs about shell to close #10"
```

The tags `close #10` and `fix #2` are referencing issues that are posted on the upstream repo where you will direct your pull request. When your PR is merged into the master branch, these messages will automatically close the issues, and further, they will link your commits directly to the issues they intend to fix. This will help future maintainers understand your contribution, or (hopefully not) revert the code back to a previous version if necessary.

#### 2.4.4 Step 4. Push your branch to your fork

When you are done with your commits, you should push your branch to your fork (and you can also continuously push commits here as you work):

```
$ git push origin new-feature
```

Note that you should always check the status of your branches to see what has been pushed (or not):

```
$ git status
```

#### 2.4.5 Step 5. Submit a Pull Request

Once you have pushed your branch, then you can go to your fork (in the web GUI on GitHub) and [submit a Pull Request](#). Regardless of the name of your branch, your PR should be submitted to the Syllabs `master` branch. Submitting your PR will open a conversation thread for the maintainers of Singularity to discuss your contribution. At this time, the continuous integration that is linked with the code base will also be executed. If there is an issue, or if the maintainers suggest changes, you can continue to push commits to your branch and they will update the Pull Request.

#### 2.4.6 Step 6. Keep your branch in sync

Cloning the repo will create an exact copy of the Singularity repository at that moment. As you work, your branch may become out of date as others merge changes into the upstream master. In the event that you need to update a branch, you will need to follow the next steps:

```
$ git remote add upstream https://github.com/syllabs/singularity.git && # to add a new ↵
↪remote named "upstream" \
    git checkout master && # or another branch to be updated \
    git pull upstream master && \
    git push origin master && # to update your fork \
    git checkout new-feature && \
    git merge master
```



## INSTALLATION

This document will guide you through the process of installing Singularity  $\geq 3.2.0$  via several different methods. (For instructions on installing earlier versions of Singularity please see [earlier versions of the docs.](#))

### 3.1 Overview

Singularity runs on Linux natively and can also be run on Windows and Mac through virtual machines (VMs). Here we cover several different methods of installing Singularity ( $\geq 3.2.0$ ) on Linux and also give methods for downloading and running VMs with singularity pre-installed from [Vagrant Cloud](#).

### 3.2 Install on Linux

Linux is the only operating system that can support containers because of kernel features like namespaces. You can use these methods to install Singularity on bare metal Linux or a Linux VM.

#### 3.2.1 Before you begin

If you have an earlier version of Singularity installed, you should *remove it* before executing the installation commands. You will also need to install some dependencies and install Go.

##### 3.2.1.1 Install Dependencies

Install these dependencies with `apt-get` or `yum/rpm` as shown below or similar with other package managers.

`apt-get`

```
$ sudo apt-get update && sudo apt-get install -y \  
  build-essential \  
  libssl-dev \  
  uuid-dev \  
  libgpgme11-dev \  
  squashfs-tools \  
  libseccomp-dev \  
  wget \  
  pkg-config
```

`yum`

```
$ sudo yum update -y && \  
  sudo yum groupinstall -y 'Development Tools' && \  
  sudo yum install -y \  
  openssl-devel \  
  libuuid-devel \  
  libseccomp-devel \  
  wget \  
  squashfs-tools
```

### 3.2.1.2 Install Go

This is one of several ways to [install and configure Go](#).

Visit the [Go download page](#) and pick a package archive to download. Copy the link address and download with `wget`. Then extract the archive to `/usr/local` (or use other instructions on go installation page).

```
$ export VERSION=1.12 OS=linux ARCH=amd64 && \  
  wget https://dl.google.com/go/go$VERSION.$OS-$ARCH.tar.gz && \  
  sudo tar -C /usr/local -xzvf go$VERSION.$OS-$ARCH.tar.gz && \  
  rm go$VERSION.$OS-$ARCH.tar.gz
```

Then, set up your environment for Go.

```
$ echo 'export GOPATH=${HOME}/go' >> ~/.bashrc && \  
  echo 'export PATH=/usr/local/go/bin:${PATH}:${GOPATH}/bin' >> ~/.bashrc && \  
  source ~/.bashrc
```

### 3.2.2 Download Singularity from a release

You can download Singularity from one of our releases. To see a full list, visit <https://github.com/sylabs/singularity/releases>. After that you can just run the following commands to proceed with the installation.

```
$ export VERSION=3.2.0 && # adjust this as necessary \  
  wget https://github.com/sylabs/singularity/releases/download/v${VERSION}/  
  ↪singularity-${VERSION}.tar.gz && \  
  tar -xzf singularity-${VERSION}.tar.gz && \  
  cd singularity
```

### 3.2.3 Download from source

The following commands will install Singularity from the [GitHub repo](#) to `/usr/local`. This method will work for `>=v3.2.0`. To install an older tagged release see [older versions of the docs](#).

When installing from source, you can decide to install from either a **tag**, a **release branch**, or from the **master branch**.

- **tag**: GitHub tags form the basis for releases, so installing from a tag is the same as downloading and installing a [specific release](#). Tags are expected to be relatively stable and well-tested.
- **release branch**: A release branch represents the latest version of a minor release with all the newest bug fixes and enhancements (even those that have not yet made it into a point release). For instance, to install v3.2 with the latest bug fixes and enhancements checkout `release-3.2`. Release branches may be less stable than code in a tagged point release.

- **master branch:** The `master` branch contains the latest, bleeding edge version of Singularity. This is the default branch when you clone the source code, so you don't have to check out any new branches to install it. The `master` branch changes quickly and may be unstable.

To ensure that the Singularity source code is downloaded to the appropriate directory use these commands.

```
$ git clone https://github.com/sylabs/singularity.git && \  
  cd singularity && \  
  git checkout v3.2.0
```

### 3.2.3.1 Compile Singularity

Singularity uses a custom build system called `makeit`. `mconfig` is called to generate a `Makefile` and then `make` is used to compile and install.

```
$ ./mconfig && \  
  make -C ./builddir && \  
  sudo make -C ./builddir install
```

By default Singularity will be installed in the `/usr/local` directory hierarchy. You can specify a custom directory with the `--prefix` option, to `mconfig` like so:

```
$ ./mconfig --prefix=/opt/singularity
```

This option can be useful if you want to install multiple versions of Singularity, install a personal version of Singularity on a shared system, or if you want to remove Singularity easily after installing it.

For a full list of `mconfig` options, run `mconfig --help`. Here are some of the most common options that you may need to use when building Singularity from source.

- `--sysconfdir`: Install read-only config files in `sysconfdir`. This option is important if you need the `singularity.conf` file or other configuration files in a custom location.
- `--localstatedir`: Set the state directory where containers are mounted. This is a particularly important option for administrators installing Singularity on a shared file system. The `--localstatedir` should be set to a directory that is present on each individual node.
- `-b`: Build Singularity in a given directory. By default this is `./builddir`.

### 3.2.3.2 Source bash completion file

To enjoy bash completion with Singularity commands and options, source the bash completion file like so. Add this command to your `~/.bashrc` file so that bash completion continues to work in new shells. (Obviously adjust this path if you installed the bash completion file in a different location.)

```
$ . /usr/local/etc/bash_completion.d/singularity
```

### 3.2.4 Build and install an RPM

Building and installing a Singularity RPM allows the installation be more easily managed, upgraded and removed. In Singularity `>=v3.0.1` you can build an RPM directly from the [release tarball](#).

---

**Note:** Be sure to download the correct asset from the [GitHub releases page](#). It should be named `singularity-<version>.tar.gz`.

---

After installing the *dependencies* and installing *Go* as detailed above, you are ready download the tarball and build and install the RPM.

```
$ export VERSION=3.2.0 && # adjust this as necessary \  
  wget https://github.com/sylabs/singularity/releases/download/v${VERSION}/  
  ↪singularity-${VERSION}.tar.gz && \  
  rpmbuild -tb singularity-${VERSION}.tar.gz && \  
  sudo rpm -ivh ~/rpmbuild/RPMS/x86_64/singularity-${VERSION}-1.el7.x86_64.rpm && \  
  rm -rf ~/rpmbuild singularity-${VERSION}*.tar.gz
```

If you encounter a failed dependency error for *golang* but installed it from source, build with this command:

```
rpmbuild -tb --nodeps singularity-${VERSION}.tar.gz
```

Options to *mconfig* can be passed using the familiar syntax to *rpmbuild*. For example, if you want to force the local state directory to */mnt* (instead of the default */var*) you can do the following:

```
rpmbuild -tb --define='_localstatedir /mnt' singularity-${VERSION}.tar.gz
```

---

**Note:** It is very important to set the local state directory to a directory that physically exists on nodes within a cluster when installing Singularity in an HPC environment with a shared file system. Thus the `_localstatedir` option should be of considerable interest to HPC admins.

---

### 3.2.5 Remove an old version

When you run `sudo make install`, the command lists files as they are installed. They must all be removed in order to completely remove Singularity.

For example, in a standard installation of Singularity 3.0.1 and beyond (when building from source) you must remove all of these files and directories to completely remove Singularity.

Obviously, this list of files may differ depending on how you install Singularity or with newer versions of Singularity released following the writing of this document.

```
$ sudo rm -rf \  
  /usr/local/libexec/singularity \  
  /usr/local/var/singularity \  
  /usr/local/etc/singularity \  
  /usr/local/bin/singularity \  
  /usr/local/bin/run-singularity \  
  /usr/local/etc/bash_completion.d/singularity
```

If you anticipate needing to remove Singularity, it might be easier to install it in a custom directory using the `--prefix` option to *mconfig*. In that case Singularity can be uninstalled simply by deleting the parent directory. Or it may be useful to install Singularity *using a package manager* so that it can be updated and/or uninstalled with ease in the future.

### 3.2.6 Distribution packages of Singularity

---

**Note:** Packaged versions of Singularity in Linux distribution repos are maintained by community members. They (necessarily) tend to be older releases of Singularity. For the latest upstream versions of Singularity it is recommended that you build from source using one of the methods detailed above.

---



### 3.2.6.1 Install the Debian/Ubuntu package using apt

Singularity is available on Debian and derivative distributions starting with Debian stretch and the Ubuntu 16.10 releases. The package is called `singularity-container`. For more recent releases of singularity and backports for older Debian and Ubuntu releases, it is recommended that you use the [NeuroDebian repository](#).

Enable the NeuroDebian repository following instructions on the [NeuroDebian](#) site. Use the dropdown menus to find the best mirror for your operating system and location. For example, after selecting Ubuntu 16.04 and selecting a mirror in CA, you are instructed to add these lists:

```
$ sudo wget -O- http://neuro.debian.net/lists/xenial.us-ca.full | sudo tee /etc/apt/
↳sources.list.d/neurodebian.sources.list && \
    sudo apt-key adv --recv-keys --keyserver hkp://pool.sks-keyservers.net:80
↳0xA5D32F012649A5A9 && \
    sudo apt-get update
```

Now singularity can be installed like so:

```
sudo apt-get install -y singularity-container
```

During the above, if you have a previously installed configuration, you might be asked if you want to define a custom configuration/init, or just use the default provided by the package, eg:

```
Configuration file '/etc/singularity/init'
==> File on system created by you or by a script.
==> File also in package provided by package maintainer.
What would you like to do about it ? Your options are:
  Y or I : install the package maintainer's version
  N or O : keep your currently-installed version
  D      : show the differences between the versions
  Z      : start a shell to examine the situation
The default action is to keep your current version.
*** init (Y/I/N/O/D/Z) [default=N] ? Y

Configuration file '/etc/singularity/singularity.conf'
==> File on system created by you or by a script.
==> File also in package provided by package maintainer.
What would you like to do about it ? Your options are:
  Y or I : install the package maintainer's version
  N or O : keep your currently-installed version
  D      : show the differences between the versions
  Z      : start a shell to examine the situation
The default action is to keep your current version.
*** singularity.conf (Y/I/N/O/D/Z) [default=N] ? Y
```

Most users should accept these defaults. For cluster admins, we recommend that you read the [admin docs](#) to get a better understanding of the configuration file options available to you.

After following this procedure, you can check the Singularity version like so:

```
$ singularity --version
2.5.2-dist
```

If you need a backport build of the recent release of Singularity on those or older releases of Debian and Ubuntu, you can see [all the various builds and other information here](#).

### 3.2.6.2 Install the CentOS/RHEL package using yum

The epel (Extra Packages for Enterprise Linux) repos contain Singularity. The singularity package is actually split into two packages called `singularity-runtime` (which simply contains the necessary bits to run singularity containers) and `singularity` (which also gives you the ability to build Singularity containers).

To install Singularity from the epel repos, first install the repos and then install Singularity. For instance, on CentOS6/7 do the following:

```
$ sudo yum update -y && \  
  sudo yum install -y epel-release && \  
  sudo yum update -y && \  
  sudo yum install -y singularity-runtime singularity
```

After following this procedure, you can check the Singularity version like so:

```
$ singularity --version  
2.6.0-dist
```

## 3.3 Install on Windows or Mac

Linux containers like Singularity cannot run natively on Windows or Mac because of basic incompatibilities with the host kernel. (Contrary to a popular misconception, Mac does not run on a Linux kernel. It runs on a kernel called Darwin originally forked from BSD.)

For this reason, the Singularity community maintains a set of Vagrant Boxes via [Vagrant Cloud](#), one of Hashicorp's open source tools. The current versions can be found under the [sylabs](#) organization.

### 3.3.1 Setup

First, install the following software:

#### 3.3.1.1 Windows

Install the following programs:

- [Git for Windows](#)
- [VirtualBox for Windows](#)
- [Vagrant for Windows](#)
- [Vagrant Manager for Windows](#)

#### 3.3.1.2 Mac

With Singularity Desktop for macOS (Alpha Preview):

The disk image file is available [here](#). More information can be found [here](#).

Singularity is also available via Vagrant (installable with [Homebrew](#) or manually) or with the Singularity Desktop for macOS (Alpha Preview).

With Homebrew:

```
$ /usr/bin/ruby -e "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/
↪master/install)"
$ brew cask install virtualbox && \
  brew cask install vagrant && \
  brew cask install vagrant-manager
```

brew cask install vagrant-manager

### 3.3.2 Singularity Vagrant Box

Run GitBash (Windows) or open a terminal (Mac) and create and enter a directory to be used with your Vagrant VM.

```
$ mkdir vm-singularity && \
  cd vm-singularity
```

If you have already created and used this folder for another VM, you will need to destroy the VM and delete the Vagrantfile.

```
$ vagrant destroy && \
  rm Vagrantfile
```

Then issue the following commands to bring up the Virtual Machine. (Substitute a different value for the `$VM` variable if you like.)

```
$ export VM=sylabs/singularity-3.2-ubuntu-bionic64 && \
  vagrant init $VM && \
  vagrant up && \
  vagrant ssh
```

You can check the installed version of Singularity with the following:

```
vagrant@vagrant:~$ singularity version
3.2.0-1
```

Of course, you can also start with a plain OS Vagrant box as a base and then install Singularity using one of the above methods for Linux.

## 3.4 Singularity on a shared resource

Perhaps you are a user who wants a few talking points and background to share with your administrator. Or maybe you are an administrator who needs to decide whether to install Singularity.

This document, and the accompanying administrator documentation provides answers to many common questions.

If you need to request an installation you may decide to draft a message similar to this:

```
Dear shared resource administrator,
```

```
We are interested in having Singularity (https://www.sylabs.io/docs/)
installed on our shared resource. Singularity containers will allow us to
build encapsulated environments, meaning that our work is reproducible and
we are empowered to choose all dependencies including libraries, operating
system, and custom software. Singularity is already in use on many of the
top HPC centers around the world. Examples include:
```

(continues on next page)

(continued from previous page)

```
Texas Advanced Computing Center
GSI Helmholtz Center for Heavy Ion Research
Oak Ridge Leadership Computing Facility
Purdue University
National Institutes of Health HPC
UFIT Research Computing at the University of Florida
San Diego Supercomputing Center
Lawrence Berkeley National Laboratory
University of Chicago
McGill HPC Centre/Calcul Québec
Barcelona Supercomputing Center
Sandia National Lab
Argonne National Lab
```

Importantly, it has a vibrant team of developers, scientists, and HPC administrators that invest heavily in the security and development of the software, and are quick to respond to the needs of the community. To help learn more about Singularity, I thought these items might be of interest:

```
- Security: A discussion of security concerns is discussed at
  https://www.sylabs.io/guides/2.5.2/user-guide/introduction.html#security-and-
  ↪privilege-escalation
```

```
- Installation:
  https://www.sylabs.io/guides/3.2/user-guide/installation.html
```

If you have questions about any of the above, you can email the open source list ([singularity@lbl.gov](mailto:singularity@lbl.gov)), join the open source slack channel ([singularity-container.slack.com](https://singularity-container.slack.com)), or contact the organization that supports Singularity directly to get a human response ([sylabs.io/contact](https://sylabs.io/contact)). I can do my best to facilitate this interaction if help is needed.

Thank you kindly for considering this request!

Best,

User

As is stated in the sample message above, you can always [reach out](#) to us for additional questions or support.

## COMMAND LINE INTERFACE

Below are links to the automatically generated CLI docs

### 4.1 singularity

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

#### 4.1.1 Synopsis

Singularity containers provide an application virtualization layer enabling mobility of compute via both application and environment portability. With Singularity one is capable of building a root file system that runs on any other Linux system where Singularity is installed.

```
singularity [global options...]
```

#### 4.1.2 Examples

```
$ singularity help <command> [<subcommand>]
$ singularity help build
$ singularity help instance start
```

#### 4.1.3 Options

```
-d, --debug      print debugging information (highest verbosity)
-h, --help      help for singularity
  --nocolor     print without color output (default False)
-q, --quiet     suppress normal output
-s, --silent    only print errors
-v, --verbose   print additional information
```

#### 4.1.4 SEE ALSO

- *singularity build* - Build a Singularity image
- *singularity cache* - Manage the local cache

- *singularity capability* - Manage Linux capabilities for users and groups
- *singularity exec* - Run a command within a container
- *singularity inspect* - Show metadata for an image
- *singularity instance* - Manage containers running as services
- *singularity key* - Manage OpenPGP keys
- *singularity oci* - Manage OCI containers
- *singularity plugin* - Manage singularity plugins
- *singularity pull* - Pull an image from a URI
- *singularity push* - Upload image to the provided library (default is “cloud.sylabs.io”)
- *singularity remote* - Manage singularity remote endpoints
- *singularity run* - Run the user-defined default command within a container
- *singularity run-help* - Show the user-defined help for an image
- *singularity search* - Search a Library for images
- *singularity shell* - Run a shell within a container
- *singularity sign* - Attach a cryptographic signature to an image
- *singularity test* - Run the user-defined tests within a container
- *singularity verify* - Verify cryptographic signatures attached to an image
- *singularity version* - Show the version for Singularity

*Auto generated by spf13/cobra on 23-May-2019*

## 4.2 singularity apps

List available apps within a container

### 4.2.1 Synopsis

List applications (apps) installed in a container, located at /scif/apps. See <http://containers-ftw.org/SCI-F/>

To access apps, use shell, exec, run, inspect with `-app <appname>`

The following environment variables are available to you when called from the shell inside the container. The top variables are relevant to the active app (`-app <app>`) and the bottom available for all apps regardless of the active app:

ACTIVE APP ENVIRONMENT:

SCIF\_APPNAME the name of the application SCIF\_APPROOT the application base (/scif/apps/<app>) SCIF\_APPMETA the application metadata folder SCIF\_APPDATA the data base folder for active app

SCIF\_APPINPUT expected input folder within data base folder SCIF\_APPOUTPUT the output data folder within data base folder

GLOBAL APP ENVIRONMENT:

SCIF\_DATA scif defined data base for all apps (/scif/data) SCIF\_APPS scif defined install bases for all apps (/scif/apps) SCIF\_APPROOT\_<app> root for application <app> SCIF\_APPDATA\_<app> data root for application <app>

For additional help, please visit our public documentation pages which are found at:

<https://www.sylabs.io/docs/>

```
singularity apps <image path>
```

## 4.2.2 Examples

```
$ singularity apps ubuntu.img
bar
foo
```

## 4.2.3 Options

```
-h, --help help for apps
```

## 4.2.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 24-Apr-2019*

## 4.3 singularity build

Build a Singularity image

### 4.3.1 Synopsis

IMAGE PATH:

When Singularity builds the container, output can be one of a few formats:

default: The compressed Singularity read only image format (default) sandbox: This is a read-write container within a directory structure

note: It is a common workflow to use the “sandbox” mode for development of the container, and then build it as a default Singularity image for production use. The default format is immutable.

BUILD SPEC:

The build spec target is a definition (def) file, local image, or URI that can be used to create a Singularity container. Several different local target formats exist:

def file : This is a recipe for building a container (examples below) directory: A directory structure containing a (ch)root file system image: A local image on your machine (will convert to sif if

it is legacy format)

Targets can also be remote and defined by a URI of the following formats:

library:// an image library (default <https://cloud.sylabs.io/library>) docker:// a Docker registry (default Docker Hub) shub:// a Singularity registry (default Singularity Hub)

```
singularity build [local options...] <IMAGE PATH> <BUILD SPEC>
```

### 4.3.2 Examples

```
DEF FILE BASE OS:

Library:
  Bootstrap: library
  From: debian:9

Docker:
  Bootstrap: docker
  From: tensorflow/tensorflow:latest
  IncludeCmd: yes # Use the CMD as runscript instead of ENTRYPOINT

Singularity Hub:
  Bootstrap: shub
  From: singularityhub/centos

YUM/RHEL:
  Bootstrap: yum
  OSVersion: 7
  MirrorURL: http://mirror.centos.org/centos-%{OSVERSION}/%{OSVERSION}/os/x86_
↪64/
  Include: yum

Debian/Ubuntu:
  Bootstrap: debootstrap
  OSVersion: trusty
  MirrorURL: http://us.archive.ubuntu.com/ubuntu/

Local Image:
  Bootstrap: localimage
  From: /home/dave/starter.img

Scratch:
  Bootstrap: scratch # Populate the container with a minimal rootfs in %setup

DEFFILE SECTIONS:

%pre
  echo "This is a scriptlet that will be executed on the host, as root before"
  echo "the container has been bootstrapped. This section is not commonly used."

%setup
  echo "This is a scriptlet that will be executed on the host, as root, after"
  echo "the container has been bootstrapped. To install things into the ↪
↪container"
  echo "reference the file system location with $SINGULARITY_ROOTFS."
```

(continues on next page)



(continued from previous page)

```

%post
    echo "This scriptlet section will be executed from within the container after"
    echo "the bootstrap/base has been created and setup."

%test
    echo "Define any test commands that should be executed after container has_
↳been"
    echo "built. This scriptlet will be executed from within the running container
↳"
    echo "as the root user. Pay attention to the exit/return value of this_
↳scriptlet"
    echo "as any non-zero exit code will be assumed as failure."
    exit 0

%runscript
    echo "Define actions for the container to be executed with the run command or"
    echo "when container is executed."

%startscript
    echo "Define actions for container to perform when started as an instance."

%labels
    HELLO MOTO
    KEY VALUE

%files
    /path/on/host/file.txt /path/on/container/file.txt
    relative_file.txt /path/on/container/relative_file.txt

%environment
    LUKE=goodguy
    VADER=badguy
    HAN=someguy
    export HAN VADER LUKE

%help
    This is a text file to be displayed with the run-help command.

COMMANDS:

Build a sif file from a Singularity recipe file:
    $ singularity build /tmp/debian0.sif /path/to/debian.def

Build a sif image from the Library:
    $ singularity build /tmp/debian1.sif library://debian:latest

Build a base sandbox from DockerHub, make changes to it, then build sif
    $ singularity build --sandbox /tmp/debian docker://debian:latest
    $ singularity exec --writable /tmp/debian apt-get install python
    $ singularity build /tmp/debian2.sif /tmp/debian

```

### 4.3.3 Options

```

--builder string    remote Build Service URL, setting this implies --remote_
↳(default "https://build.sylabs.io")

```

(continues on next page)

(continued from previous page)

```

-d, --detached          submit build job and print build ID (no real-time logs and
↳requires --remote)
  --docker-login       login to a Docker Repository interactively
-F, --force            delete and overwrite an image if it currently exists
-h, --help             help for build
  --json              interpret build definition as JSON
  --library string     container Library URL (default "https://library.sylabs.io")
  --no-cleanup        do NOT clean up bundle after failed build, can be helpful for
↳debugging
  --nohttps           do NOT use HTTPS, for communicating with local docker registry
-T, --notest          build without running tests in %test section
-r, --remote          build image remotely (does not require root)
-s, --sandbox         build image as sandbox format (chroot directory structure)
  --section strings   only run specific section(s) of deffile (setup, post, files,
↳environment, test, labels, none) (default [all])
  --tmpdir string     specify a temporary directory to use for build
-u, --update          run definition over existing container (skips header)

```

#### 4.3.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.4 singularity cache

Manage the local cache

### 4.4.1 Synopsis

Manage your local singularity cache. You can list/clean using the specific types.

```
singularity cache
```

### 4.4.2 Examples

All group commands have their own help output:

```

$ singularity cache
$ singularity cache --help

```

### 4.4.3 Options

```
-h, --help  help for cache
```

#### 4.4.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* *singularity cache clean* - Clean your local Singularity cache \* *singularity cache list* - List your local Singularity cache

*Auto generated by spf13/cobra on 23-May-2019*

## 4.5 singularity cache clean

Clean your local Singularity cache

### 4.5.1 Synopsis

This will clean your local cache (stored at \$HOME/.singularity/cache if SINGULARITY\_CACHEDIR is not set). By default only blob cache is cleaned, use ‘-all’ to clean the entire cache. Note: if you use Singularity as root, cache will be stored in ‘/root/.singularity/.cache’, to clean that cache, you will need to run ‘cache clean -all’ as root, or with ‘sudo’.

```
singularity cache clean [clean options...]
```

### 4.5.2 Examples

All group commands have their own help output:

```
$ singularity help cache clean --name cache_name.sif
$ singularity help cache clean --type=library,oci
$ singularity cache clean --help
```

### 4.5.3 Options

```
-a, --all           clean all cache (will override all other options)
-h, --help         help for clean
-N, --name string  specify a container cache to clean (will clear all cache with_
↳the same name)
-T, --type strings clean cache type, choose between: library, oci, and blob_
↳(default [blob])
```

#### 4.5.4 SEE ALSO

- *singularity cache* - Manage the local cache

*Auto generated by spf13/cobra on 23-May-2019*

## 4.6 singularity cache list

List your local Singularity cache

### 4.6.1 Synopsis

This will list your local cache (stored at `$HOME/.singularity/cache` if `SINGULARITY_CACHEDIR` is not set).

```
singularity cache list [list options...]
```

### 4.6.2 Examples

All group commands have their own help output:

```
$ singularity help cache list
$ singularity help cache list --type=library,oci
$ singularity cache list --help
```

### 4.6.3 Options

```
-a, --all           list all cache types
-h, --help         help for list
-s, --summary      display a cache summary
-T, --type strings a list of cache types to display, possible entries: library, oci,
↳ blob(s), blobSum, all (default [library,oci,blobSum])
```

### 4.6.4 SEE ALSO

- *singularity cache* - Manage the local cache

*Auto generated by spf13/cobra on 23-May-2019*

## 4.7 singularity capability

Manage Linux capabilities for users and groups

### 4.7.1 Synopsis

Capabilities allow you to have fine grained control over the permissions that your containers need to run.

NOTE: capability add/drop commands requires root to run.

```
singularity capability
```

## 4.7.2 Examples

All group commands have their own help output:

```
$ singularity help capability add
$ singularity capability add --help
```

## 4.7.3 Options

```
-h, --help  help for capability
```

## 4.7.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* *singularity capability add* - Add capabilities to a user or group (requires root) \* *singularity capability avail* - Show description for available capabilities \* *singularity capability drop* - Remove capabilities from a user or group (requires root) \* *singularity capability list* - Show capabilities for a given user or group

*Auto generated by spf13/cobra on 23-May-2019*

## 4.8 singularity capability add

Add capabilities to a user or group (requires root)

### 4.8.1 Synopsis

Add Linux capabilities to a user or group. NOTE: This command requires root to run.

The capabilities argument must be separated by commas and is not case sensitive.

To see available capabilities, type “singularity capability avail” or refer to capabilities manual “man 7 capabilities”.

```
singularity capability add [add options...] <capabilities>
```

### 4.8.2 Examples

```
$ sudo singularity capability add --user nobody AUDIT_READ, chown
$ sudo singularity capability add --group nobody cap_audit_write
```

To add all capabilities to a user:

```
$ sudo singularity capability add --user nobody all
```

### 4.8.3 Options

```
-g, --group string  manage capabilities for a group
-h, --help          help for add
-u, --user string   manage capabilities for a user
```

### 4.8.4 SEE ALSO

- *singularity capability* - Manage Linux capabilities for users and groups

*Auto generated by spf13/cobra on 23-May-2019*

## 4.9 singularity capability avail

Show description for available capabilities

### 4.9.1 Synopsis

Show description for available Linux capabilities.

```
singularity capability avail [capabilities]
```

### 4.9.2 Examples

```
Show description for all available capabilities:
$ singularity capability avail

Show CAP_CHOWN description:
$ singularity capability avail CAP_CHOWN

Show CAP_CHOWN/CAP_NET_RAW description:
$ singularity capability avail CAP_CHOWN,CAP_NET_RAW
```

### 4.9.3 Options

```
-h, --help  help for avail
```

### 4.9.4 SEE ALSO

- *singularity capability* - Manage Linux capabilities for users and groups

*Auto generated by spf13/cobra on 23-May-2019*

## 4.10 singularity capability drop

Remove capabilities from a user or group (requires root)

### 4.10.1 Synopsis

Remove Linux capabilities from an user/group. NOTE: This command requires root to run.

The capabilities argument must be separated by commas and is not case sensitive.

To see available capabilities, type “singularity capability avail” or refer to capabilities manual “man 7 capabilities”

```
singularity capability drop [drop options...] <capabilities>
```

### 4.10.2 Examples

```
$ sudo singularity capability drop --user nobody AUDIT_READ,CHOWN
$ sudo singularity capability drop --group nobody audit_write
```

To drop all capabilities for a user:

```
$ sudo singularity capability drop --user nobody all
```

### 4.10.3 Options

```
-g, --group string  manage capabilities for a group
-h, --help          help for drop
-u, --user string   manage capabilities for a user
```

### 4.10.4 SEE ALSO

- *singularity capability* - Manage Linux capabilities for users and groups

*Auto generated by spf13/cobra on 23-May-2019*

## 4.11 singularity capability list

Show capabilities for a given user or group

### 4.11.1 Synopsis

Show the capabilities for a user or group.

```
singularity capability list [user/group]
```

### 4.11.2 Examples

```
To list capabilities set for user or group nobody:
$ singularity capability list nobody

To list capabilities for all users/groups:
$ singularity capability list
```

### 4.11.3 Options

```
-h, --help  help for list
```

### 4.11.4 SEE ALSO

- *singularity capability* - Manage Linux capabilities for users and groups

*Auto generated by spf13/cobra on 23-May-2019*

## 4.12 singularity exec

Run a command within a container

### 4.12.1 Synopsis

singularity exec supports the following formats:

\*.sif Singularity Image Format (SIF). Native to Singularity 3.0+

\*.squash SquashFS format. Native to Singularity 2.4+

\*.img ext3 format. Native to Singularity versions < 2.4.

**directory/ sandbox format.** Directory containing a valid root file system and optionally Singularity meta-data.

**instance://\*** A local running instance of a container. (See the `instance` command group.)

**library://\*** A container hosted on a Library (default <https://cloud.sylabs.io/library>)

**docker://\*** A container hosted on Docker Hub

**shub://\*** A container hosted on Singularity Hub

```
singularity exec [exec options...] <container> <command>
```

### 4.12.2 Examples



```

$ singularity exec /tmp/debian.sif cat /etc/debian_version
$ singularity exec /tmp/debian.sif python ./hello_world.py
$ cat hello_world.py | singularity exec /tmp/debian.sif python
$ sudo singularity exec --writable /tmp/debian.sif apt-get update
$ singularity exec instance://my_instance ps -ef
$ singularity exec library://centos cat /etc/os-release

```

### 4.12.3 Options

```

--add-caps string      a comma separated capability list to add
--allow-setuid        allow setuid binaries in container (root only)
--app string          set an application to run inside a container
--apply-cgroups string apply cgroups from file for container processes (root_
↳only)
-B, --bind strings    a user-bind path specification. spec has the format_
↳src[:dest[:opts]], where src and dest are outside and inside paths. If dest is not_
↳given, it is set equal to src. Mount options ('opts') may be specified as 'ro'_
↳(read-only) or 'rw' (read/write, which is the default). Multiple bind paths can be_
↳given by a comma separated list.
-e, --cleanenv        clean environment before running container
-c, --contain         use minimal /dev and empty other directories (e.g. /tmp_
↳and $HOME) instead of sharing filesystems from your host
-C, --containall      contain not only file systems, but also PID, IPC, and_
↳environment
  --dns string        list of DNS server separated by commas to add in resolv.
↳conf
  --docker-login      login to a Docker Repository interactively
  --drop-caps string  a comma separated capability list to drop
-h, --help            help for exec
-H, --home string     a home directory specification. spec can either be a_
↳src path or src:dest pair. src is the source path of the home directory outside_
↳the container and dest overrides the home directory within the container. (default
↳"/root")
  --hostname string   set container hostname
-i, --ipc             run container in a new IPC namespace
  --keep-privs        let root user keep privileges in container (root only)
-n, --net             run container in a new network namespace (sets up a_
↳bridge network interface by default)
  --network string    specify desired network type separated by commas, each_
↳network will bring up a dedicated interface inside container (default "bridge")
  --network-args strings specify network arguments to pass to CNI plugins
  --no-home           do NOT mount users home directory if home is not the_
↳current working directory
  --no-init           do NOT start shim process with --pid
  --no-nv             drop all privileges from root user in container)
  --no-privs          do NOT use HTTPS, for communicating with local docker_
↳registry
  --nv               enable experimental Nvidia support
-o, --overlay strings use an overlayFS image for persistent data storage or as_
↳read-only layer of container
-p, --pid            run container in a new PID namespace
  --pwd string        initial working directory for payload process inside the_
↳container
-S, --scratch strings include a scratch directory within the container that is_
↳linked to a temporary dir (use -W to force location)

```

(continues on next page)

(continued from previous page)

```

--security strings      enable security features (SELinux, Apparmor, Seccomp)
-u, --usersns          run container in a new user namespace, allowing
↳ Singularity to run completely unprivileged on recent kernels. This disables some
↳ features of Singularity, for example it only works with sandbox images.
--uts                  run container in a new UTS namespace
--vm                   enable VM support
--vm-cpu string        Number of CPU cores to allocate to Virtual Machine
↳ (implies --vm) (default "1")
--vm-err               enable attaching stderr from VM
--vm-ram string        Amount of RAM in MiB to allocate to Virtual Machine
↳ (implies --vm) (default "1024")
-W, --workdir string   working directory to be used for /tmp, /var/tmp and
↳ $HOME (if -c/--contain was also used)
-w, --writable          by default all Singularity containers are available as
↳ read only. This option makes the file system accessible as read/write.
--writable-tmpfs       makes the file system accessible as read-write with non
↳ persistent data (with overlay support only)

```

## 4.12.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.13 singularity inspect

Show metadata for an image

### 4.13.1 Synopsis

Inspect will show you labels, environment variables, and scripts associated with the image determined by the flags you pass.

```
singularity inspect [inspect options...] <image path>
```

### 4.13.2 Examples

```
$ singularity inspect ubuntu.sif
```

### 4.13.3 Options

```

--app string          inspect a specific app
-d, --deffile         show the Singularity recipe file that was used to generate the
↳ image
-e, --environment     show the environment settings for the image

```

(continues on next page)

(continued from previous page)

```

-h, --help          help for inspect
-H, --helpfile     inspect the runscript helpfile, if it exists
-j, --json         print structured json instead of sections
-l, --labels       show the labels associated with the image (default)
  --list-apps      list all apps in a container
-r, --runscript    show the runscript for the image
-t, --test         show the test script for the image

```

#### 4.13.4 SEE ALSO

- [singularity](#) -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.14 singularity instance

Manage containers running as services

### 4.14.1 Synopsis

Instances allow you to run containers as background processes. This can be useful for running services such as web servers or databases.

```
singularity instance
```

### 4.14.2 Examples

All group commands have their own help output:

```

$ singularity help instance start
$ singularity instance start --help

```

### 4.14.3 Options

```
-h, --help  help for instance
```

#### 4.14.4 SEE ALSO

- [singularity](#) -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* [singularity instance list](#) - List all running and named Singularity instances \* [singularity instance start](#) - Start a named instance of the given container image \* [singularity instance stop](#) - Stop a named instance of a given container image

Auto generated by spf13/cobra on 23-May-2019

## 4.15 singularity instance list

List all running and named Singularity instances

### 4.15.1 Synopsis

The instance list command allows you to view the Singularity container instances that are currently running in the background.

```
singularity instance list [list options...]
```

### 4.15.2 Examples

```
$ singularity instance list
DAEMON NAME      PID      CONTAINER IMAGE
test             11963    /home/mibauer/singularity/sinstance/test.sif

$ sudo singularity instance list -u mibauer
DAEMON NAME      PID      CONTAINER IMAGE
test             11963    /home/mibauer/singularity/sinstance/test.sif
test2            16219    /home/mibauer/singularity/sinstance/test.sif
```

### 4.15.3 Options

```
-h, --help          help for list
-j, --json          Print structured json instead of list
-u, --user string   If running as root, list instances from "<username>"
```

### 4.15.4 SEE ALSO

- *singularity instance* - Manage containers running as services

Auto generated by spf13/cobra on 23-May-2019

## 4.16 singularity instance start

Start a named instance of the given container image

### 4.16.1 Synopsis

The instance start command allows you to create a new named instance from an existing container image that will begin running in the background. If a startscript is defined in the container metadata the commands in that script will be executed with the instance start command as well. You can optionally pass arguments to startscript

singularity instance start accepts the following container formats

- \*.sif Singularity Image Format (SIF). Native to Singularity 3.0+
- \*.sqsh SquashFS format. Native to Singularity 2.4+
- \*.img ext3 format. Native to Singularity versions < 2.4.

**directory/ sandbox format.** Directory containing a valid root file system and optionally Singularity meta-data.

**instance://\*** A local running instance of a container. (See the `instance` command group.)

**library://\*** A container hosted on a Library (default <https://cloud.sylabs.io/library>)

**docker://\*** A container hosted on Docker Hub

**shub://\*** A container hosted on Singularity Hub

```
singularity instance start [start options...] <container path> <instance name>
↳ [startscript args...]
```

## 4.16.2 Examples

```
$ singularity instance start /tmp/my-sql.sif mysql

$ singularity shell instance://mysql
Singularity my-sql.sif> pwd
/home/mibauer/mysql
Singularity my-sql.sif> ps
PID TTY          TIME CMD
  1 pts/0        00:00:00 sinit
  2 pts/0        00:00:00 bash
  3 pts/0        00:00:00 ps
Singularity my-sql.sif>

$ singularity instance stop /tmp/my-sql.sif mysql
Stopping /tmp/my-sql.sif mysql
```

## 4.16.3 Options

```
--add-caps string      a comma separated capability list to add
--allow-setuid         allow setuid binaries in container (root only)
--apply-cgroups string apply cgroups from file for container processes (root
↳ only)
-B, --bind strings    a user-bind path specification. spec has the format
↳ src[:dest[:opts]], where src and dest are outside and inside paths. If dest is not
↳ given, it is set equal to src. Mount options ('opts') may be specified as 'ro'
↳ (read-only) or 'rw' (read/write, which is the default). Multiple bind paths can be
↳ given by a comma separated list.
  --boot              execute /sbin/init to boot container (root only)
-e, --cleanenv        clean environment before running container
-c, --contain         use minimal /dev and empty other directories (e.g. /tmp
↳ and $HOME) instead of sharing filesystems from your host
-C, --containall      contain not only file systems, but also PID, IPC, and
↳ environment
  --dns string        list of DNS server separated by commas to add in resolv.
↳ conf
```

(continues on next page)

(continued from previous page)

```

--docker-login      login to a Docker Repository interactively
--drop-caps string  a comma separated capability list to drop
-h, --help          help for start
-H, --home string   a home directory specification.  spec can either be a
↳src path or src:dest pair.  src is the source path of the home directory outside
↳the container and dest overrides the home directory within the container. (default
↳"/root")
--hostname string   set container hostname
-i, --ipc           run container in a new IPC namespace
--keep-privs        let root user keep privileges in container (root only)
-n, --net           run container in a new network namespace (sets up a
↳bridge network interface by default)
--network string    specify desired network type separated by commas, each
↳network will bring up a dedicated interface inside container (default "bridge")
--network-args strings specify network arguments to pass to CNI plugins
--no-home           do NOT mount users home directory if home is not the
↳current working directory
--no-init           do NOT start shim process with --pid
--no-nv             do NOT start shim process with --nv
--no-privs          drop all privileges from root user in container)
--nohttps           do NOT use HTTPS, for communicating with local docker
↳registry
--nv               enable experimental Nvidia support
-o, --overlay strings use an overlayFS image for persistent data storage or as
↳read-only layer of container
-p, --pid           run container in a new PID namespace
-S, --scratch strings include a scratch directory within the container that is
↳linked to a temporary dir (use -W to force location)
--security strings enable security features (SELinux, Apparmor, Seccomp)
-u, --usersns      run container in a new user namespace, allowing
↳Singularity to run completely unprivileged on recent kernels. This disables some
↳features of Singularity, for example it only works with sandbox images.
--uts              run container in a new UTS namespace
-W, --workdir string working directory to be used for /tmp, /var/tmp and
↳$HOME (if -c/--contain was also used)
-w, --writable      by default all Singularity containers are available as
↳read only. This option makes the file system accessible as read/write.
--writable-tmpfs    makes the file system accessible as read-write with non
↳persistent data (with overlay support only)

```

#### 4.16.4 SEE ALSO

- *singularity instance* - Manage containers running as services

*Auto generated by spf13/cobra on 23-May-2019*

## 4.17 singularity instance stop

Stop a named instance of a given container image

### 4.17.1 Synopsis

The command `singularity instance stop` allows you to stop and clean up a named, running instance of a given container image.

```
singularity instance stop [stop options...] [instance]
```

### 4.17.2 Examples

```
$ singularity instance start my-sql.sif mysql1
$ singularity instance start my-sql.sif mysql2
$ singularity instance stop mysql*
Stopping mysql1 instance of my-sql.sif (PID=23845)
Stopping mysql2 instance of my-sql.sif (PID=23858)

$ singularity instance start my-sql.sif mysql1

Force instance to shutdown
$ singularity instance stop -f mysql1 (may corrupt data)

Send SIGTERM to the instance
$ singularity instance stop -s SIGTERM mysql1
$ singularity instance stop -s TERM mysql1
$ singularity instance stop -s 15 mysql1
```

### 4.17.3 Options

```
-a, --all           stop all user's instances
-F, --force        force kill instance
-h, --help        help for stop
-s, --signal string signal sent to the instance
-t, --timeout int  force kill non stopped instances after X seconds (default 10)
-u, --user string  If running as root, stop instances belonging to user
```

### 4.17.4 SEE ALSO

- *singularity instance* - Manage containers running as services

*Auto generated by spf13/cobra on 23-May-2019*

## 4.18 singularity key

Manage OpenPGP keys

### 4.18.1 Synopsis

Manage OpenPGP keys both locally via a Singularity keychain and remotely via a Sylabs Cloud Keystore. The default keyring is `~/singularity/syppg` if `SINGULARITY_SYPPGDIR` is not set.

```
singularity key [key options...]
```

## 4.18.2 Examples

All group commands have their own help output:

```
$ singularity help key newpair
$ singularity key list --help
```

## 4.18.3 Options

```
-h, --help  help for key
```

## 4.18.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* *singularity key export* - Export a public or private key into a specific file \* *singularity key import* - Import a local key into the local keyring \* *singularity key list* - List keys in your local keyring \* *singularity key newpair* - Create a new key pair \* *singularity key pull* - Download a public key from a key server \* *singularity key push* - Upload a public key to a key server \* *singularity key remove* - Remove a local public key from your keyring \* *singularity key search* - Search for keys on a key server

*Auto generated by spf13/cobra on 23-May-2019*

## 4.19 singularity key export

Export a public or private key into a specific file

### 4.19.1 Synopsis

The 'key export' command allows you to export a key and save it to a file.

```
singularity key export [export options...] <output-file>
```

### 4.19.2 Examples

```
Exporting a private key:
$ singularity key export --secret ./private.asc

Exporting a public key:
$ singularity key export ./public.asc
```



### 4.19.3 Options

```
-a, --armor      ascii armored format
-h, --help      help for export
-s, --secret     export a secret key
```

### 4.19.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.20 singularity key import

Import a local key into the local keyring

### 4.20.1 Synopsis

The 'key import' command allows you to add a key to your local keyring from a specific file.

```
singularity key import [import options...] <input-key>
```

### 4.20.2 Examples

```
$ singularity key import ./my-key.asc
```

### 4.20.3 Options

```
-h, --help      help for import
```

### 4.20.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.21 singularity key list

List keys in your local keyring

### 4.21.1 Synopsis

The ‘key list’ command allows you to list public/private keys from the default user’s local keyring location (i.e., \$HOME/.singularity/syppg).

```
singularity key list
```

### 4.21.2 Examples

```
$ singularity key list
```

### 4.21.3 Options

```
-h, --help      help for list  
-s, --secret    list private keys instead of the default which displays public ones
```

### 4.21.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.22 singularity key newpair

Create a new key pair

### 4.22.1 Synopsis

The ‘key newpair’ command allows you to create a new key or public/private keys to be stored in the default user local key store location (e.g., \$HOME/.singularity/syppg).

```
singularity key newpair
```

### 4.22.2 Examples

```
$ singularity key newpair
```

### 4.22.3 Options

```
-h, --help      help for newpair
```

## 4.22.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.23 singularity key pull

Download a public key from a key server

### 4.23.1 Synopsis

The ‘key pull’ command allows you to connect to a key server look for and download a public key. Key rings are stored into (e.g., \$HOME/.singularity/syppg).

```
singularity key pull [pull options...] <fingerprint>
```

### 4.23.2 Examples

```
$ singularity key pull 8883491F4268F173C6E5DC49EDECE4F3F38D871E
```

### 4.23.3 Options

```
-h, --help          help for pull
-u, --url string    specify the key server URL (default "https://keys.sylabs.io")
```

## 4.23.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.24 singularity key push

Upload a public key to a key server

### 4.24.1 Synopsis

The ‘key push’ command allows you to connect to a key server and upload public keys from the local key store.

```
singularity key push [push options...] <fingerprint>
```

## 4.24.2 Examples

```
$ singularity key push 8883491F4268F173C6E5DC49EDECE4F3F38D871E
```

## 4.24.3 Options

```
-h, --help      help for push  
-u, --url string specify the key server URL (default "https://keys.sylabs.io")
```

## 4.24.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.25 singularity key remove

Remove a local public key from your keyring

### 4.25.1 Synopsis

The ‘key remove’ command will remove a local public key.

```
singularity key remove <fingerprint>
```

### 4.25.2 Examples

```
$ singularity key remove D87FE3AF5C1F063FCBCC9B02F812842B5EEE5934
```

### 4.25.3 Options

```
-h, --help      help for remove
```

### 4.25.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.26 singularity key search

Search for keys on a key server

### 4.26.1 Synopsis

The 'key search' command allows you to connect to a key server and look for public keys matching the argument passed to the command line. You can also search for a key by fingerprint or key ID by adding '0x' before the fingerprint. (Maximum 100 search entities)

```
singularity key search [search options...] <search_string>
```

### 4.26.2 Examples

```
$ singularity key search sylabs.io

# note the '0x' before the fingerprint:
$ singularity key search 0x8883491F4268F173C6E5DC49EDECE4F3F38D871E

# search by key ID: (again, there's '0x' before the ID)
$ singularity key search 0xF38D871E
```

### 4.26.3 Options

```
-h, --help          help for search
-u, --url string    specify the key server URL (default "https://keys.sylabs.io")
```

### 4.26.4 SEE ALSO

- *singularity key* - Manage OpenPGP keys

*Auto generated by spf13/cobra on 23-May-2019*

## 4.27 singularity oci

Manage OCI containers

### 4.27.1 Synopsis

Allow you to manage containers from OCI bundle directories.

NOTE: all oci commands requires to run as root

### 4.27.2 Examples

```
All group commands have their own help output:

$ singularity oci create -b ~/bundle mycontainer
$ singularity oci start mycontainer
```

### 4.27.3 Options

```
-h, --help  help for oci
```

### 4.27.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* *singularity oci attach* - Attach console to a running container process (root user only) \* *singularity oci create* - Create a container from a bundle directory (root user only) \* *singularity oci delete* - Delete container (root user only) \* *singularity oci exec* - Execute a command within container (root user only) \* *singularity oci kill* - Kill a container (root user only) \* *singularity oci mount* - Mount create an OCI bundle from SIF image (root user only) \* *singularity oci pause* - Suspends all processes inside the container (root user only) \* *singularity oci resume* - Resumes all processes previously paused inside the container (root user only) \* *singularity oci run* - Create/start/attach/delete a container from a bundle directory (root user only) \* *singularity oci start* - Start container process (root user only) \* *singularity oci state* - Query state of a container (root user only) \* *singularity oci umount* - Umount delete bundle (root user only) \* *singularity oci update* - Update container cgroups resources (root user only)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.28 singularity oci attach

Attach console to a running container process (root user only)

### 4.28.1 Synopsis

Attach will attach console to a running container process running within container identified by container ID.

```
singularity oci attach <container_ID>
```

### 4.28.2 Examples

```
$ singularity oci attach mycontainer
```

### 4.28.3 Options

```
-h, --help  help for attach
```

### 4.28.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.29 singularity oci create

Create a container from a bundle directory (root user only)

### 4.29.1 Synopsis

Create invoke create operation to create a container instance from an OCI bundle directory

```
singularity oci create -b <bundle_path> [create options...] <container_ID>
```

### 4.29.2 Examples

```
$ singularity oci create -b ~/bundle mycontainer
```

### 4.29.3 Options

```
-b, --bundle string      specify the OCI bundle path (required)
  --empty-process        run container without executing container process (eg: for
↳POD container)
-h, --help              help for create
  --log-format string    specify the log file format. Available formats are basic,
↳kubernetes and json (default "kubernetes")
-l, --log-path string    specify the log file path
  --pid-file string      specify the pid file
-s, --sync-socket string specify the path to unix socket for state synchronization
```

### 4.29.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.30 singularity oci delete

Delete container (root user only)

### 4.30.1 Synopsis

Delete invoke delete operation to delete resources that were created for container identified by container ID.

```
singularity oci delete <container_ID>
```

### 4.30.2 Examples

```
$ singularity oci delete mycontainer
```

### 4.30.3 Options

```
-h, --help  help for delete
```

### 4.30.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.31 singularity oci exec

Execute a command within container (root user only)

### 4.31.1 Synopsis

Exec will execute the provided command/arguments within container identified by container ID.

```
singularity oci exec <container_ID> <command> <args>
```

### 4.31.2 Examples

```
$ singularity oci exec mycontainer id
```

### 4.31.3 Options

```
-h, --help  help for exec
```

### 4.31.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.32 singularity oci kill

Kill a container (root user only)



### 4.32.1 Synopsis

Kill invoke kill operation to kill processes running within container identified by container ID.

```
singularity oci kill [kill options...] <container_ID>
```

### 4.32.2 Examples

```
$ singularity oci kill mycontainer INT
$ singularity oci kill mycontainer -s INT
```

### 4.32.3 Options

```
-f, --force          kill container process with SIGKILL
-h, --help          help for kill
-s, --signal string signal sent to the container (default "SIGTERM")
-t, --timeout uint32 timeout in second before killing container
```

### 4.32.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.33 singularity oci mount

Mount create an OCI bundle from SIF image (root user only)

### 4.33.1 Synopsis

Mount will mount and create an OCI bundle from a SIF image.

```
singularity oci mount <sif_image> <bundle_path>
```

### 4.33.2 Examples

```
$ singularity oci mount /tmp/example.sif /var/lib/singularity/bundles/example
```

### 4.33.3 Options

```
-h, --help  help for mount
```

#### 4.33.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

### 4.34 singularity oci pause

Suspends all processes inside the container (root user only)

#### 4.34.1 Synopsis

Pause will suspend all processes for the specified container ID.

```
singularity oci pause <container_ID>
```

#### 4.34.2 Examples

```
$ singularity oci pause mycontainer
```

#### 4.34.3 Options

```
-h, --help  help for pause
```

#### 4.34.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

### 4.35 singularity oci resume

Resumes all processes previously paused inside the container (root user only)

#### 4.35.1 Synopsis

Resume will resume all processes previously paused for the specified container ID.

```
singularity oci resume <container_ID>
```

#### 4.35.2 Examples

```
$ singularity oci resume mycontainer
```

### 4.35.3 Options

```
-h, --help help for resume
```

### 4.35.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.36 singularity oci run

Create/start/attach/delete a container from a bundle directory (root user only)

### 4.36.1 Synopsis

Run will invoke equivalent of create/start/attach/delete commands in a row.

```
singularity oci run -b <bundle_path> [run options...] <container_ID>
```

### 4.36.2 Examples

```
$ singularity oci run -b ~/bundle mycontainer

is equivalent to :

$ singularity oci create -b ~/bundle mycontainer
$ singularity oci start mycontainer
$ singularity oci attach mycontainer
$ singularity oci delete mycontainer
```

### 4.36.3 Options

```
-b, --bundle string      specify the OCI bundle path (required)
-h, --help              help for run
  --log-format string    specify the log file format. Available formats are basic, ↵
↳kubernetes and json (default "kubernetes")
-l, --log-path string    specify the log file path
  --pid-file string      specify the pid file
-s, --sync-socket string specify the path to unix socket for state synchronization
```

### 4.36.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.37 singularity oci start

Start container process (root user only)

### 4.37.1 Synopsis

Start invoke start operation to start a previously created container identified by container ID.

```
singularity oci start <container_ID>
```

### 4.37.2 Examples

```
$ singularity oci start mycontainer
```

### 4.37.3 Options

```
-h, --help    help for start
```

### 4.37.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.38 singularity oci state

Query state of a container (root user only)

### 4.38.1 Synopsis

State invoke state operation to query state of a created/running/stopped container identified by container ID.

```
singularity oci state <container_ID>
```

### 4.38.2 Examples

```
$ singularity oci state mycontainer
```

### 4.38.3 Options

```
-h, --help                help for state  
-s, --sync-socket string  specify the path to unix socket for state synchronization
```

## 4.38.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.39 singularity oci umount

Umount delete bundle (root user only)

### 4.39.1 Synopsis

Umount will umount an OCI bundle previously mounted with singularity oci mount.

```
singularity oci umount <bundle_path>
```

### 4.39.2 Examples

```
$ singularity oci umount /var/lib/singularity/bundles/example
```

### 4.39.3 Options

```
-h, --help help for umount
```

## 4.39.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.40 singularity oci update

Update container cgroups resources (root user only)

### 4.40.1 Synopsis

Update will update cgroups resources for the specified container ID. Container must be in a RUNNING or CREATED state.

```
singularity oci update [update options...] <container_ID>
```

### 4.40.2 Examples

```
$ singularity oci update --from-file /tmp/cgroups-update.json mycontainer  
or to update from stdin :  
$ cat /tmp/cgroups-update.json | singularity oci update --from-file - mycontainer
```

### 4.40.3 Options

```
-f, --from-file string    specify path to OCI JSON cgroups resource file ('-' to read  
↳from STDIN)  
-h, --help                help for update
```

### 4.40.4 SEE ALSO

- *singularity oci* - Manage OCI containers

*Auto generated by spf13/cobra on 23-May-2019*

## 4.41 singularity plugin

Manage singularity plugins

### 4.41.1 Synopsis

The ‘plugin’ command allows you to manage

```
singularity plugin [plugin options...]
```

### 4.41.2 Examples

All group commands have their own help output:

```
$ singularity help plugin compile  
$ singularity plugin list --help
```

### 4.41.3 Options

```
-h, --help    help for plugin
```

### 4.41.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* *singularity plugin compile* - Compile a singularity plugin \* *singularity plugin disable* - disable an installed singularity plugin \* *singularity plugin enable* - Enable an installed singularity plugin \* *singularity plugin inspect* - Inspect a singularity plugin (either an installed one or an image) \* *singularity plugin install* - Install a singularity plugin \* *singularity plugin list* - List installed singularity plugins \* *singularity plugin uninstall* - Uninstall removes the named plugin from the system

*Auto generated by spf13/cobra on 23-May-2019*

## 4.42 singularity plugin compile

Compile a singularity plugin

### 4.42.1 Synopsis

The ‘plugin compile’ command allows a developer to compile a singularity plugin in the expected environment. The provided host directory is the location of the plugin’s source code.

```
singularity plugin compile [compile options...] <host_path>
```

### 4.42.2 Examples

```
$ singularity plugin compile $HOST_PATH
```

### 4.42.3 Options

```
-h, --help          help for compile
-o, --out string
```

### 4.42.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

## 4.43 singularity plugin disable

disable an installed singularity plugin

### 4.43.1 Synopsis

The ‘plugin disable’ command allows a user to disable a plugin that is already installed in the system and which has been previously enabled.

```
singularity plugin disable <name>
```

## 4.43.2 Examples

```
$ singularity plugin disable sylabs.io/test-plugin
```

## 4.43.3 Options

```
-h, --help  help for disable
```

## 4.43.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

## 4.44 singularity plugin enable

Enable an installed singularity plugin

### 4.44.1 Synopsis

The ‘plugin enable’ command allows a user to enable a plugin that is already installed in the system and which has been previously disabled.

```
singularity plugin enable <name>
```

### 4.44.2 Examples

```
$ singularity plugin enable sylabs.io/test-plugin
```

### 4.44.3 Options

```
-h, --help  help for enable
```

### 4.44.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

## 4.45 singularity plugin inspect

Inspect a singularity plugin (either an installed one or an image)



### 4.45.1 Synopsis

The ‘plugin inspect’ command allows a user to inspect a plugin that is already installed in the system or an image containing a plugin that is yet to be installed.

```
singularity plugin inspect (<name>|<image>)
```

### 4.45.2 Examples

```
$ singularity plugin inspect sylabs.io/test-plugin
$ singularity plugin inspect test-plugin.sif
```

### 4.45.3 Options

```
-h, --help    help for inspect
```

### 4.45.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

## 4.46 singularity plugin install

Install a singularity plugin

### 4.46.1 Synopsis

The ‘plugin install’ command installs the plugin found at plugin\_path into the appropriate directory on the host.

```
singularity plugin install [install options...] <plugin_path>
```

### 4.46.2 Examples

```
$ singularity plugin install $PLUGIN_PATH
```

### 4.46.3 Options

```
-h, --help          help for install
-n, --name string   Name to install the plugin as, defaults to the value in the u
↳manifest
```

#### 4.46.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

### 4.47 singularity plugin list

List installed singularity plugins

#### 4.47.1 Synopsis

The ‘plugin list’ command lists the singularity plugins installed on the host.

```
singularity plugin list [list options...]
```

#### 4.47.2 Examples

```
$ singularity plugin list
```

#### 4.47.3 Options

```
-h, --help help for list
```

#### 4.47.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

### 4.48 singularity plugin uninstall

Uninstall removes the named plugin from the system

#### 4.48.1 Synopsis

The ‘plugin uninstall’ command removes the named plugin from the system

```
singularity plugin uninstall <name>
```

#### 4.48.2 Examples

```
$ singularity plugin uninstall example.org/plugin
```

### 4.48.3 Options

```
-h, --help  help for uninstall
```

### 4.48.4 SEE ALSO

- *singularity plugin* - Manage singularity plugins

*Auto generated by spf13/cobra on 23-May-2019*

## 4.49 singularity pull

Pull an image from a URI

### 4.49.1 Synopsis

The ‘pull’ command allows you to download or build a container from a given URI. Supported URIs include:

**library:** Pull an image from the currently configured library `library://user/collection/container[:tag]`

**docker:** Pull an image from Docker Hub `docker://user/image:tag`

**shub:** Pull an image from Singularity Hub to CWD `shub://user/image:tag`

```
singularity pull [pull options...] [output file] <URI>
```

### 4.49.2 Examples

```
From Sylabs cloud library
$ singularity pull alpine.sif library://alpine:latest

From Docker
$ singularity pull tensorflow.sif docker://tensorflow/tensorflow:latest

From Shub
$ singularity pull singularity-images.sif shub://vsoch/singularity-images
```

### 4.49.3 Options

```
-U, --allow-unsigned  do not require a signed container
  --dir string        download images to the specific directory
  --docker-login      login to a Docker Repository interactively
-F, --force          overwrite an image file if it exists
-h, --help           help for pull
  --library string    download images from the provided library (default "https://library.sylabs.io")
  --no-cleanup        do NOT clean up bundle after failed build, can be helpul for
  --debugging         ↪ debugging
  --nohttps           do NOT use HTTPS with the docker:// transport (useful for
  ↪ local docker registries without a certificate) (continues on next page)
```

#### 4.49.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

### 4.50 singularity push

Upload image to the provided library (default is “cloud.sylabs.io”)

#### 4.50.1 Synopsis

The Singularity push command allows you to upload your sif image to a library of your choosing. It’s always good practice to sign your containers before pushing them to the library. An auth token is required to push to the remote, so you may need to configure if first with ‘singularity remote’.

```
singularity push [push options...] <image> library://user/collection/container[:tag]
```

#### 4.50.2 Examples

```
$ singularity push /home/user/my.sif library://user/collection/my.sif:latest
```

#### 4.50.3 Options

```
-U, --allow-unsigned  do not require a signed container  
-h, --help           help for push  
--library string     the library to push to (default "https://library.sylabs.io")
```

#### 4.50.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

### 4.51 singularity remote

Manage singularity remote endpoints

### 4.51.1 Synopsis

The ‘remote’ commands allow you to manage singularity remote endpoints through its subcommands. These allow you to add, log in, and use endpoints. The remote configuration is stored in `$HOME/.singularity/remotes.yaml` by default.

### 4.51.2 Examples

All group commands have their own help output:

```
$ singularity help remote list
$ singularity remote list
```

### 4.51.3 Options

```
-c, --config string  path to the file holding remote endpoint configurations_
↪ (default "/root/.singularity/remote.yaml")
-h, --help          help for remote
```

### 4.51.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC) \* *singularity remote add* - Create a new singularity remote endpoint \* *singularity remote list* - List all singularity remote endpoints that are configured \* *singularity remote login* - Log into a singularity remote endpoint using an authentication token \* *singularity remote remove* - Remove an existing singularity remote endpoint \* *singularity remote status* - Check the status of the singularity services at an endpoint \* *singularity remote use* - Set a singularity remote endpoint to be actively used

*Auto generated by spf13/cobra on 23-May-2019*

## 4.52 singularity remote add

Create a new singularity remote endpoint

### 4.52.1 Synopsis

The ‘remote add’ command allows you to create a new remote endpoint to be used for singularity remote services.

```
singularity remote add <remote_name> <remote_URI> [flags]
```

### 4.52.2 Examples

```
$ singularity remote add SylabsCloud cloud.sylabs.io
```

### 4.52.3 Options

```
-g, --global  edit the list of globally configured remote endpoints  
-h, --help   help for add
```

### 4.52.4 SEE ALSO

- *singularity remote* - Manage singularity remote endpoints

*Auto generated by spf13/cobra on 23-May-2019*

## 4.53 singularity remote list

List all singularity remote endpoints that are configured

### 4.53.1 Synopsis

The 'remote list' command lists all remote endpoints configured for use. If you have set a remote as a default, its name will be encompassed by brackets.

```
singularity remote list [flags]
```

### 4.53.2 Examples

```
$ singularity remote list
```

### 4.53.3 Options

```
-h, --help   help for list
```

### 4.53.4 SEE ALSO

- *singularity remote* - Manage singularity remote endpoints

*Auto generated by spf13/cobra on 23-May-2019*

## 4.54 singularity remote login

Log into a singularity remote endpoint using an authentication token

### 4.54.1 Synopsis

The ‘remote login’ command allows you to set an authentication token for a specific endpoint. This command will produce a link directing you to the token service you can use to generate a valid token. If no endpoint is specified, it will try the default remote.

```
singularity remote login <remote_name> [flags]
```

### 4.54.2 Examples

```
$ singularity remote login SylabsCloud
```

### 4.54.3 Options

```
-h, --help          help for login
--tokenfile string  path to the file holding token
```

### 4.54.4 SEE ALSO

- *singularity remote* - Manage singularity remote endpoints

*Auto generated by spf13/cobra on 23-May-2019*

## 4.55 singularity remote remove

Remove an existing singularity remote endpoint

### 4.55.1 Synopsis

The ‘remote remove’ command allows you to remove an existing remote endpoint from the list of potential endpoints to use.

```
singularity remote remove <remote_name> [flags]
```

### 4.55.2 Examples

```
$ singularity remote remove SylabsCloud
```

### 4.55.3 Options

```
-g, --global  edit the list of globally configured remote endpoints
-h, --help    help for remove
```

## 4.55.4 SEE ALSO

- *singularity remote* - Manage singularity remote endpoints

*Auto generated by spf13/cobra on 23-May-2019*

## 4.56 singularity remote status

Check the status of the singularity services at an endpoint

### 4.56.1 Synopsis

the ‘remote status’ command checks the status of the specified remote endpoint and reports the availability of services and their versions. If no endpoint is specified, it will check the status of the default remote.

```
singularity remote status <remote_name> [flags]
```

### 4.56.2 Examples

```
$ singularity remote status SylabsCloud
```

### 4.56.3 Options

```
-h, --help  help for status
```

## 4.56.4 SEE ALSO

- *singularity remote* - Manage singularity remote endpoints

*Auto generated by spf13/cobra on 23-May-2019*

## 4.57 singularity remote use

Set a singularity remote endpoint to be actively used

### 4.57.1 Synopsis

The ‘remote use’ command sets the remote to be used by default by any command that interacts with singularity services.

```
singularity remote use <remote_name> [flags]
```



## 4.57.2 Examples

```
$ singularity remote use SylabsCloud
```

## 4.57.3 Options

```
-g, --global  edit the list of globally configured remote endpoints
-h, --help    help for use
```

## 4.57.4 SEE ALSO

- *singularity remote* - Manage singularity remote endpoints

*Auto generated by spf13/cobra on 23-May-2019*

## 4.58 singularity run

Run the user-defined default command within a container

### 4.58.1 Synopsis

This command will launch a Singularity container and execute a runscrip if one is defined for that container. The runscrip is a metadata file within the container that contains shell commands. If the file is present (and executable) then this command will execute that file within the container automatically. All arguments following the container name will be passed directly to the runscrip.

singularity run accepts the following container formats:

\*.sif Singularity Image Format (SIF). Native to Singularity 3.0+

\*.sqsh SquashFS format. Native to Singularity 2.4+

\*.img ext3 format. Native to Singularity versions < 2.4.

**directory/ sandbox format.** Directory containing a valid root file system and optionally Singularity meta-data.

**instance://\*** A local running instance of a container. (See the `instance` command group.)

**library://\*** A container hosted on a Library (default <https://cloud.sylabs.io/library>)

**docker://\*** A container hosted on Docker Hub

**shub://\*** A container hosted on Singularity Hub

```
singularity run [run options...] <container>
```

### 4.58.2 Examples

```
# Here we see that the runscript prints "Hello world: "
$ singularity exec /tmp/debian.sif cat /singularity
#!/bin/sh
echo "Hello world: "

# It runs with our inputs when we run the image
$ singularity run /tmp/debian.sif one two three
Hello world: one two three

# Note that this does the same thing
$ ./tmp/debian.sif one two three
```

### 4.58.3 Options

```
--add-caps string      a comma separated capability list to add
--allow-setuid         allow setuid binaries in container (root only)
--app string          set an application to run inside a container
--apply-cgroups string apply cgroups from file for container processes (root_
↳only)
-B, --bind strings    a user-bind path specification. spec has the format_
↳src[:dest[:opts]], where src and dest are outside and inside paths. If dest is not_
↳given, it is set equal to src. Mount options ('opts') may be specified as 'ro'_
↳(read-only) or 'rw' (read/write, which is the default). Multiple bind paths can be_
↳given by a comma separated list.
-e, --cleanenv       clean environment before running container
-c, --contain        use minimal /dev and empty other directories (e.g. /tmp_
↳and $HOME) instead of sharing filesystems from your host
-C, --containall     contain not only file systems, but also PID, IPC, and_
↳environment
  --dns string        list of DNS server separated by commas to add in resolv_
↳conf
  --docker-login      login to a Docker Repository interactively
  --drop-caps string  a comma separated capability list to drop
-h, --help           help for run
-H, --home string    a home directory specification. spec can either be a_
↳src path or src:dest pair. src is the source path of the home directory outside_
↳the container and dest overrides the home directory within the container. (default_
↳"/root")
  --hostname string   set container hostname
-i, --ipc            run container in a new IPC namespace
  --keep-privs        let root user keep privileges in container (root only)
-n, --net           run container in a new network namespace (sets up a_
↳bridge network interface by default)
  --network string    specify desired network type separated by commas, each_
↳network will bring up a dedicated interface inside container (default "bridge")
  --network-args strings specify network arguments to pass to CNI plugins
  --no-home          do NOT mount users home directory if home is not the_
↳current working directory
  --no-init          do NOT start shim process with --pid
  --no-nv            drop all privileges from root user in container)
  --no-privs        do NOT use HTTPS, for communicating with local docker_
↳registry
  --nv              enable experimental Nvidia support
-o, --overlay strings use an overlayFS image for persistent data storage or as_
↳read-only layer of container
```

(continues on next page)

(continued from previous page)

```

-p, --pid                run container in a new PID namespace
  --pwd string           initial working directory for payload process inside the
↳container
-S, --scratch strings    include a scratch directory within the container that is
↳linked to a temporary dir (use -W to force location)
  --security strings     enable security features (SELinux, Apparmor, Seccomp)
-u, --usersns            run container in a new user namespace, allowing
↳Singularity to run completely unprivileged on recent kernels. This disables some
↳features of Singularity, for example it only works with sandbox images.
  --uts                 run container in a new UTS namespace
  --vm                 enable VM support
  --vm-cpu string       Number of CPU cores to allocate to Virtual Machine
↳(implies --vm) (default "1")
  --vm-err              enable attaching stderr from VM
  --vm-ram string       Amount of RAM in MiB to allocate to Virtual Machine
↳(implies --vm) (default "1024")
-W, --workdir string     working directory to be used for /tmp, /var/tmp and
↳$HOME (if -c/--contain was also used)
-w, --writable           by default all Singularity containers are available as
↳read only. This option makes the file system accessible as read/write.
  --writable-tmpfs       makes the file system accessible as read-write with non
↳persistent data (with overlay support only)

```

#### 4.58.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.59 singularity run-help

Show the user-defined help for an image

### 4.59.1 Synopsis

Show the help for an image.

The help text is from the ‘%help’ section of the definition file. If you are using the ‘-apps’ option, the help text is instead from that app’s ‘%apphelp’ section.

```
singularity run-help <image path>
```

### 4.59.2 Examples

```

$ cat my_container.def
Bootstrap: docker
From: busybox

```

(continues on next page)

(continued from previous page)

```
%help
    Some help for this container

%apphelp foo
    Some help for application 'foo' in this container

$ sudo singularity build my_container.sif my_container.def
Using container recipe deffile: my_container.def
[...snip...]
Cleaning up...

$ singularity run-help my_container.sif

    Some help for this container

$ singularity run-help --app foo my_container.sif

    Some help for application in this container
```

### 4.59.3 Options

```
--app string  Show the help for an app
-h, --help    help for run-help
```

### 4.59.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.60 singularity search

Search a Library for images

### 4.60.1 Synopsis

The Singularity search command allows you to search within a container library of your choosing. The container library defaults to cloud.sylabs.io when no other library argument is given.

```
singularity search [search options...] <search query>
```

### 4.60.2 Examples

```
$ singularity search lolcow
```

### 4.60.3 Options

```
-h, --help          help for search
  --library string  URI for library to search (default "https://library.sylabs.io")
```

### 4.60.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.61 singularity shell

Run a shell within a container

### 4.61.1 Synopsis

singularity shell supports the following formats:

- \*.sif Singularity Image Format (SIF). Native to Singularity 3.0+
- \*.sqsh SquashFS format. Native to Singularity 2.4+
- \*.img ext3 format. Native to Singularity versions < 2.4.

**directory/ sandbox format.** Directory containing a valid root file system and optionally Singularity meta-data.

**instance://\*** A local running instance of a container. (See the **instance** command group.)

**library://\*** A container hosted on a Library (default <https://cloud.sylabs.io/library>)

**docker://\*** A container hosted on Docker Hub

**shub://\*** A container hosted on Singularity Hub

```
singularity shell [shell options...] <container>
```

### 4.61.2 Examples

```
$ singularity shell /tmp/Debian.sif
Singularity/Debian.sif> pwd
/home/gmk/test
Singularity/Debian.sif> exit

$ singularity shell -C /tmp/Debian.sif
```

(continues on next page)

(continued from previous page)

```

Singularity/Debian.sif> pwd
/home/gmk
Singularity/Debian.sif> ls -l
total 0
Singularity/Debian.sif> exit

$ sudo singularity shell -w /tmp/Debian.sif
$ sudo singularity shell --writable /tmp/Debian.sif

$ singularity shell instance://my_instance

$ singularity shell instance://my_instance
Singularity: Invoking an interactive shell within container...
Singularity container:~> ps -ef
UID          PID  PPID  C  STIME TTY          TIME CMD
ubuntu        1    0    0  20:00 ?           00:00:00 /usr/local/bin/singularity/bin/sinit
ubuntu        2    0    0  20:01 pts/8       00:00:00 /bin/bash --norc
ubuntu        3    2    0  20:02 pts/8       00:00:00 ps -ef

```

### 4.61.3 Options

```

--add-caps string      a comma separated capability list to add
--allow-setuid         allow setuid binaries in container (root only)
--app string          set an application to run inside a container
--apply-cgroups string apply cgroups from file for container processes (root_
↳only)
-B, --bind strings    a user-bind path specification. spec has the format_
↳src[:dest[:opts]], where src and dest are outside and inside paths. If dest is not_
↳given, it is set equal to src. Mount options ('opts') may be specified as 'ro'_
↳(read-only) or 'rw' (read/write, which is the default). Multiple bind paths can be_
↳given by a comma separated list.
-e, --cleanenv        clean environment before running container
-c, --contain         use minimal /dev and empty other directories (e.g. /tmp_
↳and $HOME) instead of sharing filesystems from your host
-C, --containall     contain not only file systems, but also PID, IPC, and_
↳environment
  --dns string        list of DNS server separated by commas to add in resolv_
↳conf
  --docker-login      login to a Docker Repository interactively
  --drop-caps string  a comma separated capability list to drop
-h, --help            help for shell
-H, --home string     a home directory specification. spec can either be a_
↳src path or src:dest pair. src is the source path of the home directory outside_
↳the container and dest overrides the home directory within the container. (default_
↳"/root")
  --hostname string   set container hostname
-i, --ipc             run container in a new IPC namespace
  --keep-privs        let root user keep privileges in container (root only)
-n, --net             run container in a new network namespace (sets up a_
↳bridge network interface by default)
  --network string    specify desired network type separated by commas, each_
↳network will bring up a dedicated interface inside container (default "bridge")
  --network-args strings specify network arguments to pass to CNI plugins
  --no-home           do NOT mount users home directory if home is not the_
↳current working directory

```

(continues on next page)

(continued from previous page)

```

--no-init          do NOT start shim process with --pid
--no-nv
--no-privs        drop all privileges from root user in container)
--nohttps        do NOT use HTTPS, for communicating with local docker_
↪registry
  --nv           enable experimental Nvidia support
-o, --overlay strings use an overlayFS image for persistent data storage or as_
↪read-only layer of container
-p, --pid        run container in a new PID namespace
  --pwd string   initial working directory for payload process inside the_
↪container
-S, --scratch strings include a scratch directory within the container that is_
↪linked to a temporary dir (use -W to force location)
  --security strings enable security features (SELinux, Apparmor, Seccomp)
-s, --shell string path to program to use for interactive shell
  --syos        execute SyOS shell
-u, --usersns    run container in a new user namespace, allowing_
↪Singularity to run completely unprivileged on recent kernels. This disables some_
↪features of Singularity, for example it only works with sandbox images.
  --uts         run container in a new UTS namespace
  --vm         enable VM support
  --vm-cpu string Number of CPU cores to allocate to Virtual Machine_
↪(implies --vm) (default "1")
  --vm-err      enable attaching stderr from VM
  --vm-ram string Amount of RAM in MiB to allocate to Virtual Machine_
↪(implies --vm) (default "1024")
-W, --workdir string working directory to be used for /tmp, /var/tmp and
↪$HOME (if -c/--contain was also used)
-w, --writable    by default all Singularity containers are available as_
↪read only. This option makes the file system accessible as read/write.
  --writable-tmpfs makes the file system accessible as read-write with non_
↪persistent data (with overlay support only)

```

#### 4.61.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.62 singularity sign

Attach a cryptographic signature to an image

### 4.62.1 Synopsis

The sign command allows a user to create a cryptographic signature on either a single data object or a list of data objects within the same SIF group. By default without parameters, the command searches for the primary partition and creates a verification block that is then added to the SIF container file.

To generate a keypair, see ‘singularity help key newpair’

```
singularity sign [sign options...] <image path>
```

## 4.62.2 Examples

```
$ singularity sign container.sif
```

## 4.62.3 Options

```
-g, --groupid uint32  group ID to be signed
-h, --help           help for sign
-i, --id uint32      descriptor ID to be signed
-k, --keyidx int     private key to use (index from 'keys list') (default -1)
-u, --url string     key server URL (default "https://keys.sylabs.io")
```

## 4.62.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.63 singularity test

Run the user-defined tests within a container

### 4.63.1 Synopsis

The 'test' command allows you to execute a testscript (if available) inside of a given container

**NOTE:** For instances if there is a daemon process running inside the container, then subsequent container commands will all run within the same namespaces. This means that the `--writable` and `--contain` options will not be honored as the namespaces have already been configured by the 'singularity start' command.

```
singularity test [exec options...] <image path>
```

### 4.63.2 Examples

```
Set the '%test' section with a definition file like so:
%test
  echo "hello from test" "$@"

$ singularity test /tmp/debian.sif command
hello from test command
```

(continues on next page)



(continued from previous page)

For additional help, please visit our public documentation pages which are found at:

<https://www.sylabs.io/docs/>

### 4.63.3 Options

```

--add-caps string      a comma separated capability list to add
--allow-setuid        allow setuid binaries in container (root only)
--app string          set an application to run inside a container
--apply-cgroups string apply cgroups from file for container processes (root_
↳only)
-B, --bind strings    a user-bind path specification. spec has the format_
↳src[:dest[:opts]], where src and dest are outside and inside paths. If dest is not_
↳given, it is set equal to src. Mount options ('opts') may be specified as 'ro'_
↳(read-only) or 'rw' (read/write, which is the default). Multiple bind paths can be_
↳given by a comma separated list.
-e, --cleanenv        clean environment before running container
-c, --contain         use minimal /dev and empty other directories (e.g. /tmp_
↳and $HOME) instead of sharing filesystems from your host
-C, --containall      contain not only file systems, but also PID, IPC, and_
↳environment
  --dns string        list of DNS server separated by commas to add in resolv_
↳conf
  --docker-login      login to a Docker Repository interactively
  --drop-caps string  a comma separated capability list to drop
-h, --help            help for test
-H, --home string     a home directory specification. spec can either be a_
↳src path or src:dest pair. src is the source path of the home directory outside_
↳the container and dest overrides the home directory within the container. (default
↳"/root")
  --hostname string   set container hostname
-i, --ipc             run container in a new IPC namespace
  --keep-privs        let root user keep privileges in container (root only)
-n, --net            run container in a new network namespace (sets up a_
↳bridge network interface by default)
  --network string    specify desired network type separated by commas, each_
↳network will bring up a dedicated interface inside container (default "bridge")
  --network-args strings specify network arguments to pass to CNI plugins
  --no-home           do NOT mount users home directory if home is not the_
↳current working directory
  --no-init           do NOT start shim process with --pid
  --no-nv             do NOT start shim process with --nv
  --no-privs          drop all privileges from root user in container)
  --nohttps           do NOT use HTTPS, for communicating with local docker_
↳registry
  --nv               enable experimental Nvidia support
-o, --overlay strings use an overlayFS image for persistent data storage or as_
↳read-only layer of container
-p, --pid            run container in a new PID namespace
  --pwd string        initial working directory for payload process inside the_
↳container
-S, --scratch strings include a scratch directory within the container that is_
↳linked to a temporary dir (use -W to force location)
  --security strings  enable security features (SELinux, Apparmor, Seccomp)

```

(continues on next page)

(continued from previous page)

```

-u, --usersns          run container in a new user namespace, allowing
↳Singularity to run completely unprivileged on recent kernels. This disables some
↳features of Singularity, for example it only works with sandbox images.
  --uts                run container in a new UTS namespace
  --vm                 enable VM support
  --vm-cpu string      Number of CPU cores to allocate to Virtual Machine
↳(implies --vm) (default "1")
  --vm-err             enable attaching stderr from VM
  --vm-ram string      Amount of RAM in MiB to allocate to Virtual Machine
↳(implies --vm) (default "1024")
-W, --workdir string  working directory to be used for /tmp, /var/tmp and
↳$HOME (if -c/--contain was also used)
-w, --writable         by default all Singularity containers are available as
↳read only. This option makes the file system accessible as read/write.
  --writable-tmpfs     makes the file system accessible as read-write with non
↳persistent data (with overlay support only)

```

## 4.63.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.64 singularity verify

Verify cryptographic signatures attached to an image

### 4.64.1 Synopsis

The verify command allows a user to verify cryptographic signatures on SIF container files. There may be multiple signatures for data objects and multiple data objects signed. By default the command searches for the primary partition signature. If found, a list of all verification blocks applied on the primary partition is gathered so that data integrity (hashing) and signature verification is done for all those blocks.

```
singularity verify [verify options...] <image path>
```

### 4.64.2 Examples

```
$ singularity verify container.sif
```

### 4.64.3 Options

```
-g, --groupid uint32  group ID to be verified
-h, --help           help for verify
-i, --id uint32     descriptor ID to be verified
-l, --local         only verify with local keys
-u, --url string    key server URL (default "https://keys.sylabs.io")
```

#### 4.64.4 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*

## 4.65 singularity version

Show the version for Singularity

### 4.65.1 Synopsis

Show the version for Singularity

```
singularity version
```

### 4.65.2 Options

```
-h, --help  help for version
```

#### 4.65.3 SEE ALSO

- *singularity* -

Linux container platform optimized for High Performance Computing (HPC) and Enterprise Performance Computing (EPC)

*Auto generated by spf13/cobra on 23-May-2019*



## BUILD A CONTAINER

`build` is the “Swiss army knife” of container creation. You can use it to download and assemble existing containers from external resources like the [Container Library](#) and [Docker Hub](#). You can use it to convert containers between the formats supported by Singularity. And you can use it in conjunction with a Singularity definition file to create a container from scratch and customized it to fit your needs.

### 5.1 Overview

The `build` command accepts a target as input and produces a container as output.

The target defines the method that `build` uses to create the container. It can be one of the following:

- URI beginning with **library://** to build from the Container Library
- URI beginning with **docker://** to build from Docker Hub
- URI beginning with **shub://** to build from Singularity Hub
- path to a **existing container** on your local machine
- path to a **directory** to build from a sandbox
- path to a Singularity definition file

`build` can produce containers in two different formats that can be specified as follows.

- compressed read-only **Singularity Image File (SIF)** format suitable for production (default)
- writable **(ch)root directory** called a sandbox for interactive development ( `--sandbox` option)

Because `build` can accept an existing container as a target and create a container in either supported format you can convert existing containers from one format to another.

### 5.2 Downloading an existing container from the Container Library

You can use the `build` command to download a container from the Container Library.

```
$ sudo singularity build lolcow.simg library://sylabs-jms/testing/lolcow
```

The first argument (`lolcow.simg`) specifies a path and name for your container. The second argument (`library://sylabs-jms/testing/lolcow`) gives the Container Library URI from which to download. By default the container will be converted to a compressed, read-only SIF. If you want your container in a writable format use the `--sandbox` option.

## 5.3 Downloading an existing container from Docker Hub

You can use `build` to download layers from Docker Hub and assemble them into Singularity containers.

```
$ sudo singularity build lolcow.sif docker://godlovedc/lolcow
```

## 5.4 Creating writable `--sandbox` directories

If you wanted to create a container within a writable directory (called a sandbox) you can do so with the `--sandbox` option. It's possible to create a sandbox without root privileges, but to ensure proper file permissions it is recommended to do so as root.

```
$ sudo singularity build --sandbox lolcow/ library://sylabs-jms/testing/lolcow
```

The resulting directory operates just like a container in a SIF file. To make changes within the container, use the `--writable` flag when you invoke your container. It's a good idea to do this as root to ensure you have permission to access the files and directories that you want to change.

```
$ sudo singularity shell --writable lolcow/
```

## 5.5 Converting containers from one format to another

If you already have a container saved locally, you can use it as a target to build a new container. This allows you convert containers from one format to another. For example if you had a sandbox container called `development/` and you wanted to convert it to SIF container called `production.sif` you could:

```
$ sudo singularity build production.sif development/
```

Use care when converting a sandbox directory to the default SIF format. If changes were made to the writable container before conversion, there is no record of those changes in the Singularity definition file rendering your container non-reproducible. It is a best practice to build your immutable production containers directly from a Singularity definition file instead.

## 5.6 Building containers from Singularity definition files

Of course, Singularity definition files can be used as the target when building a container. For detailed information on writing Singularity definition files, please see the [Container Definition docs](#). Let's say you already have the following container definition file called `lolcow.def`, and you want to use it to build a SIF container.

```
Bootstrap: docker
From: ubuntu:16.04

%post
    apt-get -y update
    apt-get -y install fortune cowsay lolcat

%environment
    export LC_ALL=C
    export PATH=/usr/games:$PATH
```

(continues on next page)

(continued from previous page)

```
%runscript
  fortune | cowsay | lolcat
```

You can do so with the following command.

```
$ sudo singularity build lolcow.sif lolcow.def
```

The command requires `sudo` just as installing software on your local machine requires root privileges.

---

**Note:** Beware that it is possible to build an image on a host and have the image not work on a different host. This could be because of the default compressor supported by the host. For example, when building an image on a host in which the default compressor is `xz` and then trying to run that image on a CentOS 6 node, where the only compressor available is `gzip`.

---

## 5.7 Build options

### 5.7.1 `--builder`

Singularity 3.0 introduces the option to perform a remote build. The `--builder` option allows you to specify a URL to a different build service. For instance, you may need to specify a URL to build to an on premises installation of the remote builder. This option must be used in conjunction with `--remote`.

### 5.7.2 `--detached`

When used in combination with the `--remote` option, the `--detached` option will detach the build from your terminal and allow it to build in the background without echoing any output to your terminal.

### 5.7.3 `--force`

The `--force` option will delete and overwrite an existing Singularity image without presenting the normal interactive prompt.

### 5.7.4 `--json`

The `--json` option will force Singularity to interpret a given definition file as a json.

### 5.7.5 `--library`

This command allows you to set a different library. (The default library is "<https://library.sylabs.io>")

### 5.7.6 `--notest`

If you don't want to run the `%test` section during the container build, you can skip it with the `--notest` option. For instance, maybe you are building a container intended to run in a production environment with GPUs. But perhaps your local build resource does not have GPUs. You want to include a `%test` section that runs a short validation but you don't want your build to exit with an error because it cannot find a GPU on your system.

### 5.7.7 `--remote`

Singularity 3.0 introduces the ability to build a container on an external resource running a remote builder. (The default remote builder is located at "<https://cloud.sylabs.io/builder>".)

### 5.7.8 `--sandbox`

Build a sandbox (chroot directory) instead of the default SIF format.

### 5.7.9 `--section`

Instead of running the entire definition file, only run a specific section or sections. This option accepts a comma delimited string of definition file sections. Acceptable arguments include `all`, `none` or any combination of the following: `setup`, `post`, `files`, `environment`, `test`, `labels`.

Under normal build conditions, the Singularity definition file is saved into a container's meta-data so that there is a record showing how the container was built. Using the `--section` option may render this meta-data useless, so use care if you value reproducibility.

### 5.7.10 `--update`

You can build into the same sandbox container multiple times (though the results may be unpredictable and it is generally better to delete your container and start from scratch).

By default if you build into an existing sandbox container, the `build` command will prompt you to decide whether or not to overwrite the container. Instead of this behavior you can use the `--update` option to build `_into_` an existing container. This will cause Singularity to skip the header and build any sections that are in the definition file into the existing container.

The `--update` option is only valid when used with sandbox containers.

## 5.8 More Build topics

- If you want to **customize the cache location** (where Docker layers are downloaded on your system), specify Docker credentials, or any custom tweaks to your build environment, see [build environment](#).
- If you want to make internally **modular containers**, check out the getting started guide [here](#)
- If you want to **build your containers** on the Remote Builder, (because you don't have root access on a Linux machine or want to host your container on the cloud) check out [this site](#)



## DEFINITION FILES

A Singularity Definition File (or “def file” for short) is like a set of blueprints explaining how to build a custom container. It includes specifics about the base OS to build or the base container to start from, software to install, environment variables to set at runtime, files to add from the host system, and container metadata.

### 6.1 Overview

A Singularity Definition file is divided into two parts:

1. **Header:** The Header describes the core operating system to build within the container. Here you will configure the base operating system features needed within the container. You can specify, the Linux distribution, the specific version, and the packages that must be part of the core install (borrowed from the host system).
2. **Sections:** The rest of the definition is comprised of sections, (sometimes called scriptlets or blobs of data). Each section is defined by a % character followed by the name of the particular section. All sections are optional, and a def file may contain more than one instance of a given section. Sections that are executed at build time are executed with the `/bin/sh` interpreter and can accept `/bin/sh` options. Similarly, sections that produce scripts to be executed at runtime can accept options intended for `/bin/sh`

For more in-depth and practical examples of def files, see the [Sylabs examples repository](#)

### 6.2 Header

The header should be written at the top of the def file. It tells Singularity about the base operating system that it should use to build the container. It is composed of several keywords.

The only keyword that is required for every type of build is `Bootstrap`. It determines the *bootstrap agent* that will be used to create the base operating system you want to use. For example, the `library` bootstrap agent will pull a container from the [Container Library](#) as a base. Similarly, the `docker` bootstrap agent will pull docker layers from [Docker Hub](#) as a base OS to start your image.

Starting with Singularity 3.2, the `Bootstrap` keyword needs to be the first entry in the header section. This breaks compatibility with older versions that allow the parameters of the header to appear in any order.

Depending on the value assigned to `Bootstrap`, other keywords may also be valid in the header. For example, when using the `library` bootstrap agent, the `From` keyword becomes valid. Observe the following example for building a Debian container from the Container Library:

```
Bootstrap: library
From: debian:7
```

A def file that uses an official mirror to install Centos-7 might look like this:

```
Bootstrap: yum
OSVersion: 7
MirrorURL: http://mirror.centos.org/centos-%{OSVERSION}/%{OSVERSION}/os/$basearch/
Include: yum
```

Each bootstrap agent enables its own options and keywords. You can read about them and see examples in the appendix:

- *library* (images hosted on the Container Library)
- *docker* (images hosted on Docker Hub)
- *shub* (images hosted on Singularity Hub)
- *localimage* (images saved on your machine)
- *yum* (yum based systems such as CentOS and Scientific Linux)
- *debootstrap* (apt based systems such as Debian and Ubuntu)
- *arch* (Arch Linux)
- *busybox* (BusyBox)
- *zypper* (zypper based systems such as Suse and OpenSuse)

## 6.3 Sections

The main content of the bootstrap file is broken into sections. Different sections add different content or execute commands at different times during the build process. Note that if any command fails, the build process will halt.

Here is an example definition file that uses every available section. We will discuss each section in turn. It is not necessary to include every section (or any sections at all) within a def file. Furthermore, the order of the sections in the def file is unimportant and multiple sections of the same name can be included and will be appended to one another during the build process.

```
Bootstrap: library
From: ubuntu:18.04
Stage: build

%setup
    touch /file1
    touch ${SINGULARITY_ROOTFS}/file2

%files
    /file1
    /file1 /opt

%environment
    export LISTEN_PORT=12345
    export LC_ALL=C

%post
    apt-get update && apt-get install -y netcat
    NOW=`date`
    echo "export NOW=\"${NOW}\"" >> $SINGULARITY_ENVIRONMENT

%runscript
```

(continues on next page)

(continued from previous page)

```

echo "Container was created $NOW"
echo "Arguments received: $*"
exec echo "$@"

%startscript
nc -lp $LISTEN_PORT

%test
grep -q NAME="Ubuntu\" /etc/os-release
if [ $? -eq 0 ]; then
    echo "Container base is Ubuntu as expected."
else
    echo "Container base is not Ubuntu."
fi

%labels
Author d@sylabs.io
Version v0.0.1

%help
This is a demo container used to illustrate a def file that uses all
supported sections.

```

### 6.3.1 %setup

Commands in the %setup section are executed on the host system outside of the container after the base OS has been installed. You can reference the container file system with the \$SINGULARITY\_ROOTFS environment variable in the %setup section.

---

**Note:** Be careful with the %setup section! This scriptlet is executed outside of the container on the host system itself, and is executed with elevated privileges. Commands in %setup can alter and potentially damage the host.

---

Consider the example from the definition file above:

```

%setup
touch /file1
touch ${SINGULARITY_ROOTFS}/file2

```

Here, file1 is created at the root of the file system **on the host**. We'll use file1 to demonstrate the usage of the %files section below. The file2 is created at the root of the file system **within the container**.

In later versions of Singularity the %files section is provided as a safer alternative to copying files from the host system into the container during the build. Because of the potential danger involved in running the %setup scriptlet with elevated privileges on the host system during the build, it's use is generally discouraged.

### 6.3.2 %files

The %files section allows you to copy files from your host system into the container with greater safety than using the %setup section. Each line is a <source> and <destination> pair, where the source is a path on your host system, and the destination is a path in the container. The <destination> specification can be omitted and will be assumed to be the same path as the <source> specification.

Consider the example from the definition file above:

```
%files
  /file1
  /file1 /opt
```

`file1` was created in the root of the host file system during the `%setup` section (see above). The `%files` scriptlet will copy `file1` to the root of the container file system and then make a second copy of `file1` within the container in `/opt`.

Files can be copied from other stages by providing the source location in the previous stage and the destination in the current container.

```
%files from stage_name
  /root/hello /bin/hello
```

Files in the `%files` section are copied before the `%post` section is executed so that they are available during the build and configuration process.

### 6.3.3 %environment

The `%environment` section allows you to define environment variables that will be set at runtime. Note that these variables are not made available at build time by their inclusion in the `%environment` section. This means that if you need the same variables during the build process, you should also define them in your `%post` section. Specifically:

- **during build:** The `%environment` section is written to a file in the container metadata directory. This file is not sourced.
- **during runtime:** The file in the container metadata directory is sourced.

You should use the same conventions that you would use in a `.bashrc` or `.profile` file. Consider this example from the `def` file above:

```
%environment
  export LISTEN_PORT=12345
  export LC_ALL=C
```

The `$LISTEN_PORT` variable will be used in the `%startscript` section below. The `$LC_ALL` variable is useful for many programs (often written in Perl) that complain when no locale is set.

After building this container, you can verify that the environment variables are set appropriately at runtime with the following command:

```
$ singularity exec my_container.sif env | grep -E 'LISTEN_PORT|LC_ALL'
LISTEN_PORT=12345
LC_ALL=C
```

In the special case of variables generated at build time, you can also add environment variables to your container in the `%post` section (see below).

At build time, the content of the `%environment` section is written to a file called `/.singularity.d/env/90-environment.sh` inside of the container. Text redirected to the `$SINGULARITY_ENVIRONMENT` variable during `%post` (see below) is added to a file called `/.singularity.d/env/91-environment.sh`.

At runtime, scripts in `/.singularity/env` are sourced in order. This means that variables in the `%post` section take precedence over those added via `%environment`.

See [Environment and Metadata](#) for more information about the Singularity container environment.

### 6.3.4 %post

Commands in the %post section are executed within the container after the base OS has been installed at build time. This is where you will download files from the internet with tools like `git` and `wget`, install new software and libraries, write configuration files, create new directories, etc.

Consider the example from the definition file above:

```
%post
  apt-get update && apt-get install -y netcat
  NOW=`date`
  echo "export NOW=\"${NOW}\"" >> $SINGULARITY_ENVIRONMENT
```

This %post scriptlet uses the Ubuntu package manager `apt` to update the container and install the program `netcat` (that will be used in the %startscript section below).

The script is also setting an environment variable at build time. Note that the value of this variable cannot be anticipated, and therefore cannot be set during the %environment section. For situations like this, the `$SINGULARITY_ENVIRONMENT` variable is provided. Redirecting text to this variable will cause it to be written to a file called `./singularity.d/env/91-environment.sh` that will be sourced at runtime. Note that variables set in %post take precedence over those set in the %environment section as explained above.

### 6.3.5 %runscript

The contents of the %runscript section are written to a file within the container that is executed when the container image is run (either via the `singularity run` command or by executing the container directly as a command). When the container is invoked, arguments following the container name are passed to the runscript. This means that you can (and should) process arguments within your runscript.

Consider the example from the def file above:

```
%runscript
  echo "Container was created $NOW"
  echo "Arguments received: $*"
  exec echo "$@"
```

In this runscript, the time that the container was created is echoed via the `$NOW` variable (set in the %post section above). The options passed to the container at runtime are printed as a single string (`$*`) and then they are passed to `echo` via a quoted array (`$@`) which ensures that all of the arguments are properly parsed by the executed command. The `exec` preceding the final `echo` command replaces the current entry in the process table (which originally was the call to Singularity). Thus the runscript shell process ceases to exist, and only the process running within the container remains.

Running the container built using this def file will yield the following:

```
$ ./my_container.sif
Container was created Thu Dec  6 20:01:56 UTC 2018
Arguments received:

$ ./my_container.sif this that and the other
Container was created Thu Dec  6 20:01:56 UTC 2018
Arguments received: this that and the other
this that and the other
```

### 6.3.6 %startscript

Similar to the %runscript section, the contents of the %startscript section are written to a file within the container at build time. This file is executed when the instance start command is issued.

Consider the example from the def file above.

```
%startscript
nc -lp $LISTEN_PORT
```

Here the netcat program is used to listen for TCP traffic on the port indicated by the \$LISTEN\_PORT variable (set in the %environment section above). The script can be invoked like so:

```
$ singularity instance start my_container.sif instance1
INFO:      instance started successfully

$ lsof | grep LISTEN
nc        19061          vagrant    3u      IPv4        107409      0t0      
↪ TCP *:12345 (LISTEN)

$ singularity instance stop instance1
Stopping instance1 instance of /home/vagrant/my_container.sif (PID=19035)
```

### 6.3.7 %test

The %test section runs at the very end of the build process to validate the container using a method of your choice. You can also execute this scriptlet through the container itself, using the test command.

Consider the example from the def file above:

```
%test
grep -q NAME=\"Ubuntu\" /etc/os-release
if [ $? -eq 0 ]; then
    echo "Container base is Ubuntu as expected."
else
    echo "Container base is not Ubuntu."
fi
```

This (somewhat silly) script tests if the base OS is Ubuntu. You could also write a script to test that binaries were appropriately downloaded and built, or that software works as expected on custom hardware. If you want to build a container without running the %test section (for example, if the build system does not have the same hardware that will be used on the production system), you can do so with the --notest build option:

```
$ sudo singularity build --notest my_container.sif my_container.def
```

Running the test command on a container built with this def file yields the following:

```
$ singularity test my_container.sif
Container base is Ubuntu as expected.
```

### 6.3.8 %labels

The %labels section is used to add metadata to the file /.singularity.d/labels.json within your container. The general format is a name-value pair.

Consider the example from the def file above:

**%labels**

```
Author d@sylabs.io
Version v0.0.1
```

The easiest way to see labels is to inspect the image:

```
$ singularity inspect my_container.sif
{
  "Author": "d@sylabs.io",
  "Version": "v0.0.1",
  "org.label-schema.build-date": "Thursday_6_December_2018_20:1:56_UTC",
  "org.label-schema.schema-version": "1.0",
  "org.label-schema.usage": "/.singularity.d/runscript.help",
  "org.label-schema.usage.singularity.deffile.bootstrap": "library",
  "org.label-schema.usage.singularity.deffile.from": "ubuntu:18.04",
  "org.label-schema.usage.singularity.runscript.help": "/.singularity.d/runscript.
↪help",
  "org.label-schema.usage.singularity.version": "3.0.1"
}
```

Some labels that are captured automatically from the build process. You can read more about labels and metadata [here](#).

### 6.3.9 %help

Any text in the %help section is transcribed into a metadata file in the container during the build. This text can then be displayed using the run-help command.

Consider the example from the def file above:

**%help**

```
This is a demo container used to illustrate a def file that uses all
supported sections.
```

After building the help can be displayed like so:

```
$ singularity run-help my_container.sif
This is a demo container used to illustrate a def file that uses all
supported sections.
```

## 6.4 Multi-Stage Builds

Singularity 3.2 introduces multi-stage builds where one environment can be used for compilation, then the resulting binary can be copied into a final environment. This allows a slimmer final image that does not require the entire development stack.

```
Bootstrap: docker
From: golang:1.12.3-alpine3.9
Stage: build

%post
  # prep environment
```

(continues on next page)

(continued from previous page)

```

export PATH="/go/bin:/usr/local/go/bin:$PATH"
export HOME="/root"
cd /root

# insert source code, could also be copied from host with %files
cat << EOF > hello.go
package main
import "fmt"

func main() {
    fmt.Printf("Hello World!\n")
}
EOF

go build -o hello hello.go

# Install binary into final image
Bootstrap: library
From: alpine:3.9
Stage: final

# install binary from stage one
%files from build
    /root/hello /bin/hello

```

The names of stages are arbitrary. Files can only be copied from stages declared before the current stage in the definition. E.g., the `devel` stage in the above definition cannot copy files from the `final` stage, but the `final` stage can copy files from the `devel` stage.

## 6.5 Apps

In some circumstances, it may be redundant to build different containers for each app with nearly equivalent dependencies. Singularity supports installing apps within internal modules based on the concept of [Standard Container Integration Format \(SCI-F\)](#)

The following runsript demonstrates how to build 2 different apps into the same container using SCI-F modules:

```

Bootstrap: docker
From: ubuntu

%environment
    GLOBAL=variables
    AVAILABLE="to all apps"

#####
# foo
#####

%apprun foo
    exec echo "RUNNING FOO"

%applabels foo

```

(continues on next page)



(continued from previous page)

```

BESTAPP FOO

%appinstall foo
  touch foo.exec

%appenv foo
  SOFTWARE=foo
  export SOFTWARE

%apphelp foo
  This is the help for foo.

%appfiles foo
  foo.txt

#####
# bar
#####

%apphelp bar
  This is the help for bar.

%applabels bar
  BESTAPP BAR

%appinstall bar
  touch bar.exec

%appenv bar
  SOFTWARE=bar
  export SOFTWARE

```

An `%appinstall` section is the equivalent of `%post` but for a particular app. Similarly, `%appenv` equates to the app version of `%environment` and so on.

The `%app*` sections can exist alongside any of the primary sections (i.e. `%post`, `%runscript`, `%environment`, etc.). As with the other sections, the ordering of the `%app*` sections isn't important.

After installing apps into modules using the `%app*` sections, the `--app` option becomes available allowing the following functions:

To run a specific app within the container:

```
% singularity run --app foo my_container.sif
RUNNING FOO
```

The same environment variable, `$SOFTWARE` is defined for both apps in the def file above. You can execute the following command to search the list of active environment variables and `grep` to determine if the variable changes depending on the app we specify:

```
$ singularity exec --app foo my_container.sif env | grep SOFTWARE
SOFTWARE=foo

$ singularity exec --app bar my_container.sif env | grep SOFTWARE
SOFTWARE=bar
```

## 6.6 Best Practices for Build Recipes

When crafting your recipe, it is best to consider the following:

1. Always install packages, programs, data, and files into operating system locations (e.g. not `/home`, `/tmp`, or any other directories that might get commonly binded on).
2. Document your container. If your runsript doesn't supply help, write a `%help` or `%apphelp` section. A good container tells the user how to interact with it.
3. If you require any special environment variables to be defined, add them to the `%environment` and `%appenv` sections of the build recipe.
4. Files should always be owned by a system account (UID less than 500).
5. Ensure that sensitive files like `/etc/passwd`, `/etc/group`, and `/etc/shadow` do not contain secrets.
6. Build production containers from a definition file instead of a sandbox that has been manually changed. This ensures greatest possibility of reproducibility and mitigates the “black box” effect.

## BUILD ENVIRONMENT

### 7.1 Overview

You may wish to customize your build environment by doing things such as specifying a custom cache directory for images or sending your Docker Credentials to the registry endpoint. Here we will discuss these and other topics related to the build environment.

### 7.2 Cache Folders

To make downloading images for build and pull faster and less redundant, Singularity uses a caching strategy. By default, Singularity will create a set of folders in your `$HOME` directory for docker layers, Cloud library images, and metadata, respectively:

```
$HOME/.singularity/cache/library
$HOME/.singularity/cache/oci
$HOME/.singularity/cache/oci-tmp
```

If you want to cache in a different directory, set `SINGULARITY_CACHEDIR` to the desired path. By using the `-E` option with the `sudo` command, `SINGULARITY_CACHEDIR` will be passed along to root's environment and respected during the build. Remember that when you run commands as root images will be cached in root's home at `/root` and not your user's home.

### 7.3 Cache commands

Singularity 3.1 comes with new commands for cleaning and listing the cache image files generated.

---

**Note:** Running the cache commands with `sudo` privilege will consider cache location as `/root/.singularity/cache`. The default location for cache without `sudo` privilege is `~/.singularity/cache`. Make sure that if you build a container with `sudo` privilege, you will need to consider the `sudo` location from the cache, and not the default.

For example, running the following command with `sudo` privilege (considering the `sudo` privilege location for cache `/root/.singularity/cache`):

```
$ sudo singularity cache list
NAME                DATE CREATED      SIZE      TYPE
ubuntu_latest.sif  2019-01-31 14:59:32  28.11 Mb  library
ubuntu_18.04.sif   2019-01-31 14:58:44  27.98 Mb  library
```

and then cleaning the cache without sudo privilege (`singularity cache clean -a`) will not work, since the default cache location is `~/.singularity/cache`. In this case you would need to run the clean command with sudo privilege:

```
$ sudo singularity cache clean -a
NAME                                DATE CREATED          SIZE                   TYPE
There 0 containers using: 0.00 kB, 0 oci blob file(s) using 0.00 kB of space.
Total space used: 0.00 kB
```

### 7.3.1 Listing Cache

For example, you can list cache image files and check which type they belong to: Library or oci.

```
$ singularity cache list
NAME                                DATE CREATED          SIZE                   TYPE
ubuntu_latest.sif                 2019-01-31 14:59:32  28.11 Mb              library
ubuntu_18.04.sif                  2019-01-31 14:58:44  27.98 Mb              library
alpine_latest.sif                 2019-01-31 14:58:24  2.18 Mb               library
centos_latest.sif                 2019-01-31 14:59:07  72.96 Mb              library
centos_latest.sif                 2019-01-31 14:59:26  73.45 Mb              oci
ubuntu_18.04.sif                  2019-01-31 14:58:58  27.99 Mb              oci
ubuntu_latest.sif                 2019-01-31 14:59:41  27.99 Mb              oci
alpine_latest.sif                 2019-01-31 14:58:30  2.72 Mb               oci

There are 15 oci blob file(s) using 112.51 Mb of space. Use: '-T=blob' to list
```

You can also clean a specific cache type, choosing between: `library`, `oci`, `blob` (separated by commas)

```
# clean only library cache
$ singularity cache clean --type=library

# clean only oci cache
$ singularity cache clean --type=oci

# clean only blob cache
$ singularity cache clean --type=blob

# clean only library, and oci cache
$ singularity cache clean --type=library,oci
```

---

**Note:** This feature of passing additional flags with comma-separated arguments can also be used with the `singularity cache clean` command we will see below.

---

### 7.3.2 Cleaning the Cache

Most of the `cache clean` and `cache list` flags can be interchanged, (`--name` is only reserved for `cache clean`).

It's worth noting that by running the following command: (with no flags)

```
$ singularity cache clean
```

By default will just clean the blob cache, but if you do:

```
$ singularity cache clean --all
```

It will clean all the cache.

## 7.4 Temporary Folders

Singularity uses a temporary directory to build the squashfs filesystem, and this temp space needs to be large enough to hold the entire resulting Singularity image. By default this happens in `/tmp` but the location can be configured by setting `SINGULARITY_TMPDIR` to the full path where you want the sandbox and squashfs temp files to be stored. Remember to use `-E` option to pass the value of `SINGULARITY_TMPDIR` to root's environment when executing the `build` command with `sudo`.

When you run one of the action commands (i.e. `run`, `exec`, or `shell`) with a container from the container library or an OCI registry, Singularity builds the container in the temporary directory caches it and runs it from the cached location.

Consider the following command:

```
$ singularity exec docker://busybox /bin/sh
```

This container is first built in `/tmp`. Since all the oci blobs are converted into SIF format, by default a temporary runtime directory is created in:

```
$HOME/.singularity/cache/oci-tmp/<sha256-code>/busybox_latest.sif
```

In this case, the `SINGULARITY_TMPDIR` and `SINGULARITY_CACHEDIR` variables will also be respected.

## 7.5 Pull Folder

To customize your pull default location you can do so by specifying Singularity in which folder to pull the image, assuming you own a folder called `mycontainers` inside your `$HOME` folder, you would need to do something like the following:

```
$ singularity pull $HOME/mycontainers library://library/default/alpine
```

Singularity also allows you to modify the default cache location for pulling images. By default, the location of the pull folder is given by the environment variable `SINGULARITY_CACHEDIR`. `SINGULARITY_CACHEDIR` by default points to `$HOME/.singularity/cache` but this path can be modified. You would need to set and export the `SINGULARITY_CACHEDIR` environment variable before pulling the image, like so:

And that will successfully pull that container image inside your new `SINGULARITY_CACHEDIR` location.

## 7.6 Environment Variables

1. If a flag is represented by both a CLI option and an environment variable, and both are set, the CLI option will always take precedence. This is true for all environment variables except for `SINGULARITY_BIND` and `SINGULARITY_BINDPATH` which is combined with the `--bind` option, argument pair if both are present.

2. Environment variables overwrite default values in the CLI code
3. Any defaults in the CLI code are applied.

### 7.6.1 Defaults

The following variables have defaults that can be customized by you via environment variables at runtime.

#### 7.6.1.1 Docker

**SINGULARITY\_DOCKER\_LOGIN** Used for the interactive login for Docker Hub.

**SINGULARITY\_DOCKER\_USERNAME** Your Docker username.

**SINGULARITY\_DOCKER\_PASSWORD** Your Docker password.

**RUNSCRIPT\_COMMAND** Is not obtained from the environment, but is a hard coded default (“/bin/bash”). This is the fallback command used in the case that the docker image does not have a **CMD** or **ENTRYPOINT**. **TAG** Is the default tag, `latest`.

**SINGULARITY\_NOHTTPS** This is relevant if you want to use a registry that doesn’t have https, and it speaks for itself. If you export the variable `SINGULARITY_NOHTTPS` you can force the software to not use https when interacting with a Docker registry. This use case is typically for use of a local registry.

#### 7.6.1.2 Library

**SINGULARITY\_BUILDER** Used to specify the remote builder service URL. The default value is our remote builder.

**SINGULARITY\_LIBRARY** Used to specify the library to pull from. Default is set to our Cloud Library.

**SINGULARITY\_REMOTE** Used to build an image remotely (This does not require root). The default is set to false.

## CLOUD LIBRARY

This page will cover how to use the Singularity Cloud Services (SCS) with Singularity.

### 8.1 Overview

The Cloud Library is the place to *push* your containers to the cloud so other users can download, *verify*, and use the containers.

Sylabs also provides a *Remote Builder*, used to build your containers containers without root access within the cloud.

### 8.2 Make an Account

Making an account is easy, and straightforward:

1. Go to: <https://cloud.sylabs.io/library>.
2. Click “Sign in to Sylabs” (top right corner).
3. Select your method to sign in, with Google, GitHub, GitLab, or Microsoft.
4. Type your passwords, and that’s it!

### 8.3 Creating a Access token

Access tokens for pushing a container, and remote builder.

To generate a access token, do the following steps:

1. Go to: <https://cloud.sylabs.io/library>.
2. Click the top right button with your username.
3. From the dropdown menu, click “Access Tokens”.
4. Enter a token name, like `my token for linux ws`, then click “Create a new token”.
5. Copy the token and paste it to `~/.singularity/sylabs-token`.

Now that you have your token, you are ready to push your container!

## 8.4 Pushing a Container

The `singularity push` command will push a container to the container library with the given URL. Here's an example of a typical push command:

```
$ singularity push my-container.sif library://your-name/project-dir/my-  
→container:latest
```

The `:latest` is the container tag. Tags are used to have different version of the same container.

---

**Note:** When pushing your container, there's no need to add a `.sif` (Singularity Image Format) to the end of the container name, (like on your local machine), because all containers on the library are SIF containers.

---

Let's assume you have your container (v1.0.1), and you want to push that container without deleting your `:latest` container, then you can add a version tag to that container, like so:

```
$ singularity push my-container.sif library://your-name/project-dir/my-container:1.0.1
```

You can download the container with that tag by replacing the `:latest`, with the tagged container you want to download.

## 8.5 Pulling a container

The `singularity pull` command will download a container from the [Library](#) (`library://`), [Docker Hub](#) (`docker://`), and also [Shub](#) (`shub://`).

---

**Note:** When pulling from Docker, the container will automatically be converted to a SIF (Singularity Image Format) container.

---

Here's a typical pull command:

```
$ singularity pull file-out.sif library://alpine:latest  
  
# or pull from docker:  
  
$ singularity pull file-out.sif docker://alpine:latest
```

---

**Note:** If there's no tag after the container name, Singularity automatically will pull the container with the `:latest` tag.

---

To pull a container with a specific tag, just add the tag to the library URL:

```
$ singularity pull file-out.sif library://alpine:3.8
```

Of course, you can pull your own containers. Here's what that will look like:

### 8.5.1 Pulling your own container

Pulling your own container is just like pulling from Github, Docker, etc...



```
$ singularity pull out-file.sif library://your-name/project-dir/my-container:latest
# or use a different tag:
$ singularity pull out-file.sif library://your-name/project-dir/my-container:1.0.1
```

---

**Note:** You *don't* have to specify a output file, one will be created automatically, but it's good practice to always specify your output file.

---

## 8.6 Verify/Sign your Container

Verify containers that you pull from the library, ensuring they are bit-for-bit reproductions of the original image. Check out [this page](#) on how to: *verify a container, making PGP key, and sign your own containers*.

## 8.7 Searching the Library for Containers

When it comes to searching the library, you could always go to: <https://cloud.sylabs.io/library> and search from there through the web GUI. Or you can use `singularity search <container/user>`, this will search the library for the `<container/user>`.

### 8.7.1 Using the CLI Search

Here is an example for searching the library for centos:

```
$ singularity search centos
No users found for 'centos'

No collections found for 'centos'

Found 6 containers for 'centos'
  library://dtrudg/linux/centos
    Tags: 6 7 centos6 centos7 latest
  library://library/default/centos
    Tags: 6 7 latest
  library://gmk/demo/centos-vim
    Tags: latest
  library://mroche/baseline/centos
    Tags: 7 7.5 7.5.1804 7.6 7.6.1810 latest
  library://gmk/default/centos7-devel
    Tags: latest
  library://emmeff/default/centos7-python36
    Tags: 1.0
```

Notice there are different tags for the same container.

## 8.8 Remote Builder

The remote builder service can build your container in the cloud removing the requirement for root access.

Here's a typical remote build command:

```
$ singularity build --remote file-out.sif docker://ubuntu:18.04
```

### 8.8.1 Building from a definition file:

This is our definition file. Let's call it `ubuntu.def`:

```
bootstrap: library
from: ubuntu:18.04

%runscript
    echo "hello world from ubuntu container!"
```

Now, to build the container, use the `--remote` flag, and without `sudo`:

```
$ singularity build --remote ubuntu.sif ubuntu.def
```

---

**Note:** Make sure you have a *access token*, otherwise the build will fail.

---

After building, you can test your container like so:

```
$ ./ubuntu.sif
hello world from ubuntu container!
```

You can also use the web GUI to build containers remotely. First, go to <https://cloud.sylabs.io/builder> (make sure you are signed in). Then you can copy and paste, upload, or type your definition file. When you are finished, click build. Then you can download the container with the URL.

## SUPPORT FOR DOCKER AND OCI

### 9.1 Overview

Effort has been expended in developing [Docker](#) containers. Deconstructed into one or more compressed archives (typically split across multiple segments, or **layers** as they are known in Docker parlance) plus some metadata, images for these containers are built from specifications known as `Dockerfiles`. The public [Docker Hub](#), as well as various private registries, host images for use as Docker containers. Singularity has from the outset emphasized the importance of interoperability with Docker. As a consequence, this section of the Singularity User Docs first makes its sole focus interoperability with Docker. In so doing, the following topics receive attention here:

- Application of Singularity action commands on ephemeral containers derived from public Docker images
- Converting public Docker images into Singularity's native format for containerization, namely the Singularity Image Format (SIF)
- Authenticated application of Singularity commands to containers derived from private Docker images
- Authenticated application of Singularity commands to containers derived from private Docker images originating from private registries
- Building SIF containers for Singularity via the command line or definition files from a variety of sources for Docker images and image archives

The second part of this section places emphasis upon Singularity's interoperability with open standards emerging from the [Open Containers Initiative](#) (OCI). Specifically, in documenting Singularity interoperability as it relates to the OCI Image Specification, the following topics are covered:

- Compliance with the OCI Image Layout Specification
- OCI-compliant caching in Singularity
- Acquiring OCI images and image archives via Singularity
- Building SIF containers for Singularity via the command line or definition files from a variety of sources for OCI images and image archives

The section closes with a brief enumeration of emerging best practices plus consideration of troubleshooting common issues.

### 9.2 Running action commands on public images from Docker Hub

`godlovedc/lolcow` is a whimsical example of a publicly accessible image hosted via [Docker Hub](#). Singularity can execute this image as follows:

```

$ singularity run docker://godlovedc/lolcow
INFO:   Converting OCI blobs to SIF format
INFO:   Starting build...
Getting image source signatures
Copying blob sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
 45.33 MiB / 45.33 MiB [=====] 1s
Copying blob sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
 848 B / 848 B [=====] 0s
Copying blob sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
 621 B / 621 B [=====] 0s
Copying blob sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
 853 B / 853 B [=====] 0s
Copying blob sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
 169 B / 169 B [=====] 0s
Copying blob sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
 53.75 MiB / 53.75 MiB [=====] 2s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: /home/vagrant/.singularity/cache/oci-tmp/
↪a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb/lolcow_latest.sif
INFO:   Image cached as SIF at /home/vagrant/.singularity/cache/oci-tmp/
↪a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb/lolcow_latest.sif

-----
/ Repartee is something we think of \
| twenty-four hours too late.      |
|                                  |
\ -- Mark Twain                    /
-----
      \   ^__^
       \  (oo)\_______
          (__)\       )\/\
             ||----w |
              ||     ||

```

Here `docker` is prepended to ensure that the `run` command of Singularity is instructed to bootstrap container creation based upon this Docker image, thus creating a complete URI for Singularity. Singularity subsequently downloads *all the OCI blobs that comprise this image*, and converts them into a *single* SIF file - the native format for Singularity containers. Because this image from Docker Hub is cached locally in the `$HOME/.singularity/cache/oci-tmp/<org.opencontainers.image.ref.name>/lolcow_latest.sif` directory, where `<org.opencontainers.image.ref.name>` will be replaced by the appropriate hash for the container, the image does not need to be downloaded again (from Docker Hub) the next time a Singularity run is executed. In other words, the cached copy is sensibly reused:

```

$ singularity run docker://godlovedc/lolcow

-----
/ Soap and education are not as sudden as \
| a massacre, but they are more deadly in |
| the long run.                          |
|                                  |
\ -- Mark Twain                          /
-----
      \   ^__^
       \  (oo)\_______
          (__)\       )\/\

```

(continues on next page)

(continued from previous page)

```

  ||----w |
  ||     ||

```

**Note:** Image caching is *documented in detail below*.

**Note:** Use is made of the `$HOME/.singularity` directory by default to cache images. To cache images elsewhere, use of the environment variable `SINGULARITY_CACHEDIR` can be made.

As the runtime of this container is encapsulated as a single SIF file, it is possible to

```

cd /home/vagrant/.singularity/cache/oci-tmp/
↪ a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb/

```

and then execute the SIF file directly:

```

./lolcow_latest.sif
-----
/ The secret source of humor is not joy \
| but sorrow; there is no humor in    |
| Heaven.                             |
|                                     |
\ -- Mark Twain                       /
-----
      \      ^__^
       \      (oo)\_______
            (__)\       )\/\
                ||----w |
                ||     ||

```

**Note:** SIF files abstract Singularity containers as a single file. As with any executable, a SIF file can be executed directly.

`fortune | cowsay | lolcat` is executed by *default* when this container is run by Singularity. Singularity's `exec` command allows a different command to be executed; for example:

```

$ singularity exec docker://godlovedc/lolcow fortune
Don't go around saying the world owes you a living. The world owes you
nothing. It was here first.
-- Mark Twain

```

**Note:** The *same* cached copy of the `lolcow` container is reused here by Singularity `exec`, and immediately below here by shell.

**Note:** Execution defaults are documented below - see *Directing Execution* and *Container Metadata*.

In addition to non-interactive execution of an image from Docker Hub, Singularity provides support for an *interactive* shell session:

```
$ singularity shell docker://godlovedc/lolcow
Singularity lolcow_latest.sif:~> cat /etc/os-release
NAME="Ubuntu"
VERSION="16.04.3 LTS (Xenial Xerus)"
ID=ubuntu
ID_LIKE=debian
PRETTY_NAME="Ubuntu 16.04.3 LTS"
VERSION_ID="16.04"
HOME_URL="http://www.ubuntu.com/"
SUPPORT_URL="http://help.ubuntu.com/"
BUG_REPORT_URL="http://bugs.launchpad.net/ubuntu/"
VERSION_CODENAME=xenial
UBUNTU_CODENAME=xenial
Singularity lolcow_latest.sif:~>
```

From this it is evident that use is being made of Ubuntu 16.04 *within* this container, whereas the shell *external* to the container is running a more recent release of Ubuntu (not illustrated here).

`inspect` reveals the metadata for a Singularity container encapsulated via SIF; *Container Metadata* is documented below.

---

**Note:** `singularity search [search options...] <search query>` does *not* support Docker registries like [Docker Hub](#). Use the search box at Docker Hub to locate Docker images. Docker pull commands, e.g., `docker pull godlovedc/lolcow`, can be easily translated into the corresponding command for Singularity. The Docker pull command is available under “DETAILS” for a given image on Docker Hub.

---

## 9.3 Making use of public images from Docker Hub

Singularity can make use of public images available from the [Docker Hub](#). By specifying the `docker://` URI for an image that has already been located, Singularity can pull it - e.g.:

```
$ singularity pull docker://godlovedc/lolcow
WARNING: Authentication token file not found : Only pulls of public images will
↳succeed
INFO: Starting build...
Getting image source signatures
Copying blob sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
 45.33 MiB / 45.33 MiB [=====] 2s
Copying blob sha256:3b61febd4ae982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
 848 B / 848 B [=====] 0s
Copying blob sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
 621 B / 621 B [=====] 0s
Copying blob sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
 853 B / 853 B [=====] 0s
Copying blob sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
 169 B / 169 B [=====] 0s
Copying blob sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
 53.75 MiB / 53.75 MiB [=====] 3s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: lolcow_latest.sif
```

This pull results in a *local* copy of the Docker image in SIF, the Singularity Image Format:

```
$ file lolcow_latest.sif
lolcow_latest.sif: a /usr/bin/env run-singularity script executable (binary data)
```

In converting to SIF, individual layers of the Docker image have been *combined* into a single, native file for use by Singularity; there is no need to subsequently *build* the image for Singularity. For example, you can now `exec`, `run` or `shell` into the SIF version via Singularity, *as described above*.

---

**Note:** The above authentication warning originates from a check for the existence of `/${HOME}/.singularity/sylabs-token`. It can be ignored when making use of Docker Hub, or silenced by issuing `touch ${HOME}/.singularity/sylabs-token` once.

---

`inspect` reveals metadata for the container encapsulated via SIF:

```
$ singularity inspect lolcow_latest.sif

{
  "org.label-schema.build-date": "Thursday_6_December_2018_17:29:48_UTC",
  "org.label-schema.schema-version": "1.0",
  "org.label-schema.usage.singularity.deffile.bootstrap": "docker",
  "org.label-schema.usage.singularity.deffile.from": "godlovedc/lolcow",
  "org.label-schema.usage.singularity.version": "3.0.1-40.g84083b4f"
}
```

---

**Note:** *Container Metadata* is documented below.

---

SIF files built from Docker images are *not* cryptographically signed:

```
$ singularity verify lolcow_latest.sif
Verifying image: lolcow_latest.sif
ERROR: verification failed: error while searching for signature blocks: no_
↳signatures found for system partition
```

The `sign` command allows a cryptographic signature to be added. Refer to *Signing and Verifying Containers* for details. But caution should be exercised in signing images from Docker Hub because, unless you build an image from scratch (OS mirrors) you are probably not really sure about the complete contents of that image.

---

**Note:** `pull` is a one-time-only operation that builds a SIF file corresponding to the image retrieved from Docker Hub. Updates to the image on Docker Hub will *not* be reflected in the *local* copy.

---

In our example `docker://godlovedc/lolcow`, `godlovedc` specifies a Docker Hub user, whereas `lolcow` is the name of the repository. Adding the option to specify an image tag, the generic version of the URI is `docker://<user>/<repo-name>[:<tag>]`. [Repositories on Docker Hub](#) provides additional details.

## 9.4 Making use of private images from Docker Hub

After successful authentication, Singularity can also make use of *private* images available from the [Docker Hub](#). The two means available for authentication follow here. Before describing these means, it is instructive to illustrate the error generated when attempting access a private image *without* credentials:

```
$ singularity pull docker://ilumb/mylolcow
INFO: Starting build...
FATAL: Unable to pull docker://ilumb/mylolcow: conveyor failed to get: Error_
↳reading manifest latest in docker.io/ilumb/mylolcow: errors:
denied: requested access to the resource is denied
unauthorized: authentication required
```

In this case, the `mylolcow` repository of user `ilumb` **requires** authentication through specification of a valid user-name and password.

### 9.4.1 Authentication via Interactive Login

Interactive login is the first of two means provided for authentication with Docker Hub. It is enabled through use of the `--docker-login` option of Singularity's `pull` command; for example:

```
$ singularity pull --docker-login docker://ilumb/mylolcow
Enter Docker Username: ilumb
Enter Docker Password:
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↳sha256:7b8b6451c85f072fd0d7961c97be3fe6e2f772657d471254f6d52ad9f158a580
Skipping fetch of repeat blob_
↳sha256:ab4d1096d9ba178819a3f71f17add95285b393e96d08c8a6bfc3446355bcdc49
Skipping fetch of repeat blob_
↳sha256:e6797d1788acd741d33f4530106586ffee568be513d47e6e20a4c9bc3858822e
Skipping fetch of repeat blob_
↳sha256:e25c5c290bde5267364aa9f59a18dd22a8b776d7658a41ffabbf691d8104e36
Skipping fetch of repeat blob_
↳sha256:258e068bc5e36969d3ba4b47fd3ca0d392c6de465726994f7432b14b0414d23b
Copying config sha256:8a8f815257182b770d32dffff7f185013b4041d076e065893f9dd1e89ad8a671
 3.12 KiB / 3.12 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: mylolcow_latest.sif
```

After successful authentication, the private Docker image is pulled and converted to SIF as described above.

---

**Note:** For interactive sessions, `--docker-login` is *recommended* as use of plain-text passwords in your environment is *avoided*. Encoded authentication data is communicated with Docker Hub via secure HTTP.

---

### 9.4.2 Authentication via Environment Variables

Environment variables offer an alternative means for authentication with Docker Hub. The **required** exports are as follows:

```
export SINGULARITY_DOCKER_USERNAME=ilumb
export SINGULARITY_DOCKER_PASSWORD=<redacted>
```

Of course, the `<redacted>` plain-text password needs to be replaced by a valid one to be of practical use.

Based upon these exports, `$ singularity pull docker://ilumb/mylolcow` allows for the retrieval of this private image.



---

**Note:** This approach for authentication supports both interactive and non-interactive sessions. However, the requirement for a plain-text password assigned to an environment variable, is the security compromise for this flexibility.

---



---

**Note:** When specifying passwords, ‘special characters’ (e.g., \$, #, .) need to be ‘escaped’ to avoid interpretation by the shell.

---

## 9.5 Making use of private images from Private Registries

Authentication is required to access *private* images that reside in Docker Hub. Of course, private images can also reside in **private registries**. Accounting for locations *other* than Docker Hub is easily achieved.

In the complete command line specification

```
docker://<registry>/<user>/<repo-name>[:<tag>]
```

registry defaults to `index.docker.io`. In other words,

```
$ singularity pull docker://godlovedc/lolcow
```

is functionally equivalent to

```
$ singularity pull docker://index.docker.io/godlovedc/lolcow
```

From the above example, it is evident that

```
$ singularity pull docker://nvcr.io/nvidia/pytorch:18.11-py3
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↪sha256:18d680d616571900d78ee1c8fff0310f2a2afe39c6ed0ba2651ff667af406c3e
<blob fetching details deleted>
Skipping fetch of repeat blob_
↪sha256:c71aeebc266c779eb4e769c98c935356a930b16d881d7dde4db510a09cfa4222
Copying config sha256:b77551af8073c85588088ab2a39007d04bc830831ba1eef4127b2d39aaf3a6b1
21.28 KiB / 21.28 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: pytorch_18.11-py3.sif
```

will retrieve a specific version of the **PyTorch platform** for Deep Learning from the NVIDIA GPU Cloud (NGC). Because NGC is a private registry, the above `pull` assumes *authentication via environment variables* when the blobs that collectively comprise the Docker image have not already been cached locally. In the NGC case, the required environment variable are set as follows:

```
export SINGULARITY_DOCKER_USERNAME='$oauthtoken'
export SINGULARITY_DOCKER_PASSWORD=<redacted>
```

Upon use, these environment-variable settings allow for authentication with NGC.

---

**Note:** `$oauthtoken` is to be taken literally - it is not, for example, an environment variable.

---

The password provided via these means is actually an API token. This token is generated via your NGC account, and is **required** for use of the service.

For additional details regarding authentication with NGC, and much more, please consult the NGC [Getting Started](#) documentation.

---

Alternatively, for purely interactive use, `--docker-login` is recommended:

```
$ singularity pull --docker-login docker://nvcr.io/nvidia/pytorch:18.11-py3
Enter Docker Username: $oauthtoken
Enter Docker Password:
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↪sha256:18d680d616571900d78eelc8fff0310f2a2afe39c6ed0ba2651ff667af406c3e
<blob fetching details deleted>
Skipping fetch of repeat blob_
↪sha256:c71aeebc266c779eb4e769c98c935356a930b16d881d7dde4db510a09cfa4222
Copying config sha256:b77551af8073c85588088ab2a39007d04bc830831ba1eef4127b2d39aaf3a6b1
21.28 KiB / 21.28 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: pytorch_18.11-py3.sif
```

Authentication aside, the outcome of the `pull` command is the Singularity container `pytorch_18.11-py3.sif` - i.e., a locally stored copy, that has been converted to SIF.

## 9.6 Building images for Singularity from Docker Registries

The `build` command is used to **create** Singularity containers. Because it is documented extensively *elsewhere in this manual*, only specifics relevant to Docker are provided here - namely, working with Docker Hub via *the Singularity command line* and through *Singularity definition files*.

### 9.6.1 Working from the Singularity Command Line

#### 9.6.1.1 Remotely Hosted Images

In the simplest case, `build` is functionally equivalent to `pull`:

```
$ singularity build mylolcow_latest.sif docker://godlovedc/lolcow
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↪sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
Skipping fetch of repeat blob_
↪sha256:3b61febd4ae982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
Skipping fetch of repeat blob_
↪sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
Skipping fetch of repeat blob_
↪sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
Skipping fetch of repeat blob_
↪sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
```

(continues on next page)

(continued from previous page)

```

Skipping fetch of repeat blob_
↳sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: mylolcow_latest.sif

```

This build results in a *local* copy of the Docker image in SIF, as did pull *above*. Here, build has named the Singularity container `mylolcow_latest.sif`.

**Note:** `docker://godlovedc/lolcow` is the **target** provided as input for `build`. Armed with this target, `build` applies the appropriate bootstrap agent to create the container - in this case, one appropriate for Docker Hub.

In addition to a read-only container image in SIF (**default**), `build` allows for the creation of a writable (ch)root *directory* called a **sandbox** for interactive development via the `--sandbox` option:

```

$ singularity build --sandbox mylolcow_latest_sandbox docker://godlovedc/lolcow
INFO:   Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↳sha256:9fb6c798fa41e509b58bccc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
Skipping fetch of repeat blob_
↳sha256:3b61febd4ae9e982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
Skipping fetch of repeat blob_
↳sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
Skipping fetch of repeat blob_
↳sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
Skipping fetch of repeat blob_
↳sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
Skipping fetch of repeat blob_
↳sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating sandbox directory...
INFO:   Build complete: mylolcow_latest_sandbox

```

After successful execution, the above command results in creation of the `mylolcow_latest_sandbox` directory with contents:

```

bin boot core dev environment etc home lib lib64 media mnt opt proc root_
↳ run sbin singularity srv sys tmp usr var

```

The `build` command of Singularity allows (e.g., development) sandbox containers to be converted into (e.g., production) read-only SIF containers, and vice-versa. Consult the *Build a container* documentation for the details.

Implicit in the above command-line interactions is use of public images from Docker Hub. To make use of **private** images from Docker Hub, authentication is required. Available means for authentication were described above. Use of environment variables is functionally equivalent for Singularity `build` as it is for `pull`; see *Authentication via Environment Variables* above. For purely interactive use, authentication can be added to the `build` command as follows:

```
singularity build --docker-login mylolcow_latest_il.sif docker://ilumb/mylolcow
```

(Recall that `docker://ilumb/mylolcow` is a private image available via Docker Hub.) See [Authentication via Interactive Login](#) above regarding use of `--docker-login`.

### 9.6.1.2 Building Containers Remotely

By making use of the [Sylabs Cloud Remote Builder](#), it is possible to build SIF containers *remotely* from images hosted at Docker Hub. The Sylabs Cloud Remote Builder is a **service** that can be used from the Singularity command line or via its Web interface. Here use of the Singularity CLI is emphasized.

Once you have an account for Sylabs Cloud, and have logged in to the portal, select [Remote Builder](#). The right-hand side of this page is devoted to use of the Singularity CLI. Self-generated API tokens are used to enable authenticated access to the Remote Builder. To create a token, follow the [instructions provided](#). Once the token has been created, store it in the file `$HOME/.singularity/sylabs-token`.

The above token provides *authenticated* use of the Sylabs Cloud Remote Builder when `--remote` is *appended* to the Singularity `build` command. For example, for remotely hosted images:


```
$ singularity build --remote lolcow_rb.sif docker://godlovedc/lolcow
searching for available build agent.....INFO:   Starting build...
Getting image source signatures
Copying blob sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
 45.33 MiB / 45.33 MiB  0s
Copying blob sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
 848 B / 848 B  0s
Copying blob sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
 621 B / 621 B  0s
Copying blob sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
 853 B / 853 B  0s
Copying blob sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
 169 B / 169 B  0s
Copying blob sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
 53.75 MiB / 53.75 MiB  0s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB  0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: /tmp/image-341891107
INFO:   Now uploading /tmp/image-341891107 to the library
 87.94 MiB / 87.94 MiB 100.00% 38.96 MiB/s 2s
INFO:   Setting tag latest
 87.94 MiB / 87.94 MiB
↪ [=====]
↪ 100.00% 17.23 MiB/s 5s
```

---

**Note:** Elevated privileges (e.g., via `sudo`) are *not* required when use is made of the Sylabs Cloud Remote Builder.

---

During the build process, progress can be monitored in the Sylabs Cloud portal on the Remote Builder page - as illustrated upon completion by the screenshot below. Once complete, this results in a *local* copy of the SIF file `lolcow_rb.sif`. From the [Sylabs Cloud Singularity Library](#) it is evident that the ‘original’ SIF file remains available via this portal.

My Builds							
Build	Build Recipe	Submit Time	Started Time	Duration	Library	Status	
docker:godlovedc/lolcow 87.94 MB	Build Recipe	2019-01-23T11:31:21-05:00	2019-01-23T11:33:26-05:00	30s	ilumb/remote-builds/rb-5c4896d9e2126600194b30f	Completed	

### 9.6.1.3 Locally Available Images: Cached by Docker

Singularity containers can be built at the command line from images cached *locally* by Docker. Suppose, for example:

```
$ sudo docker images
```

REPOSITORY	TAG	IMAGE ID	CREATED	SIZE
godlovedc/lolcow	latest	577c1fe8e6d8	16 months ago	241MB

This indicates that `godlovedc/lolcow:latest` has been cached locally by Docker. Then

```
$ sudo singularity build lolcow_from_docker_cache.sif docker-daemon://godlovedc/
↳lolcow:latest
INFO: Starting build...
Getting image source signatures
Copying blob sha256:a2022691bf950a72f9d2d84d557183cb9eee07c065a76485f1695784855c5193
 119.83 MiB / 119.83 MiB [=====] 6s
Copying blob sha256:ae620432889d2553535199dbdd8ba5a264ce85fcdcd5a430974d81fc27c02b45
 15.50 KiB / 15.50 KiB [=====] 0s
Copying blob sha256:c561538251751e3685c7c6e7479d488745455ad7f84e842019dcb452c7b6fecc
 14.50 KiB / 14.50 KiB [=====] 0s
Copying blob sha256:f96e6b25195f1b36ad02598b5d4381e41997c93ce6170cab1b81d9c68c514db0
 5.50 KiB / 5.50 KiB [=====] 0s
Copying blob sha256:7f7a065d245a6501a782bf674f4d7e9d0a62fa6bd212edbf1f17bad0d5cd0bfc
 3.00 KiB / 3.00 KiB [=====] 0s
Copying blob sha256:70ca7d49f8e9c44705431e3dade0636a2156300ae646ff4f09c904c138728839
 116.56 MiB / 116.56 MiB [=====] 6s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: lolcow_from_docker_cache.sif
```

results in `lolcow_from_docker_cache.sif` for native use by Singularity. There are two important differences in syntax evident in the above build command:

1. The `docker` part of the URI has been appended by `daemon`. This ensures Singularity seek an image locally cached by Docker to bootstrap the conversion process to SIF, as opposed to attempting to retrieve an image remotely hosted via Docker Hub.
2. `sudo` is prepended to the `build` command for Singularity; this is required as the Docker daemon executes as `root`. However, if the user issuing the `build` command is a member of the `docker` Linux group, then `sudo` need not be prepended.

---

**Note:** The image tag, in this case `latest`, is **required** when bootstrapping creation of a container for Singularity from an image locally cached by Docker.

---



---

**Note:** The Sylabs Cloud Remote Builder *does not* interoperate with local Docker daemons; therefore, images cached locally by Docker, *cannot* be used to bootstrap creation of SIF files via the Remote Builder service. Of course, a SIF

---

file could be created locally as detailed above. Then, in a separate, manual step, *pushed to the Sylabs Cloud Singularity Library*.

#### 9.6.1.4 Locally Available Images: Stored Archives

Singularity containers can also be built at the command line from Docker images stored locally as tar files.

The `lolcow.tar` file employed below in this example can be produced by making use of an environment in which Docker is available as follows:

1. Obtain a local copy of the image from Docker Hub via `sudo docker pull godlovedc/lolcow`. Issuing the following command confirms that a copy of the desired image is available locally:

```
$ sudo docker images
REPOSITORY          TAG                 IMAGE ID            CREATED             SIZE
godlovedc/lolcow    latest             577c1fe8e6d8       17 months ago      241MB
```

2. Noting that the image identifier above is `577c1fe8e6d8`, the required archive can be created by `sudo docker save 577c1fe8e6d8 -o lolcow.tar`.

Thus `lolcow.tar` is a locally stored archive in the *current* working directory with contents:

```
$ sudo tar tvf lolcow.tar
drwxr-xr-x 0/0          0 2017-09-21 19:37
↳ 02aefa059d08482d344293d0ad27182a0a9d330ebc73abd92a1f9744844f91e9/
-rw-r--r-- 0/0          3 2017-09-21 19:37
↳ 02aefa059d08482d344293d0ad27182a0a9d330ebc73abd92a1f9744844f91e9/VERSION
-rw-r--r-- 0/0          1417 2017-09-21 19:37
↳ 02aefa059d08482d344293d0ad27182a0a9d330ebc73abd92a1f9744844f91e9/json
-rw-r--r-- 0/0          122219008 2017-09-21 19:37
↳ 02aefa059d08482d344293d0ad27182a0a9d330ebc73abd92a1f9744844f91e9/layer.tar
drwxr-xr-x 0/0          0 2017-09-21 19:37
↳ 3762e087ebbb895fd9c38981c1f7bfc76c9879fd3fdadef64df49e92721bb527/
-rw-r--r-- 0/0          3 2017-09-21 19:37
↳ 3762e087ebbb895fd9c38981c1f7bfc76c9879fd3fdadef64df49e92721bb527/VERSION
-rw-r--r-- 0/0          482 2017-09-21 19:37
↳ 3762e087ebbb895fd9c38981c1f7bfc76c9879fd3fdadef64df49e92721bb527/json
-rw-r--r-- 0/0          14848 2017-09-21 19:37
↳ 3762e087ebbb895fd9c38981c1f7bfc76c9879fd3fdadef64df49e92721bb527/layer.tar
-rw-r--r-- 0/0          4432 2017-09-21 19:37
↳ 577c1fe8e6d84360932b51767b65567550141af0801ff6d24ad10963e40472c5.json
drwxr-xr-x 0/0          0 2017-09-21 19:37
↳ 5bad884501c0e760bc0c9ca3ae3dca3f12c4abeb7d18194c364fec522b91b4f9/
-rw-r--r-- 0/0          3 2017-09-21 19:37
↳ 5bad884501c0e760bc0c9ca3ae3dca3f12c4abeb7d18194c364fec522b91b4f9/VERSION
-rw-r--r-- 0/0          482 2017-09-21 19:37
↳ 5bad884501c0e760bc0c9ca3ae3dca3f12c4abeb7d18194c364fec522b91b4f9/json
-rw-r--r-- 0/0          3072 2017-09-21 19:37
↳ 5bad884501c0e760bc0c9ca3ae3dca3f12c4abeb7d18194c364fec522b91b4f9/layer.tar
drwxr-xr-x 0/0          0 2017-09-21 19:37
↳ 81ce2fd011bc8241ae72eaae9146116b7c289e941467ff276397720171e6c576/
-rw-r--r-- 0/0          3 2017-09-21 19:37
↳ 81ce2fd011bc8241ae72eaae9146116b7c289e941467ff276397720171e6c576/VERSION
```

(continues on next page)

(continued from previous page)

```

-rw-r--r-- 0/0          406 2017-09-21 19:37 ↵
↳81ce2fd011bc8241ae72eaae9146116b7c289e941467ff276397720171e6c576/json
-rw-r--r-- 0/0          125649920 2017-09-21 19:37 ↵
↳81ce2fd011bc8241ae72eaae9146116b7c289e941467ff276397720171e6c576/layer.tar
drwxr-xr-x 0/0          0 2017-09-21 19:37 ↵
↳a10239905b060fd8b17ab31f37957bd126774f52f5280767d3b2639692913499/
-rw-r--r-- 0/0          3 2017-09-21 19:37 ↵
↳a10239905b060fd8b17ab31f37957bd126774f52f5280767d3b2639692913499/VERSION
-rw-r--r-- 0/0          482 2017-09-21 19:37 ↵
↳a10239905b060fd8b17ab31f37957bd126774f52f5280767d3b2639692913499/json
-rw-r--r-- 0/0          15872 2017-09-21 19:37 ↵
↳a10239905b060fd8b17ab31f37957bd126774f52f5280767d3b2639692913499/layer.tar
drwxr-xr-x 0/0          0 2017-09-21 19:37 ↵
↳ab6e1ca3392b2f4dbb60157cf99434b6975f37a767f530e293704a7348407634/
-rw-r--r-- 0/0          3 2017-09-21 19:37 ↵
↳ab6e1ca3392b2f4dbb60157cf99434b6975f37a767f530e293704a7348407634/VERSION
-rw-r--r-- 0/0          482 2017-09-21 19:37 ↵
↳ab6e1ca3392b2f4dbb60157cf99434b6975f37a767f530e293704a7348407634/json
-rw-r--r-- 0/0          5632 2017-09-21 19:37 ↵
↳ab6e1ca3392b2f4dbb60157cf99434b6975f37a767f530e293704a7348407634/layer.tar
-rw-r--r-- 0/0          574 1970-01-01 01:00 manifest.json

```

In other words, it is evident that this ‘tarball’ is a Docker-format image comprised of multiple layers along with metadata in a JSON manifest.

Through use of the `docker-archive` bootstrap agent, a SIF file (`lolcow_tar.sif`) for use by Singularity can be created via the following build command:

```

$ singularity build lolcow_tar.sif docker-archive://lolcow.tar
INFO: Starting build...
Getting image source signatures
Copying blob sha256:a2022691bf950a72f9d2d84d557183cb9eee07c065a76485f1695784855c5193
119.83 MiB / 119.83 MiB [=====] 6s
Copying blob sha256:ae620432889d2553535199dbdd8ba5a264ce85fcdcd5a430974d81fc27c02b45
15.50 KiB / 15.50 KiB [=====] 0s
Copying blob sha256:c561538251751e3685c7c6e7479d488745455ad7f84e842019dcb452c7b6fecc
14.50 KiB / 14.50 KiB [=====] 0s
Copying blob sha256:f96e6b25195f1b36ad02598b5d4381e41997c93ce6170cab1b81d9c68c514db0
5.50 KiB / 5.50 KiB [=====] 0s
Copying blob sha256:7f7a065d245a6501a782bf674f4d7e9d0a62fa6bd212edbf1f17bad0d5cd0bfc
3.00 KiB / 3.00 KiB [=====] 0s
Copying blob sha256:70ca7d49f8e9c44705431e3dade0636a2156300ae646ff4f09c904c138728839
116.56 MiB / 116.56 MiB [=====] 6s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: lolcow_tar.sif

```

There are two important differences in syntax evident in the above build command:

1. The `docker` part of the URI has been appended by `archive`. This ensures Singularity seek a Docker-format image archive stored locally as `lolcow.tar` to bootstrap the conversion process to SIF, as opposed to attempting to retrieve an image remotely hosted via Docker Hub.
2. `sudo` is *not* prepended to the build command for Singularity. This is *not* required if the executing user has the appropriate access privileges to the stored file.

**Note:** The `docker-archive` bootstrap agent handles archives (`.tar` files) as well as compressed archives (`.tar.gz`) when containers are built for Singularity via its `build` command.

---

**Note:** The Sylabs Cloud Remote Builder *does not* interoperate with locally stored Docker-format images; therefore, images cached locally by Docker, *cannot* be used to bootstrap creation of SIF files via the Remote Builder service. Of course, a SIF file could be created locally as detailed above. Then, in a separate, manual step, *pushed to the Sylabs Cloud Singularity Library*.

---

### 9.6.1.5 Pushing Locally Available Images to a Library

The outcome of bootstrapping from an image cached locally by Docker, or one stored locally as an archive, is of course a *locally* stored SIF file. As noted above, this is the *only* option available, as the Sylabs Cloud Remote Builder *does not* interoperate with the Docker daemon or locally stored archives in the Docker image format. Once produced, however, it may be desirable to make the resulting SIF file available through the Sylabs Cloud Singularity Library; therefore, the procedure to `push` a locally available SIF file to the Library is detailed here.

From the Sylabs Cloud Singularity Library, select `Create a new Project`. In this first of two steps, the publicly accessible project is created as illustrated below:

1 Choose Your Project      2 Push Image

Create a new Project

Owner:  / Project name:

Description(optional):

Public  
 Private

Initialize the project with a markdown README.

Next

Because an access token for the cloud service already exists, attention can be focused on the `push` command prototyped towards the bottom of the following screenshot:



1

Choose Your Project

2

**Push Image**

Push Image: You can use **Singularity CLI** or **Remote Builder** to push your image.

Singularity CLI

Remote Builder

## 1. Create Access Token

You already issued access token, you are good to go!

Create New Token

## 2. Input version detail of the image (what's new or release notes)

Image Tags:

Enter image version or tag

Image Description:

Enter image description (such as Release Note)

## 3. Push your image with Singularity CLI!

```
#Push with Singularity
$ singularity push image.sif library://ilumb/default/lolcow_tar
```

## 4. After you push your image, you can click Next to check the details of your images.

Next

In fact, by simply replacing `image.sif` with `lolcow_tar.sif`, the following upload is executed:

```
$ singularity push lolcow_tar.sif library://ilumb/default/lolcow_tar
INFO:   Now uploading lolcow_tar.sif to the library
 87.94 MiB / 87.94 MiB
↪ [=====] 100.
↪ 00% 1.25 MiB/s 1m10s
INFO:   Setting tag latest
```

Finally, from the perspective of the Library, the *hosted* version of the SIF file appears as illustrated below. Directions on how to *pull* this file are included from the portal.

## Project - library://ilumb/default/lolcow\_tar



Description Images

Push a new Image

latest

Edit

Unique ID: sha256.68fc39ebfb6f47d8960297119ae8d521556cfd6ddd584b4722b4106d0edf6d44

Created At: 2019/01/25 06:02:18

Release Notes:

```
# Pull with Singularity
$ singularity pull library://ilumb/default/lolcow_tar:latest
# Pull by unique ID (reproducible even if tags change)
$ singularity pull library://ilumb/default/lolcow_tar:sha256.68fc39ebfb6f47d8960297119ae8d521556cfd6ddd584b4722b4106d0edf6d44
```

[Direct Download lolcow\\_tar\\_latest.sif \(87.94 MB\)](#)
[Delete this version](#)

**Note:** The hosted version of the SIF file in the Sylabs Cloud Singularity Library is maintainable. In other words, if the image is updated locally, the update can be pushed to the Library and tagged appropriately.

## 9.6.2 Working with Definition Files

### 9.6.2.1 Mandatory Header Keywords: Remotely Bootstrapped

Akin to a set of blueprints that explain how to build a custom container, Singularity definition files (or “def files”) are considered in detail *elsewhere in this manual*. Therefore, only def file nuances specific to interoperability with Docker receive consideration here.

Singularity definition files are comprised of two parts - a **header** plus **sections**.

When working with repositories such as Docker Hub, Bootstrap and From are **mandatory** keywords within the header; for example, if the file `lolcow.def` has contents

```
Bootstrap: docker
From: godlovedc/lolcow
```

then

```
sudo singularity build lolcow.sif lolcow.def
```

creates a Singularity container in SIF by bootstrapping from the public `godlovedc/lolcow` image from Docker Hub.

In the above definition file, `docker` is one of numerous, possible bootstrap agents; this, and other bootstrap agents receive attention *in the appendix*.

Through *the means for authentication described above*, definition files permit use of private images hosted via Docker Hub. For example, if the file `mylolcow.def` has contents

```
Bootstrap: docker
From: ilumb/mylolcow
```

then

```
sudo singularity build --docker-login mylolcow.sif mylolcow.def
```

creates a Singularity container in SIF by bootstrapping from the *private* `ilumb/mylolcow` image from Docker Hub after successful *interactive authentication*.

Alternatively, if *environment variables have been set as above*, then

```
$ sudo -E singularity build mylolcow.sif mylolcow.def
```

enables authenticated use of the private image.

**Note:** The `-E` option is required to preserve the user's existing environment variables upon `sudo` invocation - a privileged escalation *required* to create Singularity containers via the `build` command.

### 9.6.2.2 Remotely Bootstrapped and Built Containers

Consider again *the definition file used the outset of the section above*:

```
Bootstrap: docker
From: godlovedc/lolcow
```

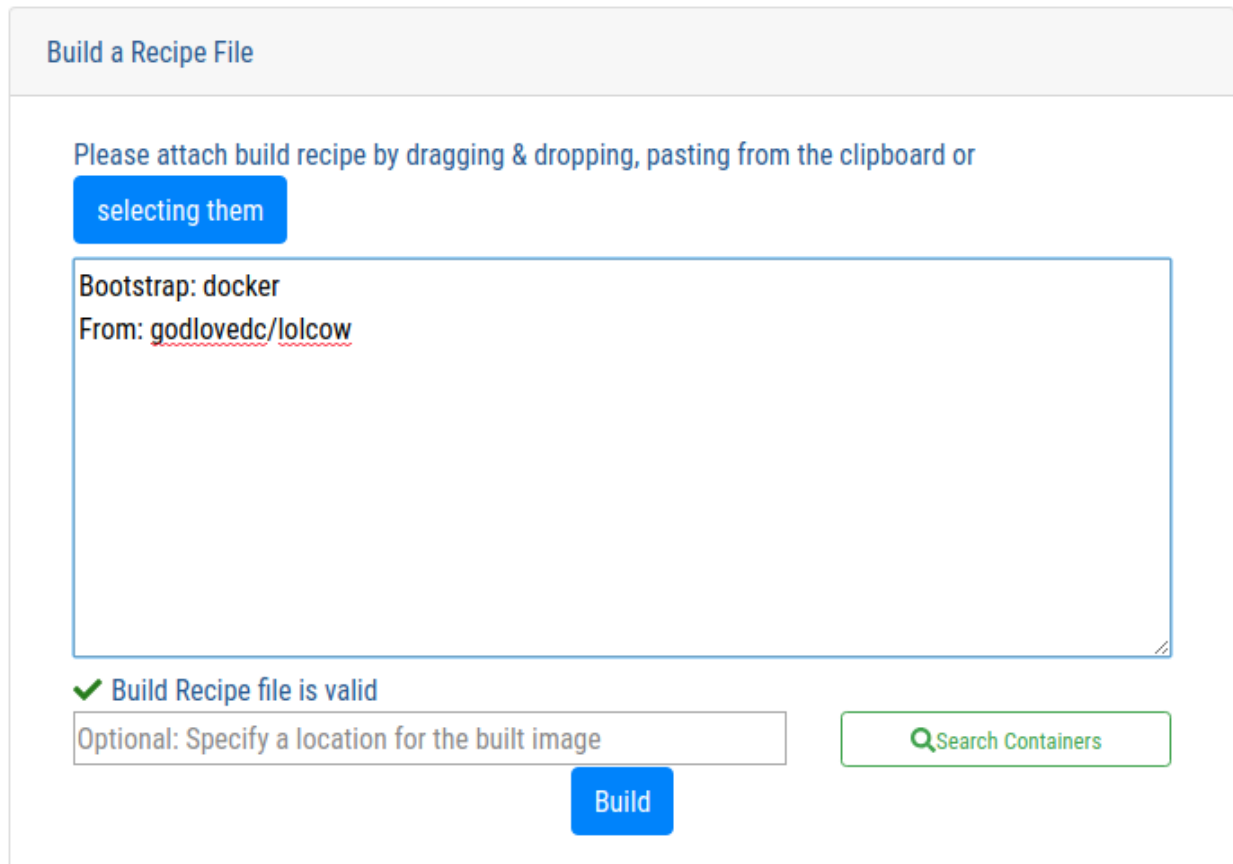
With two small adjustments to the Singularity `build` command, the Sylabs Cloud Remote Builder can be utilized:

```
$ singularity build --remote lolcow_rb_def.sif lolcow.def
searching for available build agent.....INFO: Starting build...
Getting image source signatures
Copying blob sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
 45.33 MiB / 45.33 MiB 0s
Copying blob sha256:3b61febd4ae9e982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
 848 B / 848 B 0s
Copying blob sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
 621 B / 621 B 0s
Copying blob sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
 853 B / 853 B 0s
Copying blob sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
 169 B / 169 B 0s
Copying blob sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
 53.75 MiB / 53.75 MiB 0s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: /tmp/image-994007654
INFO: Now uploading /tmp/image-994007654 to the library
 87.94 MiB / 87.94 MiB 100.00% 41.76 MiB/s 2s
INFO: Setting tag latest
 87.94 MiB / 87.94 MiB
↪[=====]
↪100.00% 19.08 MiB/s 4s
```

In the above, `--remote` has been added as the `build` option that causes use of the Remote Builder service. A much more subtle change, however, is the *absence* of `sudo` ahead of `singularity build`. Though subtle here, this absence is notable, as users can build containers via the Remote Builder with *escalated privileges*; in other words,

steps in container creation that *require* `root` access *are* enabled via the Remote Builder even for (DevOps) users *without* administrative privileges locally.

In addition to the command-line support described above, the Sylabs Cloud Remote Builder also allows definition files to be copied and pasted into its Graphical User Interface (GUI). After pasting a definition file, and having that file validated by the service, the build-centric part of the GUI appears as illustrated below. By clicking on the `Build` button, creation of the container is initiated.



**Build a Recipe File**

Please attach build recipe by dragging & dropping, pasting from the clipboard or selecting them

```
Bootstrap: docker
From: godlovedc/lolcow
```

✓ Build Recipe file is valid

Optional: Specify a location for the built image

Search Containers

Build

Once the build process has been completed, the corresponding SIF file can be retrieved from the service - as shown below. A log file for the `build` process is provided by the GUI, and made available for download as a text file (not shown here).

### Output - `builder://5c49ecbce21266000194b336`

<b>Success</b>	Build Definition: docker:godlovedc/lolcow	Duration: 01m 34s	Library: itumb/remote-builds/rb-5c49ecbce21266000194b336	Download: <a href="#">Direct Download rb-5c49ecbce21266000194b336_sha256_e96...sif (87.94 MB)</a>
----------------	--	----------------------	---	--

A copy of the SIF file created by the service remains in the Sylabs Cloud Singularity Library as illustrated below.

## User/Group - library://ilumb

Description:

Projects for this User/Group:

Name	Description	Private	Downloads	Size	Stars	
> <a href="#">rb-5c49eedce21266000194b337</a>	No description	true	0	1 images, 87.94 MB	☆ 0	
> <a href="#">rb-5c49ecbce21266000194b336</a>	No description	true	0	1 images, 87.94 MB	☆ 0	
> <a href="#">rb-5c49e852e21266000194b335</a>	No description	true	1	1 images, 87.94 MB	☆ 0	
> <a href="#">rb-5c48a350e21266000194b313</a>	No description	true	1	1 images, 87.94 MB	☆ 0	
> <a href="#">rb-5c4896d9e21266000194b30f</a>	No description	true	1	1 images, 87.94 MB	☆ 0	

**Note:** The Sylabs Cloud is currently available as an Alpha Preview. In addition to the Singularity Library and Remote Builder, a Keystore service is also available. All three services make use of a *freemium* pricing model in supporting Singularity Community Edition. In contrast, all three services are included in SingularityPRO - an enterprise grade subscription for Singularity that is offered for a fee from Sylabs. For additional details regarding the different offerings available for Singularity, please [consult the Sylabs website](#).

### 9.6.2.3 Mandatory Header Keywords: Locally Bootstrapped

When `docker-daemon` is the bootstrap agent in a Singularity definition file, SIF containers can be created from images cached locally by Docker. Suppose the definition file `lolcow-d.def` has contents:

```
Bootstrap: docker-daemon
From: godlovedc/lolcow:latest
```

**Note:** Again, the image tag `latest` is **required** when bootstrapping creation of a container for Singularity from an image locally cached by Docker.

Then,

```
$ sudo singularity build lolcow_from_docker_cache.sif lolcow-d.def
Build target already exists. Do you want to overwrite? [N/y] y
INFO: Starting build...
Getting image source signatures
Copying blob sha256:a2022691bf950a72f9d2d84d557183cb9eee07c065a76485f1695784855c5193
119.83 MiB / 119.83 MiB [=====] 6s
Copying blob sha256:ae620432889d2553535199dbdd8ba5a264ce85fcdcd5a430974d81fc27c02b45
15.50 KiB / 15.50 KiB [=====] 0s
Copying blob sha256:c561538251751e3685c7c6e7479d488745455ad7f84e842019dcb452c7b6fecc
14.50 KiB / 14.50 KiB [=====] 0s
Copying blob sha256:f96e6b25195f1b36ad02598b5d4381e41997c93ce6170cab1b81d9c68c514db0
5.50 KiB / 5.50 KiB [=====] 0s
Copying blob sha256:7f7a065d245a6501a782bf674f4d7e9d0a62fa6bd212edbf1f17bad0d5cd0bfc
```

(continues on next page)

(continued from previous page)

```

3.00 KiB / 3.00 KiB [=====] 0s
Copying blob sha256:70ca7d49f8e9c44705431e3dade0636a2156300ae646ff4f09c904c138728839
116.56 MiB / 116.56 MiB [=====] 6s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: lolcow_from_docker_cache.sif

```

In other words, this is the definition-file counterpart to *the command-line invocation provided above*.

**Note:** The `sudo` requirement in the above build request originates from Singularity; it is the standard requirement when use is made of definition files. In other words, membership of the issuing user in the `docker` Linux group is of no consequence in this context.

Alternatively when `docker-archive` is the bootstrap agent in a Singularity definition file, SIF containers can be created from images stored locally by Docker. Suppose the definition file `lolcow-da.def` has contents:

```

Bootstrap: docker-archive
From: lolcow.tar

```

Then,

```

$ sudo singularity build lolcow_tar_def.sif lolcow-da.def
INFO:   Starting build...
Getting image source signatures
Copying blob sha256:a2022691bf950a72f9d2d84d557183cb9eee07c065a76485f1695784855c5193
119.83 MiB / 119.83 MiB [=====] 6s
Copying blob sha256:ae620432889d2553535199dbdd8ba5a264ce85fcdcd5a430974d81fc27c02b45
15.50 KiB / 15.50 KiB [=====] 0s
Copying blob sha256:c561538251751e3685c7c6e7479d488745455ad7f84e842019dcb452c7b6fecc
14.50 KiB / 14.50 KiB [=====] 0s
Copying blob sha256:f96e6b25195f1b36ad02598b5d4381e41997c93ce6170cab1b81d9c68c514db0
5.50 KiB / 5.50 KiB [=====] 0s
Copying blob sha256:7f7a065d245a6501a782bf674f4d7e9d0a62fa6bd212edb1f17bad0d5cd0bfc
3.00 KiB / 3.00 KiB [=====] 0s
Copying blob sha256:70ca7d49f8e9c44705431e3dade0636a2156300ae646ff4f09c904c138728839
116.56 MiB / 116.56 MiB [=====] 6s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: lolcow_tar_def.sif

```

through build results in the SIF file `lolcow_tar_def.sif`. In other words, this is the definition-file counterpart to *the command-line invocation provided above*.

#### 9.6.2.4 Optional Header Keywords

In the two-previous examples, the `From` keyword specifies *both* the user and `repo-name` in making use of Docker Hub. *Optional* use of `Namespace` permits the more-granular split across two keywords:

```

Bootstrap: docker
Namespace: godlovedc
From: lolcow

```

**Note:** In their documentation, “`Docker ID namespace`” and `user` are employed as synonyms in the text and examples, respectively.

**Note:** The default value for the optional keyword `Namespace` is `library`.

### 9.6.2.5 Private Images and Registries

Thus far, use of Docker Hub has been assumed. To make use of a different repository of Docker images the **optional** `Registry` keyword can be added to the Singularity definition file. For example, to make use of a Docker image from the NVIDIA GPU Cloud (NGC) corresponding definition file is:

```

Bootstrap: docker
From: nvidia/pytorch:18.11-py3
Registry: nvcr.io

```

This def file `ngc_pytorch.def` can be passed as a specification to build as follows:

```

$ sudo singularity build --docker-login mypytorch.sif ngc_pytorch.def
Enter Docker Username: $oauthtoken
Enter Docker Password: <obscured>
INFO:      Starting build...
Getting image source signatures
Copying blob sha256:18d680d616571900d78ee1c8fff0310f2a2afe39c6ed0ba2651ff667af406c3e
 41.34 MiB / 41.34 MiB [=====] 2s
<blob copying details deleted>
Copying config sha256:b77551af8073c85588088ab2a39007d04bc830831ba1eef4127b2d39aaf3a6b1
21.28 KiB / 21.28 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:      Creating SIF file...
INFO:      Build complete: mypytorch.sif

```

After successful authentication via interactive use of the `--docker-login` option, output as the SIF container `mypytorch.sif` is (ultimately) produced. As above, *use of environment variables* is another option available for authenticating private Docker type repositories such as NGC; once set, the build command is as above save for the absence of the `--docker-login` option.

### 9.6.2.6 Directing Execution

The Dockerfile corresponding to `godlovedc/lolcow` (and available [here](#)) is as follows:

```

FROM ubuntu:16.04

RUN apt-get update && apt-get install -y fortune cowsay lolcat

ENV PATH /usr/games:${PATH}

```

(continues on next page)

(continued from previous page)

```
ENV LC_ALL=C

ENTRYPOINT fortune | cowsay | lolcat
```

The execution-specific part of this Dockerfile is the ENTRYPOINT - “... an optional definition for the first part of the command to be run ...” according to [the available documentation](#). After conversion to SIF, execution of `fortune | cowsay | lolcat` *within* the container produces the output:

```
$ ./mylolcow.sif

/ Q: How did you get into artificial \
| intelligence? A: Seemed logical -- I |
\ didn't have any real intelligence. /
-----
      \   ^__^
       \  (oo)\_______
            (__)\       )\/\
                ||----w |
                ||     ||
```

In addition, CMD allows an arbitrary string to be *appended* to the ENTRYPOINT. Thus, multiple commands or flags can be passed together through combined use.

Suppose now that a Singularity `%runscript` **section** is added to the definition file as follows:

```
Bootstrap: docker
Namespace: godlovedc
From: lolcow

%runscript

    fortune
```

After conversion to SIF via the Singularity `build` command, execution of the resulting container produces the output:

```
$ ./lolcow.sif
This was the most unkindest cut of all.
    -- William Shakespeare, "Julius Caesar"
```

In other words, introduction of a `%runscript` section into the Singularity definition file causes the ENTRYPOINT of the Dockerfile to be *bypassed*. The presence of the `%runscript` section would also bypass a CMD entry in the Dockerfile.

To *preserve* use of ENTRYPOINT and/or CMD as defined in the Dockerfile, the `%runscript` section must be *absent* from the Singularity definition. In this case, and to favor execution of CMD *over* ENTRYPOINT, a non-empty assignment of the *optional* IncludeCmd should be included in the header section of the Singularity definition file as follows:

```
Bootstrap: docker
Namespace: godlovedc
From: lolcow
IncludeCmd: yes
```

---

**Note:** Because only a non-empty IncludeCmd is required, *either* `yes` (as above) or `no` results in execution of CMD *over* ENTRYPOINT.

---



To summarize execution precedence:

1. If present, the `%runscript` section of the Singularity definition file is executed
2. If `IncludeCmd` is a non-empty keyword entry in the header of the Singularity definition file, then `CMD` from the `Dockerfile` is executed
3. If present in the `Dockerfile`, `ENTRYPOINT` appended by `CMD` (if present) are executed in sequence
4. Execution of the `bash` shell is defaulted to

### 9.6.2.7 Container Metadata

Singularity's `inspect` command displays container metadata - data about data that is encapsulated *within* a SIF file. Default output (assumed via the `--labels` option) from the command was *illustrated above*. `inspect`, however, provides a number of options that are *detailed elsewhere*; in the remainder of this section, Docker-specific use to establish execution precedence is emphasized.

As stated above (i.e., *the first case of execution precedence*), the very existence of a `%runscript` section in a Singularity definition file *takes precedence* over commands that might exist in the `Dockerfile`.

When the `%runscript` section is *removed* from the Singularity definition file, the result is (once again):

```
$ singularity inspect --deffile lolcow.sif

from: lolcow
bootstrap: docker
namespace: godlovedc
```

The runscript 'inherited' from the `Dockerfile` is:

```
$ singularity inspect --runscript lolcow.sif

#!/bin/sh
OCI_ENTRYPOINT="/bin/sh" "-c" "fortune | cowsay | lolcat"
OCI_CMD=''
# ENTRYPOINT only - run entrypoint plus args
if [ -z "$OCI_CMD" ] && [ -n "$OCI_ENTRYPOINT" ]; then
    SINGULARITY_OCI_RUN="${OCI_ENTRYPOINT} $@"
fi

# CMD only - run CMD or override with args
if [ -n "$OCI_CMD" ] && [ -z "$OCI_ENTRYPOINT" ]; then
    if [ $# -gt 0 ]; then
        SINGULARITY_OCI_RUN="$@"
    else
        SINGULARITY_OCI_RUN="${OCI_CMD}"
    fi
fi

# ENTRYPOINT and CMD - run ENTRYPOINT with CMD as default args
# override with user provided args
if [ $# -gt 0 ]; then
    SINGULARITY_OCI_RUN="${OCI_ENTRYPOINT} $@"
else
    SINGULARITY_OCI_RUN="${OCI_ENTRYPOINT} ${OCI_CMD}"
fi

eval ${SINGULARITY_OCI_RUN}
```

From this Bourne shell script, it is evident that only an ENTRYPOINT is detailed in the Dockerfile; thus the ENTRYPOINT only – run entrypoint plus args conditional block is executed. In this case then, *the third case of execution precedence* has been illustrated.

The above Bourne shell script also illustrates how the following scenarios will be handled:

- A CMD only entry in the Dockerfile
- **Both** ENTRYPOINT *and* CMD entries in the Dockerfile

From this level of detail, use of ENTRYPOINT *and/or* CMD in a Dockerfile has been made **explicit**. These remain examples within *the third case of execution precedence*.

## 9.7 OCI Image Support

### 9.7.1 Overview

OCI is an acronym for the [Open Containers Initiative](#) - an independent organization whose mandate is to develop open standards relating to containerization. To date, standardization efforts have focused on container formats and runtimes; it is the former that is emphasized here. Stated simply, an **OCI blob** is content that can be addressed; in other words, *each* layer of a Docker image is rendered as an OCI blob as illustrated in the (revisited) pull example below.

---

**Note:** To facilitate interoperability with Docker Hub, the Singularity core makes use of the `containers/image library` - "... a set of Go libraries aimed at working in various way[s] with containers' images and container image registries."

---

#### 9.7.1.1 Image Pulls Revisited

After describing various *action commands that could be applied to images hosted remotely via Docker Hub*, the notion of having *a local copy in Singularity's native format for containerization (SIF)* was introduced:

```
$ singularity pull docker://godlovedc/lolcow
INFO:   Starting build...
Getting image source signatures
Copying blob sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
 45.33 MiB / 45.33 MiB [=====] 1s
Copying blob sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
 848 B / 848 B [=====] 0s
Copying blob sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
 621 B / 621 B [=====] 0s
Copying blob sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
 853 B / 853 B [=====] 0s
Copying blob sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
 169 B / 169 B [=====] 0s
Copying blob sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
 53.75 MiB / 53.75 MiB [=====] 2s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: lolcow_latest.sif
```

Thus use of Singularity's `pull` command results in the *local* file copy in SIF, namely `lolcow_latest.sif`. Layers of the image from Docker Hub are copied locally as OCI blobs.

### 9.7.1.2 Image Caching in Singularity

If the *same* `pull` command is issued a *second* time, the output is different:

```
$ singularity pull docker://godlovedc/lolcow
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↪sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
Skipping fetch of repeat blob_
↪sha256:3b61febd4ae9e982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
Skipping fetch of repeat blob_
↪sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
Skipping fetch of repeat blob_
↪sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
Skipping fetch of repeat blob_
↪sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
Skipping fetch of repeat blob_
↪sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
 3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: lolcow_latest.sif
```

As the copy operation has clearly been *skipped*, it is evident that a copy of all OCI blobs **must** be cached locally. Indeed, Singularity has made an entry in its local cache as follows:

```
$ tree .singularity/
.singularity/
├── cache
│   └── oci
│       ├── blobs
│       │   └── sha256
│       │       ├── 3b61febd4ae9e982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
│       │       ├── 73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
│       │       ├── 7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
│       │       ├── 8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
│       │       ├── 9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
│       │       ├── 9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
│       │       ├── d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
│       │       └── f2a852991b0a36a9f3d6b2a33b98a461e9ede8393482f0deb5287afcbae2ce10
│       ├── index.json
│       └── oci-layout
4 directories, 10 files
```

### 9.7.1.3 Compliance with the OCI Image Layout Specification

From the perspective of the directory `$HOME/.singularity/cache/oci`, this cache implementation in Singularity complies with the [OCI Image Layout Specification](#):

- blobs directory - contains content addressable data, that is otherwise considered opaque
- oci-layout file - a mandatory JSON object file containing both mandatory and optional content
- index.json file - a mandatory JSON object file containing an index of the images

Because one or more images is 'bundled' here, the directory `$HOME/.singularity/cache/oci` is referred to as the `$OCI_BUNDLE_DIR`.

For additional details regarding this specification, consult the [OCI Image Format Specification](#).

### 9.7.1.4 OCI Compliance and the Singularity Cache

As required by the layout specification, OCI blobs are *uniquely* named by their contents:

```
$ shasum -a 256 ./blobs/sha256/
↪9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118  ./blobs/sha256/
↪9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
```

They are also otherwise opaque:

```
$ file ./blobs/sha256/
↪9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118  ./blobs/sha256/
↪9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118: gzip compressed_
↪data
```

The content of the `oci-layout` file in this example is:

```
$ cat oci-layout | jq
{
  "imageLayoutVersion": "1.0.0"
}
```

This is as required for compliance with the layout standard.

---

**Note:** In rendering the above JSON object files, use has been made of `jq` - the command-line JSON processor.

---

The index of images in this case is:

```
$ cat index.json | jq
{
  "schemaVersion": 2,
  "manifests": [
    {
      "mediaType": "application/vnd.oci.image.manifest.v1+json",
      "digest":
↪"sha256:f2a852991b0a36a9f3d6b2a33b98a461e9ede8393482f0deb5287afcbae2ce10",
      "size": 1125,
      "annotations": {
        "org.opencontainers.image.ref.name":
↪"a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb"
      },
      "platform": {
        "architecture": "amd64",
        "os": "linux"
      }
    }
  ]
}
```

(continues on next page)

(continued from previous page)

```

    }
  ]
}

```

The digest blob in this index file includes the details for all of the blobs that collectively comprise the `godlovedc/lolcow` image:

```

$ cat ./blobs/sha256/
↪f2a852991b0a36a9f3d6b2a33b98a461e9ede8393482f0deb5287afcbae2ce10 | jq
{
  "schemaVersion": 2,
  "config": {
    "mediaType": "application/vnd.oci.image.config.v1+json",
    "digest": "sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
↪",
    "size": 3410
  },
  "layers": [
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "digest":
↪"sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118",
      "size": 47536248
    },
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "digest":
↪"sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a",
      "size": 848
    },
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "digest":
↪"sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2",
      "size": 621
    },
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "digest":
↪"sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e",
      "size": 853
    },
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "digest":
↪"sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9",
      "size": 169
    },
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "digest":
↪"sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945",
      "size": 56355961
    }
  ]
}

```

The digest blob referenced in the `index.json` file references the following configuration file:

```
$ cat ./blobs/sha256/73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82_
→ | jq
{
  "created": "2017-09-21T18:37:47.278336798Z",
  "architecture": "amd64",
  "os": "linux",
  "config": {
    "Env": [
      "PATH=/usr/games:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
      "LC_ALL=C"
    ],
    "Entrypoint": [
      "/bin/sh",
      "-c",
      "fortune | cowsay | lolcat"
    ]
  },
  "rootfs": {
    "type": "layers",
    "diff_ids": [
      "sha256:a2022691bf950a72f9d2d84d557183cb9eee07c065a76485f1695784855c5193",
      "sha256:ae620432889d2553535199dbdd8ba5a264ce85fcdcd5a430974d81fc27c02b45",
      "sha256:c561538251751e3685c7c6e7479d488745455ad7f84e842019dcb452c7b6fecc",
      "sha256:f96e6b25195f1b36ad02598b5d4381e41997c93ce6170cab1b81d9c68c514db0",
      "sha256:7f7a065d245a6501a782bf674f4d7e9d0a62fa6bd212edbf1f17bad0d5cd0bfc",
      "sha256:70ca7d49f8e9c44705431e3dade0636a2156300ae646ff4f09c904c138728839"
    ]
  },
  "history": [
    {
      "created": "2017-09-18T23:31:37.453092323Z",
      "created_by": "/bin/sh -c #(nop) ADD_
→ file:5ed435208da6621b45db657dd6549ee132cde58c4b6763920030794c2f31fbc0 in / "
    },
    {
      "created": "2017-09-18T23:31:38.196268404Z",
      "created_by": "/bin/sh -c set -xe \t\t&& echo '#!/bin/sh' > /usr/sbin/policy-rc.
→ d \t\t&& echo 'exit 101' >> /usr/sbin/policy-rc.d \t\t&& chmod +x /usr/sbin/policy-rc.d_
→ \t\t&& dpkg-divert --local --rename --add /sbin/initctl \t\t&& cp -a /usr/sbin/policy-
→ rc.d /sbin/initctl \t\t&& sed -i 's/^exit.*/exit 0/' /sbin/initctl \t\t&& echo 'force-
→ unsafe-io' > /etc/dpkg/dpkg.cfg.d/docker-apt-speedup \t\t&& echo 'DPkg::Post-Invoke
→ { \"rm -f /var/cache/apt/archives/*.deb /var/cache/apt/archives/partial/*.deb /var/
→ cache/apt/*.bin || true\"; };' > /etc/apt/apt.conf.d/docker-clean \t\t&& echo
→ 'APT::Update::Post-Invoke { \"rm -f /var/cache/apt/archives/*.deb /var/cache/apt/
→ archives/partial/*.deb /var/cache/apt/*.bin || true\"; };' >> /etc/apt/apt.conf.d/
→ docker-clean \t\t&& echo 'Dir::Cache::pkgcache \"\"; Dir::Cache::srcpkgcache \"\"; ' >>
→ /etc/apt/apt.conf.d/docker-clean \t\t&& echo 'Acquire::Languages \"none\";' > /etc/
→ apt/apt.conf.d/docker-no-languages \t\t&& echo 'Acquire::GzipIndexes \"true\";_
→ Acquire::CompressionTypes::Order:: \"gz\";' > /etc/apt/apt.conf.d/docker-gzip-
→ indexes \t\t&& echo 'Apt::AutoRemove::SuggestsImportant \"false\";' > /etc/apt/apt.
→ conf.d/docker-autoremove-suggests"
    },
    {
      "created": "2017-09-18T23:31:38.788043199Z",
      "created_by": "/bin/sh -c rm -rf /var/lib/apt/lists/*"
    }
  ],
}
```

(continues on next page)

(continued from previous page)

```

{
  "created": "2017-09-18T23:31:39.411670721Z",
  "created_by": "/bin/sh -c sed -i 's/^#\s*\((deb.*universe\))$/\1/g' /etc/apt/
↪sources.list"
},
{
  "created": "2017-09-18T23:31:40.055188541Z",
  "created_by": "/bin/sh -c mkdir -p /run/systemd && echo 'docker' > /run/systemd/
↪container"
},
{
  "created": "2017-09-18T23:31:40.215057796Z",
  "created_by": "/bin/sh -c #(nop)  CMD [\"/bin/bash\"]",
  "empty_layer": true
},
{
  "created": "2017-09-21T18:37:46.483638061Z",
  "created_by": "/bin/sh -c apt-get update && apt-get install -y fortune cowsay_
↪lolcat"
},
{
  "created": "2017-09-21T18:37:47.041333952Z",
  "created_by": "/bin/sh -c #(nop)  ENV PATH=/usr/games:/usr/local/sbin:/usr/
↪local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
  "empty_layer": true
},
{
  "created": "2017-09-21T18:37:47.170535967Z",
  "created_by": "/bin/sh -c #(nop)  ENV LC_ALL=C",
  "empty_layer": true
},
{
  "created": "2017-09-21T18:37:47.278336798Z",
  "created_by": "/bin/sh -c #(nop)  ENTRYPOINT [\"/bin/sh\" \"-c\" \"fortune |_
↪cowsay | lolcat\""],
  "empty_layer": true
}
]
}

```

Even when all OCI blobs are already in Singularity's local cache, repeated image pulls cause *both* these last-two JSON object files, as well as the `oci-layout` and `index.json` files, to be updated.

## 9.7.2 Building Containers for Singularity from OCI Images

### 9.7.2.1 Working Locally from the Singularity Command Line: `oci Bootstrap Agent`

The example detailed in the previous section can be used to illustrate how a SIF file for use by Singularity can be created from the local cache - an albeit contrived example, that works because the Singularity cache is compliant with the OCI Image Layout Specification.

---

**Note:** Of course, the `oci` bootstrap agent can be applied to *any* **bundle** that is compliant with the OCI Image Layout Specification - not *just* the Singularity cache, as created by executing a Singularity `pull` command.

---

In this local case, the build command of Singularity makes use of the `oci` bootstrap agent as follows:

```
$ singularity build ~/lolcow_oci_cache.sif oci://$HOME/.singularity/cache/
↳oci:a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob
↳sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
Skipping fetch of repeat blob
↳sha256:3b61febd4ae9e982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
Skipping fetch of repeat blob
↳sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
Skipping fetch of repeat blob
↳sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
Skipping fetch of repeat blob
↳sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
Skipping fetch of repeat blob
↳sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: /home/vagrant/lolcow_oci_cache.sif
```

As can be seen, this results in the SIF file `lolcow_oci_cache.sif` in the user's home directory.

The syntax for the `oci` bootstrap agent requires some elaboration, however. In this case, and as illustrated above, `$HOME/.singularity/cache/oci` has content:

```
$ ls
blobs  index.json  oci-layout
```

In other words, it is the `$OCI_BUNDLE_DIR` containing the data and metadata that collectively comprise the image layed out in accordance with the OCI Image Layout Specification *as discussed previously* - the same data and metadata that are assembled into a single SIF file through the `build` process. However,

```
$ singularity build ~/lolcow_oci_cache.sif oci://$HOME/.singularity/cache/oci
INFO: Starting build...
FATAL: While performing build: conveyor failed to get: more than one image in oci,
↳choose an image
```

does not *uniquely* specify an image from which to bootstrap the build process. In other words, there are multiple images referenced via `org.opencontainers.image.ref.name` in the `index.json` file. By appending `:a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb` to `oci` in this example, the image is uniquely specified, and the container created in SIF (as illustrated previously).

---

**Note:** Executing the Singularity `pull` command multiple times on the same image produces multiple `org.opencontainers.image.ref.name` entries in the `index.json` file. Appending the value of the unique `org.opencontainers.image.ref.name` allows for use of the `oci` bootstrap agent.

---

### 9.7.2.2 Working Locally from the Singularity Command Line: `oci-archive` Bootstrap Agent

OCI archives, i.e., `tar` files obeying the OCI Image Layout Specification *as discussed previously*, can seed creation of a container for Singularity. In this case, use is made of the `oci-archive` bootstrap agent.



To illustrate this agent, it is convenient to build the archive from the Singularity cache. After a single pull of the `godlovedc/lolcow` image from Docker Hub, a tar format archive can be generated from the `$HOME/.singularity/cache/oci` directory as follows:

```
$ tar cvf $HOME/godlovedc_lolcow.tar *
blobs/
blobs/sha256/
blobs/sha256/73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
blobs/sha256/8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
blobs/sha256/9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
blobs/sha256/3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
blobs/sha256/9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
blobs/sha256/d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
blobs/sha256/f2a852991b0a36a9f3d6b2a33b98a461e9ede8393482f0deb5287afcbae2ce10
blobs/sha256/7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
index.json
oci-layout
```

The native container `lolcow_oci_tarfile.sif` for use by Singularity can be created by issuing the `build` command as follows:

```
$ singularity build lolcow_oci_tarfile.sif oci-archive://godlovedc_lolcow.tar
Build target already exists. Do you want to overwrite? [N/y] y
INFO: Starting build...
Getting image source signatures
Skipping fetch of repeat blob_
↳sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
Skipping fetch of repeat blob_
↳sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
Skipping fetch of repeat blob_
↳sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
Skipping fetch of repeat blob_
↳sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
Skipping fetch of repeat blob_
↳sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
Skipping fetch of repeat blob_
↳sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: lolcow_oci_tarfile.sif
```

This assumes that the tar file exists in the current working directory.

---

**Note:** Cache maintenance is a manual process at the current time. In other words, the cache can be cleared by **carefully** issuing the command `rm -rf $HOME/.singularity/cache`. Of course, this will clear the local cache of all downloaded images.

---



---

**Note:** Because the layers of a Docker image as well as the blobs of an OCI image are already `gzip` compressed, there is a minimal advantage to having compressed archives representing OCI images. For this reason, the `build` detailed above bootstraps a SIF file for use by Singularity from only a tar file, and not a `tar.gz` file.

---

### 9.7.2.3 Working from the Singularity Command Line with Remotely Hosted Images

In the previous section, an OCI archive was created from locally available OCI blobs and metadata; the resulting `tar` file served to bootstrap the creation of a container for Singularity in SIF via the `oci-archive` agent. Typically, however, OCI archives of interest are remotely hosted. Consider, for example, an Alpine Linux OCI archive stored in Amazon S3 storage. Because such an archive can be retrieved via secure HTTP, the following `pull` command results in a local copy as follows:

```
$ singularity pull https://s3.amazonaws.com/singularity-ci-public/alpine-oci-archive.
↳tar
 1.98 MiB / 1.98 MiB
↳[=====]
↳100.00% 7.48 MiB/s 0s
```

Thus `https` (and `http`) are additional bootstrap agents available to seed development of containers for Singularity.

It is worth noting that the OCI image specification compliant contents of this archive are:

```
$ tar tvf alpine-oci-archive.tar
drwxr-xr-x 1000/1000      0 2018-06-25 14:45 blobs/
drwxr-xr-x 1000/1000      0 2018-06-25 14:45 blobs/sha256/
-rw-r--r-- 1000/1000      585 2018-06-25 14:45 blobs/sha256/
↳bla7f144ece0194921befe57ab30ed1fd98c5950db7996719429020986092058
-rw-r--r-- 1000/1000      348 2018-06-25 14:45 blobs/sha256/
↳d0ff39a54244ba25ac7447f19941765bee97b05f37ceb438a72e80c9ed39854a
-rw-r--r-- 1000/1000    2065537 2018-06-25 14:45 blobs/sha256/
↳ff3a5c916c92643ff77519ffa742d3ec61b7f591b6b7504599d95a4a41134e28
-rw-r--r-- 1000/1000      296 2018-06-25 14:45 index.json
-rw-r--r-- 1000/1000      31 2018-06-25 14:45 oci-layout
```

Proceeding as before, for a (now) locally available OCI archive, a SIF file can be produced by executing:

```
$ singularity build alpine_oci_archive.sif oci-archive://alpine-oci-archive.tar
INFO: Starting build...
Getting image source signatures
Copying blob sha256:ff3a5c916c92643ff77519ffa742d3ec61b7f591b6b7504599d95a4a41134e28
 1.97 MiB / 1.97 MiB [=====] 0s
Copying config sha256:bla7f144ece0194921befe57ab30ed1fd98c5950db7996719429020986092058
 585 B / 585 B [=====] 0s
Writing manifest to image destination
Storing signatures
INFO: Creating SIF file...
INFO: Build complete: alpine_oci_archive.sif
```

The resulting SIF file can be validated as follows, for example:

```
$ ./alpine_oci_archive.sif
Singularity> cat /etc/os-release
NAME="Alpine Linux"
ID=alpine
VERSION_ID=3.7.0
PRETTY_NAME="Alpine Linux v3.7"
HOME_URL="http://alpinelinux.org"
BUG_REPORT_URL="http://bugs.alpinelinux.org"
Singularity>
$
```

---

**Note:** The `http` and `https` bootstrap agents can only be used to pull OCI archives from where they are hosted.

In working with remotely hosted OCI image archives then, a two-step workflow is *required* to produce SIF files for native use by Singularity:

1. Transfer of the image to local storage via the `https` (or `http`) bootstrap agent. The Singularity `pull` command achieves this.
  2. Creation of a SIF file via the `oci-archive` bootstrap agent. The Singularity `build` command achieves this.
- 

**Note:** Though a frequently asked question, the distribution of OCI images remains *out of scope*. In other words, there is no OCI endorsed distribution method or registry.

Established with nothing more than a Web server then, any individual, group or organization, *could* host OCI archives. This might be particularly appealing, for example, for organizations having security requirements that preclude access to public registries such as Docker Hub. Other than having a very basic hosting capability, OCI archives need only comply to the OCI Image Layout Specification *as discussed previously*.

---

#### 9.7.2.4 Working with Definition Files: Mandatory Header Keywords

Three, new bootstrap agents have been introduced as a consequence of compliance with the OCI Image Specification - assuming `http` and `https` are considered together. In addition to bootstrapping images for Singularity completely from the command line, definition files can be employed.

As *above*, the OCI image layout compliant Singularity cache can be employed to create SIF containers; the definition file, `lolcow-oci.def`, equivalent is:

```
Bootstrap: oci
From: .singularity/cache/
↳oci:a692b57abc43035b197b10390ea2c12855d21649f2ea2cc28094d18b93360eeb
```

Recall that the colon-appended string in this file uniquely specifies the `org.opencontainers.image.ref.name` of the desired image, as more than one possibility exists in the `index.json` file. The corresponding build command is:

```
$ sudo singularity build ~/lolcow_oci_cache.sif lolcow-oci.def
WARNING: Authentication token file not found : Only pulls of public images will
↳succeed
Build target already exists. Do you want to overwrite? [N/y] y
INFO: Starting build...
Getting image source signatures
Copying blob sha256:9fb6c798fa41e509b58bcc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
 45.33 MiB / 45.33 MiB [=====] 0s
Copying blob sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
 848 B / 848 B [=====] 0s
Copying blob sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
 621 B / 621 B [=====] 0s
Copying blob sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
 853 B / 853 B [=====] 0s
Copying blob sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
 169 B / 169 B [=====] 0s
Copying blob sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
 53.75 MiB / 53.75 MiB [=====] 0s
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
```

(continues on next page)

(continued from previous page)

```

3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: /home/vagrant/lolcow_oci_cache.sif

```

Required use of `sudo` allows Singularity to build the SIF container `lolcow_oci_cache.sif`.

When it comes to OCI archives, the definition file, `lolcow-ocia.def` corresponding to the command-line invocation above is:

```

Bootstrap: oci-archive
From: godlovedc_lolcow.tar

```

Applying build as follows

```

$ sudo singularity build lolcow_oci_tarfile.sif lolcow-ocia.def
WARNING: Authentication token file not found : Only pulls of public images will
↳succeed
INFO:   Starting build...
Getting image source signatures
Skipping fetch of repeat blob
↳sha256:9fb6c798fa41e509b58bccc5c29654c3ff4648b608f5daa67c1aab6a7d02c118
Skipping fetch of repeat blob
↳sha256:3b61febd4aefe982e0cb9c696d415137384d1a01052b50a85aae46439e15e49a
Skipping fetch of repeat blob
↳sha256:9d99b9777eb02b8943c0e72d7a7baec5c782f8fd976825c9d3fb48b3101aacc2
Skipping fetch of repeat blob
↳sha256:d010c8cf75d7eb5d2504d5ffa0d19696e8d745a457dd8d28ec6dd41d3763617e
Skipping fetch of repeat blob
↳sha256:7fac07fb303e0589b9c23e6f49d5dc1ff9d6f3c8c88cabe768b430bdb47f03a9
Skipping fetch of repeat blob
↳sha256:8e860504ff1ee5dc7953672d128ce1e4aa4d8e3716eb39fe710b849c64b20945
Copying config sha256:73d5b1025fbfa138f2cacf45bbf3f61f7de891559fa25b28ab365c7d9c3cbd82
3.33 KiB / 3.33 KiB [=====] 0s
Writing manifest to image destination
Storing signatures
INFO:   Creating SIF file...
INFO:   Build complete: lolcow_oci_tarfile.sif

```

results in the SIF container `lolcow_oci_tarfile.sif`.

### 9.7.2.5 Working with Definition Files: Additional Considerations

In working with definition files, the following additional considerations arise:

- In addition to the mandatory header keywords documented above, *optional header keywords* are possible additions to OCI bundle and/or archive bootstrap definition files.
- As distribution of OCI bundles and/or archives is out of the Initiative's scope, so is the authentication required to access private images and/or registries.
- The direction of execution follows along the same lines *as described above*. Of course, the SIF container's metadata will make clear the `runscript` through application of the `inspect` command *as described previously*.
- Container metadata will also reveal whether or not a given SIF file was bootstrapped from an OCI bundle or archive; for example, below it is evident that an OCI archive was employed to bootstrap creation of the SIF file:

```
$ singularity inspect --labels lolcow_oci_tarfile.sif | jq
{
  "org.label-schema.build-date": "Sunday_27_January_2019_0:5:29_UTC",
  "org.label-schema.schema-version": "1.0",
  "org.label-schema.usage.singularity.deffile.bootstrap": "oci-archive",
  "org.label-schema.usage.singularity.deffile.from": "godlovedc_lolcow.tar",
  "org.label-schema.usage.singularity.version": "3.0.3-1"
}
```

## 9.8 Best Practices

Singularity can make use of most Docker and OCI images without complication. However, there exist known cases where complications can arise. Thus a brief compilation of best practices follows below.

### 1. Accounting for trust

Docker containers *allow for* privilege escalation. In a `Dockerfile`, for example, the `USER` instruction allows for user and/or group settings to be made in the Linux operating environment. The trust model in Singularity is completely different: Singularity allows untrusted users to run untrusted containers in a trusted way. Because Singularity containers embodied as SIF files execute in *user* space, there is no possibility for privilege escalation. In other words, those familiar with Docker, should *not* expect access to elevated user permissions; and as a corollary, use of the `USER` instruction must be *avoided*.

Singularity does, however, allow for fine-grained control over the permissions that containers require for execution. Given that Singularity executes in user space, it is not surprising that permissions need to be externally established *for* the container through use of the `capability` command. *Detailed elsewhere in this documentation*, Singularity allows users and/or groups to be granted/revoked authorized capabilities. Owing to Singularity’s trust model, this fundamental best practice can be stated as follows:

“Employ `singularity capability` to manage execution privileges for containers”

### 2. Maintaining containers built from Docker and OCI images

SIF files created by bootstrapping from Docker or OCI images are, of course, only as current as the most recent Singularity `pull`. Subsequent retrievals *may* result in containers that are built and/or behave differently, owing to changes in the corresponding `Dockerfile`. A prudent practice then, for maintaining containers of value, is based upon use of Singularity definition files. Styled and implemented after a `Dockerfile` retrieved at some point in time, use of `diff` on subsequent versions of this same file, can be employed to inform maintenance of the corresponding Singularity definition file. Understanding build specifications at this level of detail places container creators in a much more sensible position prior to signing with an encrypted key. Thus the best practice is:

“Maintain detailed build specifications for containers, rather than opaque runtimes”

### 3. Working with environment variables

In a `Dockerfile`, **environment variables are declared** as key-value pairs through use of the `ENV` instruction. Declaration in the build specification for a container is advised, rather than relying upon user (e.g., `.bashrc`, `.profile`) or system-wide configuration files for interactive shells. Should a `Dockerfile` be converted into a definition file for Singularity, as suggested in the container-maintenance best practice above, *environment variables can be explicitly represented* as `ENV` instructions that have been converted into entries in the `%environment` section, respectively. This best practice can be stated as follows:

“Define environment variables in container specifications, not interactive shells”

### 4. Installation to `/root`

Docker and OCI container's are typically run as the `root` user; therefore, `/root` (this user's `$HOME` directory) will be the installation target when `$HOME` is specified. Installation to `/root` may prove workable in some circumstances - e.g., while the container is executing, or if read-only access is required to this directory after installation. In general, however, because this is the `root` directory conventional wisdom suggests this practice be avoided. Thus the best practice is:

“Avoid installations that make use of `/root`.”

### 5. Read-only / filesystem

Singularity mounts a container's `/` filesystem in read-only mode. To ensure a Docker container meets Singularity's requirements, it may prove useful to execute `docker run --read-only --tmpfs /run --tmpfs /tmp godlovedc/lolcow`. The best practice here is:

“Ensure Docker containers meet Singularity's read-only / filesystem requirement”

### 6. Installation to `$HOME` or `$TMP`

In making use of Singularity, it is common practice for `$USER` to be automatically mounted on `$HOME`, and for `$TMP` also to be mounted. To avoid the side effects (e.g., ‘missing’ or conflicting files) that might arise as a consequence of executing `mount` commands then, the best practice is:

“Avoid placing container ‘valuables’ in `$HOME` or `$TMP`.”

A detailed review of the container's build specification (e.g., its `Dockerfile`) may be required to ensure this best practice is adhered to.

### 7. Current library caches

Irrespective of containers, a **common runtime error** stems from failing to locate shared libraries required for execution. Suppose now there exists a requirement for symbolically linked libraries *within* a Singularity container. If the build process that creates the container fails to update the cache, then it is quite likely that (read-only) execution of this container will result in the common error of missing libraries. Upon investigation, it is likely revealed that the library exists, just not the required symbolic links. Thus the best practice is:

“Ensure calls to `ldconfig` are executed towards the *end* of build specifications (e.g., `Dockerfile`), so that the library cache is updated when the container is created.”

### 8. Use of plain-text passwords for authentication

For obvious reasons, it is desirable to completely *avoid* use of plain-text passwords. Therefore, for interactive sessions requiring authentication, use of the `--docker-login` option for Singularity's `pull` and `build` commands is *recommended*. At the present time, the *only* option available for non-interactive use is to *embed plain-text passwords into environment variables*. Because the Sylabs Cloud Singularity Library employs **time-limited API tokens for authentication**, use of SIF containers hosted through this service provides a more secure means for both interactive *and* non-interactive use. This best practice is:

“Avoid use of plain-text passwords”

### 9. Execution ambiguity

Short of converting an *entire* `Dockerfile` into a Singularity definition file, informed specification of the `%runscript` entry in the `def` file *removes* any ambiguity associated with `ENTRYPOINT` *versus* `CMD` and ultimately *execution precedence*. Thus the best practice is:

“Employ Singularity's `%runscript` by default to avoid execution ambiguity”

Note that the `ENTRYPOINT` can be bypassed completely, e.g., `docker run -i -t --entrypoint /bin/bash godlovedc/lolcow`. This allows for an interactive session within the container, that may prove useful in validating the built runtime.

Best practices emerge from experience. Contributions that allow additional experiences to be shared as best practices are always encouraged. Please refer to *Contributing* for additional details.

## 9.9 Troubleshooting

In making use of Docker and OCI images through Singularity the need to troubleshoot may arise. A brief compilation of issues and their resolution is provided here.

### 1. Authentication issues

Authentication is required to make use of Docker-style private images and/or private registries. Examples involving private images hosted by the public Docker Hub were *provided above*, whereas the NVIDIA GPU Cloud was used to *illustrate access to a private registry*. Even if the intended use of containers is non-interactive, issues in authenticating with these image-hosting services are most easily addressed through use of the `--docker-login` option that can be appended to a Singularity `pull` request. As soon as image signatures and blobs start being received, authentication credentials have been validated, and the image `pull` can be cancelled.

### 2. Execution mismatches

Execution intentions are detailed through specification files - i.e., the `Dockerfile` in the case of Docker images. However, intentions and precedence aside, the reality of executing a container may not align with expectations. To alleviate this mismatch, use of `singularity inspect --runscript <somecontainer>.sif` details the *effective* runscript - i.e., the one that is actually being executed. Of course, the ultimate solution to this issue is to develop and maintain Singularity definition files for containers of interest.

### 3. More than one image in the OCI bundle directory

*As illustrated above*, and with respect to the bootstrap agent `oci://$OCI_BUNDLE_DIR`, a fatal error is generated when *more* than one image is referenced in the `$OCI_BUNDLE_DIR/index.json` file. The workaround shared previously was to append the bootstrap directive with the unique reference name for the image of interest - i.e., `oci://$OCI_BUNDLE_DIR:org.opencontainers.image.ref.name`. Because it may take some effort to locate the reference name for an image of interest, an even simpler solution is to ensure that each `$OCI_BUNDLE_DIR` contains at most a single image.

### 4. Cache maintenance

Maintenance of the Singularity cache (i.e., `$HOME/.singularity/cache`) requires manual intervention at this time. By **carefully** issuing the command `rm -rf $HOME/.singularity/cache`, its local cache will be cleared of all downloaded images.

### 5. The `http` and `https` are `pull` only bootstrap agents

`http` and `https` are the only examples of `pull` only bootstrap agents. In other words, when used with Singularity's `pull` command, the result is a local copy of, for example, an OCI archive image. This means that a subsequent step is necessary to actually create a SIF container for use by Singularity - a step involving the `oci-archive` bootstrap agent in the case of an OCI image archive.

Like *best practices*, troubleshooting scenarios and solutions emerge from experience. Contributions that allow additional experiences to be shared are always encouraged. Please refer to *Contributing* for additional details.





## BIND PATHS AND MOUNTS

If [enabled by the system administrator](#), Singularity allows you to map directories on your host system to directories within your container using bind mounts. This allows you to read and write data on the host system with ease.

### 10.1 Overview

When Singularity ‘swaps’ the host operating system for the one inside your container, the host file systems becomes inaccessible. But you may want to read and write files on the host system from within the container. To enable this functionality, Singularity will bind directories back into the container via two primary methods: system-defined bind paths and user-defined bind paths.

### 10.2 System-defined bind paths

The system administrator has the ability to define what bind paths will be included automatically inside each container. Some bind paths are automatically derived (e.g. a user’s home directory) and some are statically defined (e.g. bind paths in the Singularity configuration file). In the default configuration, the system default bind points are `$HOME`, `/sys:/sys`, `/proc:/proc`, `/tmp:/tmp`, `/var/tmp:/var/tmp`, `/etc/resolv.conf:/etc/resolv.conf`, `/etc/passwd:/etc/passwd`, and `$PWD`. Where the first path before `:` is the path from the host and the second path is the path in the container.

### 10.3 User-defined bind paths

If the system administrator has [enabled user control of binds](#), you will be able to request your own bind paths within your container.

The Singularity action commands (`run`, `exec`, `shell`, and `instance start`) will accept the `--bind/-B` command-line option to specify bind paths, and will also honor the `$$SINGULARITY_BIND` (or `$$SINGULARITY_BINDPATH`) environment variable. The argument for this option is a comma-delimited string of bind path specifications in the format `src[:dest[:opts]]`, where `src` and `dest` are paths outside and inside of the container respectively. If `dest` is not given, it is set equal to `src`. Mount options (`opts`) may be specified as `ro` (read-only) or `rw` (read/write, which is the default). The `--bind/-B` option can be specified multiple times, or a comma-delimited string of bind path specifications can be used.

#### 10.3.1 Specifying bind paths

Here’s an example of using the `--bind` option and binding `/data` on the host to `/mnt` in the container (`/mnt` does not need to already exist in the container):

```
$ ls /data
bar  foo

$ singularity exec --bind /data:/mnt my_container.sif ls /mnt
bar  foo
```

You can bind multiple directories in a single command with this syntax:

```
$ singularity shell --bind /opt,/data:/mnt my_container.sif
```

This will bind `/opt` on the host to `/opt` in the container and `/data` on the host to `/mnt` in the container.

Using the environment variable instead of the command line argument, this would be:

```
$ export SINGULARITY_BIND="/opt,/data:/mnt"
$ singularity shell my_container.sif
```

Using the environment variable `$SINGULARITY_BIND`, you can bind paths even when you are running your container as an executable file with a runscript. If you bind many directories into your Singularity containers and they don't change, you could even benefit by setting this variable in your `.bashrc` file.

### 10.3.2 A note on using `--bind` with the `--writable` flag

To mount a bind path inside the container, a *bind point* must be defined within the container. The bind point is a directory within the container that Singularity can use as a destination to bind a directory on the host system.

Starting in version 3.0, Singularity will do its best to bind mount requested paths into a container regardless of whether the appropriate bind point exists within the container. Singularity can often carry out this operation even in the absence of the “overlay fs” feature.

However, binding paths to non-existent points within the container can result in unexpected behavior when used in conjunction with the `--writable` flag, and is therefore disallowed. If you need to specify bind paths in combination with the `--writable` flag, please ensure that the appropriate bind points exist within the container. If they do not already exist, it will be necessary to modify the container and create them.

## PERSISTENT OVERLAYS

Persistent overlay directories allow you to overlay a writable file system on an immutable read-only container for the illusion of read-write access.

### 11.1 Overview

A persistent overlay is a directory that “sits on top” of your compressed, immutable SIF container. When you install new software or create and modify files the overlay directory stores the changes.

If you want to use a SIF container as though it were writable, you can create a directory to use as a persistent overlay. Then you can specify that you want to use the directory as an overlay at runtime with the `--overlay` option.

You can use a persistent overlays with the following commands:

- `run`
- `exec`
- `shell`
- `instance.start`

### 11.2 Usage

To use a persistent overlay, you must first have a container.

```
$ sudo singularity build ubuntu.sif library://ubuntu
```

Then you must create a directory. (You can also use the `--overlay` option with a legacy writable ext3 image.)

```
$ mkdir my_overlay
```

Now you can use this overlay directory with your container. Note that it is necessary to be root to use an overlay directory.

```
$ sudo singularity shell --overlay my_overlay/ ubuntu.sif
Singularity ubuntu.sif:~> touch /foo
Singularity ubuntu.sif:~> apt-get update && apt-get install -y vim
Singularity ubuntu.sif:~> which vim
```

(continues on next page)

(continued from previous page)

```
/usr/bin/vim
Singularity ubuntu.sif:~> exit
```

You will find that your changes persist across sessions as though you were using a writable container.

```
$ sudo singularity shell --overlay my_overlay/ ubuntu.sif
Singularity ubuntu.sif:~> ls /foo
/foo
Singularity ubuntu.sif:~> which vim
/usr/bin/vim
Singularity ubuntu.sif:~> exit
```

If you mount your container without the `--overlay` directory, your changes will be gone.

```
$ sudo singularity shell ubuntu.sif
Singularity ubuntu.sif:~> ls /foo
ls: cannot access 'foo': No such file or directory
Singularity ubuntu.sif:~> which vim
Singularity ubuntu.sif:~> exit
```

## RUNNING SERVICES

There are *different ways* in which you can run Singularity containers. If you use commands like `run`, `exec` and `shell` to interact with processes in the container, you are running Singularity containers in the foreground. Singularity, also lets you run containers in a “detached” or “daemon” mode which can run different services in the background. A “service” is essentially a process running in the background that multiple different clients can use. For example, a web server or a database. To run services in a Singularity container one should use *instances*. A container instance is a persistent and isolated version of the container image that runs in the background.

### 12.1 Overview

Singularity v2.4 introduced the concept of *instances* allowing users to run services in Singularity. This page will help you understand instances using an elementary example followed by a more useful example running an NGINX web server using instances. In the end, you will find a more detailed example of running an instance of an API that converts URL to PDFs.

To begin with, suppose you want to run an NGINX web server outside of a container. On Ubuntu, you can simply install NGINX and start the service by:

```
$ sudo apt-get update && sudo apt-get install -y nginx
$ sudo service nginx start
```

If you were to do something like this from within a container you would also see the service start, and the web server running. But then if you were to exit the container, the process would continue to run within an unreachable mount namespace. The process would still be running, but you couldn’t easily kill or interface with it. This is called an orphan process. Singularity instances give you the ability to handle services properly.

### 12.2 Container Instances in Singularity

For demonstration, let’s use an easy (though somewhat useless) example of `alpine_latest.sif` image from the [container library](#):

```
$ singularity pull library://alpine
```

The above command will save the alpine image from the Container Library as `alpine_latest.sif`.

To start an instance, you should follow this procedure :

```
[command]           [image]           [name of instance]
$ singularity instance start  alpine_latest.sif  instance1
```

This command causes Singularity to create an isolated environment for the container services to live inside. One can confirm that an instance is running by using the `instance list` command like so:

```
$ singularity instance list

INSTANCE NAME   PID      IMAGE
instance1      12715    /home/ysub/alpine_latest.sif
```

**Note:** The instances are linked with your user. So make sure to run *all* the instance commands either with or without the `sudo` privilege. If you start an instance with `sudo` and then you must `list` it with `sudo`, as well or you will not be able to locate the instance.

If you want to run multiple instances from the same image, it's as simple as running the command multiple times with different instance names. The instance name uniquely identify instances, so they cannot be repeated.

```
$ singularity instance start alpine_latest.sif instance2
$ singularity instance start alpine_latest.sif instance3
```

And again to confirm that the instances are running as we expected:

```
$ singularity instance list

INSTANCE NAME   PID      IMAGE
instance1      12715    /home/ysub/alpine_latest.sif
instance2      12795    /home/ysub/alpine_latest.sif
instance3      12837    /home/ysub/alpine_latest.sif
```

You can use the `singularity run/exec` commands on instances:

```
$ singularity run instance://instance1
$ singularity exec instance://instance2 cat /etc/os-release
```

When using `run` with an instance URI, the `runscript` will be executed inside of the instance. Similarly with `exec`, it will execute the given command in the instance.

If you want to poke around inside of your instance, you can do a normal `singularity shell` command, but give it the instance URI:

```
$ singularity shell instance://instance3
Singularity>
```

When you are finished with your instance you can clean it up with the `instance stop` command as follows:

```
$ singularity instance stop instance1
```

If you have multiple instances running and you want to stop all of them, you can do so with a wildcard or the `--all` flag. The following three commands are all identical.

```
$ singularity instance stop \*
$ singularity instance stop --all
$ singularity instance stop --a
```

---

**Note:** Note that you must escape the wildcard with a backslash like this `\*` to pass it properly.

---

## 12.3 Nginx “Hello-world” in Singularity

The above example, although not very useful, should serve as a fair introduction to the concept of Singularity instances and running services in the background. The following illustrates a more useful example of setting up a sample NGINX web server using instances. First we will create a basic *definition file* (let’s call it `nginx.def`):

```
Bootstrap: docker
From: nginx
Includecmd: no

%startscript
  nginx
```

This downloads the official NGINX Docker container, converts it to a Singularity image, and tells it to run NGINX when you start the instance. Since we’re running a web server, we’re going to run the following commands as root.

```
$ sudo singularity build nginx.sif nginx.def

$ sudo singularity instance start --writable-tmpfs nginx.sif web
```

---

**Note:** The above `start` command requires `sudo` because we are running a web server. Also, to let the instance write temporary files during execution, you should use `--writable-tmpfs` while starting the instance.

---

Just like that we’ve downloaded, built, and run an NGINX Singularity image. And to confirm that it’s correctly running:

```
$ curl localhost

<!DOCTYPE html>
<html>
<head>
<title>Welcome to nginx!</title>
<style>
  body {
    width: 35em;
    margin: 0 auto;
    font-family: Tahoma, Verdana, Arial, sans-serif;
  }
</style>
</head>
<body>
<h1>Welcome to nginx!</h1>
<p>If you see this page, the nginx web server is successfully installed and
working. Further configuration is required.</p>

<p>For online documentation and support please refer to
<a href="http://nginx.org/">nginx.org</a>.<br/>
Commercial support is available at
<a href="http://nginx.com/">nginx.com</a>.</p>
```

(continues on next page)

(continued from previous page)

```
<p><em>Thank you for using nginx.</em></p>
</body>
</html>
```

Visit localhost on your browser, you should see a Welcome message!

## 12.4 Putting all together

In this section, we will demonstrate an example of packaging a service into a container and running it. The service we will be packaging is an API server that converts a web page into a PDF, and can be found [here](#). You can build the image by following the steps described below or you can just download the final image directly from Container Library, simply run:

```
$ singularity pull url-to-pdf.sif library://sylabs/doc-examples/url-to-pdf:latest
```

### 12.4.1 Building the image

This section will describe the requirements for creating the definition file (`url-to-pdf.def`) that will be used to build the container image. `url-to-pdf-api` is based on a Node 8 server that uses a headless version of Chromium called [Puppeteer](#). Let's first choose a base from which to build our container, in this case the docker image `node:8` which comes pre-installed with Node 8 has been used:

```
Bootstrap: docker
From: node:8
Includecmd: no
```

Puppeteer also requires a slew of dependencies to be manually installed in addition to Node 8, so we can add those into the `post` section as well as the installation script for the `url-to-pdf`:

```
%post

apt-get update && apt-get install -yq gconf-service libasound2 \
libatk1.0-0 libc6 libcairo2 libcups2 libdbus-1-3 libexpat1 \
libfontconfig1 libgcc1 libgconf-2-4 libgdk-pixbuf2.0-0 \
libglib2.0-0 libgtk-3-0 libnspr4 libpango-1.0-0 \
libpangocairo-1.0-0 libstdc++6 libx11-6 libx11-xcb1 libxcb1 \
libxcomposite1 libxcursor1 libxdamage1 libxext6 libxfixed3 libxi6 \
libxrandr2 libxrender1 libxss1 libxtst6 ca-certificates \
fonts-liberation libappindicator1 libnss3 lsb-release xdg-utils \
wget curl && rm -r /var/lib/apt/lists/*
git clone https://github.com/alvarcarto/url-to-pdf-api.git pdf_server
cd pdf_server
npm install
chmod -R 0755 .
```

And now we need to define what happens when we start an instance of the container. In this situation, we want to run the commands that starts up the `url-to-pdf` service:

```
%startscript
cd /pdf_server
```

(continues on next page)



(continued from previous page)

```
# Use nohup and /dev/null to completely detach server process from terminal
nohup npm start > /dev/null 2>&1 < /dev/null &
```

Also, the `url-to-pdf` service requires some environment variables to be set, which we can do in the environment section:

```
%environment
NODE_ENV=development
PORT=9000
ALLOW_HTTP=true
URL=localhost
export NODE_ENV PORT ALLOW_HTTP URL
```

The complete definition file will look like this:

```
Bootstrap: docker
From: node:8
Includecmd: no

%post

apt-get update && apt-get install -yq gconf-service libasound2 \
  libatk1.0-0 libc6 libcairo2 libcups2 libdbus-1-3 libexpat1 \
  libfontconfig1 libgcc1 libgconf-2-4 libgdk-pixbuf2.0-0 \
  libgl1.0-0 libgtk-3-0 libnspr4 libpango-1.0-0 \
  libpangocairo-1.0-0 libstdc++6 libx11-6 libx11-xcb1 libxcb1 \
  libxcomposite1 libxcursor1 libxdamage1 libxext6 libxfixed3 libxi6 \
  libxrandr2 libxrender1 libxss1 libxtst6 ca-certificates \
  fonts-liberation libappindicator1 libnss3 lsb-release xdg-utils \
  wget curl && rm -r /var/lib/apt/lists/*
git clone https://github.com/alvarcarto/url-to-pdf-api.git pdf_server
cd pdf_server
npm install
chmod -R 0755 .

%startscript
cd /pdf_server
# Use nohup and /dev/null to completely detach server process from terminal
nohup npm start > /dev/null 2>&1 < /dev/null &

%environment
NODE_ENV=development
PORT=9000
ALLOW_HTTP=true
URL=localhost
export NODE_ENV PORT ALLOW_HTTP URL
```

The container can be built like so:

```
$ sudo singularity build url-to-pdf.sif url-to-pdf.def
```

## 12.4.2 Running the Service

We can now start an instance and run the service:

```
$ sudo singularity instance start url-to-pdf.sif pdf
```



**Note:** If there occurs an error related to port connection being refused while starting the instance or while using it later, you can try specifying different port numbers in the `%environment` section of the definition file above.

We can confirm it's working by sending the server an http request using curl:

```
$ curl -o sylabs.pdf localhost:9000/api/render?url=http://sylabs.io/docs
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total     Spent    Left     Speed
100 73750  100 73750    0     0  14583      0  0:00:05  0:00:05  --:--:-- 19130
```

You should see a PDF file being generated like the one shown below:

**sylabs.pdf**

**3.0**

Documentation for user: [HTML](#) | [PDF](#)

Documentation for admins: [HTML](#) | [PDF](#)

**2.6**

Documentation for users: [HTML](#) | [PDF](#)

Documentation for admins: [HTML](#) | [PDF](#)

**2.5**

Documentation for users: [HTML](#) | [PDF](#)

Documentation for admins: [HTML](#) | [PDF](#)

---

**About Sylabs**

Singularity was born out of the need to properly containerize and support workflows related to artificial intelligence, machine/deep learning.

---

**Products and Services**

Singularity

If you shell into the instance, you can see the running processes:

```
$ sudo singularity shell instance://pdf
Singularity: Invoking an interactive shell within container...

Singularity final.sif:/home/ysub> ps auxf
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root      461  0.0  0.0  18204  3188 pts/1    S   17:58   0:00 /bin/bash --norc
root      468  0.0  0.0  36640  2880 pts/1    R+  17:59   0:00  \_ ps auxf
root         1  0.0  0.1 565392 12144 ?        S1   15:10   0:00 sinit
root       16  0.0  0.4 1113904 39492 ?        S1   15:10   0:00 npm
root       26  0.0  0.0   4296   752 ?        S   15:10   0:00  \_ sh -c nodemon --
↳watch ./src -e js src/index.js
root       27  0.0  0.5 1179476 40312 ?        S1   15:10   0:00  \_ node /pdf_
↳server/node_modules/.bin/nodemon --watch ./src -e js src/index.js
root       39  0.0  0.7 936444 61220 ?        S1   15:10   0:02  \_ /usr/
↳local/bin/node src/index.js

Singularity final.sif:/home/ysub> exit
```

### 12.4.3 Making it Fancy

Now that we have confirmation that the server is working, let's make it a little cleaner. It's difficult to remember the exact `curl` command and URL syntax each time you want to request a PDF, so let's automate it. To do that, we can use Scientific Filesystem (SCIF) apps, that are integrated directly into singularity. If you haven't already, check out the [Scientific Filesystem documentation](#) to come up to speed.

First off, we're going to move the installation of the `url-to-pdf` into an app, so that there is a designated spot to place output files. To do that, we want to add a section to our definition file to build the server:

```
%appinstall pdf_server
  git clone https://github.com/alvarcarto/url-to-pdf-api.git pdf_server
  cd pdf_server
  npm install
  chmod -R 0755 .
```

And update our startscript to point to the app location:

```
%startscript
  cd /scif/apps/pdf_server/scif/pdf_server
  # Use nohup and /dev/null to completely detach server process from terminal
  nohup npm start > /dev/null 2>&1 < /dev/null &
```

Now we want to define the `pdf_client` app, which we will run to send the requests to the server:

```
%apprun pdf_client
  if [ -z "${1:-}" ]; then
    echo "Usage: singularity run --app pdf <instance://name> <URL> [output file]"
    exit 1
  fi
  curl -o "${SINGULARITY_APPDATA}/output/${2:-output.pdf}" "${URL}:${PORT}/api/
↳render?url=${1}"
```

As you can see, the `pdf_client` app checks to make sure that the user provides at least one argument.

The full def file will look like this:

```

Bootstrap: docker
From: node:8
Includecmd: no

%post

apt-get update && apt-get install -yq gconf-service libasound2 \
  libatk1.0-0 libc6 libcairo2 libcups2 libdbus-1-3 libexpat1 \
  libfontconfig1 libgcc1 libgconf-2-4 libgdk-pixbuf2.0-0 \
  libglib2.0-0 libgtk-3-0 libnspr4 libpango-1.0-0 \
  libpangocairo-1.0-0 libstdc++6 libx11-6 libx11-xcb1 libxcb1 \
  libxcompositel libxcursor1 libxdamage1 libxext6 libxfixes3 libxi6 \
  libxrandr2 libxrender1 libxss1 libxtst6 ca-certificates \
  fonts-liberation libappindicator1 libnss3 lsb-release xdg-utils \
  wget curl && rm -r /var/lib/apt/lists/*

%appinstall pdf_server
git clone https://github.com/alvarcarto/url-to-pdf-api.git pdf_server
cd pdf_server
npm install
chmod -R 0755 .

%startscript
cd /scif/apps/pdf_server/scif/pdf_server
# Use nohup and /dev/null to completely detach server process from terminal
nohup npm start > /dev/null 2>&1 < /dev/null &

%environment
NODE_ENV=development
PORT=9000
ALLOW_HTTP=true
URL=localhost
export NODE_ENV PORT ALLOW_HTTP URL

%apprun pdf_client
if [ -z "${1:-}" ]; then
  echo "Usage: singularity run --app pdf <instance://name> <URL> [output file]"
  exit 1
fi
curl -o "${SINGULARITY_APPDATA}/output/${2:-output.pdf}" "${URL}:${PORT}/api/
↪render?url=${1}"

```

Create the container as before. The `--force` option will overwrite the old container:

```
$ sudo singularity build --force url-to-pdf.sif url-to-pdf.def
```

Now that we have an output directory in the container, we need to expose it to the host using a bind mount. Once we've rebuilt the container, make a new directory called `/tmp/out` for the generated PDFs to go.

```
$ mkdir /tmp/out
```

After building the image from the edited definition file we simply start the instance:

```
$ singularity instance start --bind /tmp/out/:/output url-to-pdf.sif pdf
```

To request a pdf simply do:

```
$ singularity run --app pdf_client instance://pdf http://sylabs.io/docs sylabs.pdf
```

To confirm that it worked:

```
$ ls /tmp/out/  
sylabs.pdf
```

When you are finished, use the instance stop command to close all running instances.

```
$ singularity instance stop --all
```

---

**Note:** If the service you want to run in your instance requires a bind mount, then you must pass the `--bind` option when calling `instance start`. For example, if you wish to capture the output of the `web` container instance which is placed at `/output/` inside the container you could do:

```
$ singularity instance start --bind output/dir/outside:/output/ nginx.sif web
```

---

## ENVIRONMENT AND METADATA

Singularity containers support environment variables and labels that you can add to your container during the build process. If you are looking for environment variables to set up the environment on the host system during build time, see the *build environment section*.

### 13.1 Overview

Environment variables can be included in your container by adding them in your definition file:

- In the `%environment` section of your definition file.

```
Bootstrap: library
From: default/alpine

%environment
    VARIABLE_ONE = hello
    VARIABLE_TWO = world
    export VARIABLE_ONE VARIABLE_TWO
```

- Or in the `%post` section of your definition file.

```
Bootstrap: library
From: default/alpine

%post
    echo 'export VARIABLE_NAME=variable_value' >>$$SINGULARITY_ENVIRONMENT
```

You can also add labels to your container using the `%labels` section like so:

```
Bootstrap: library
From: default/alpine

%labels
    OWNER = Joana
```

To view the labels within your container you use the `inspect` command:

```
$ singularity inspect mysifimage.sif
```

This will give you the following output:

```
{
  "OWNER": "Joana",
  "org.label-schema.build-date": "Monday_07_January_2019_0:01:50_CET",
  "org.label-schema.schema-version": "1.0",
  "org.label-schema.usage": "/.singularity.d/runscript.help",
  "org.label-schema.usage.singularity.deffile.bootstrap": "library",
  "org.label-schema.usage.singularity.deffile.from": "debian:9",
  "org.label-schema.usage.singularity.runscript.help": "/.singularity.d/runscript.
↪help",
  "org.label-schema.usage.singularity.version": "3.0.1-236.g2453fdfe"
}
```

Many of these labels are created by default, but you can also see the custom label that was added in the example above.

The `inspect` command has *additional options* that are useful for viewing the container's metadata.

## 13.2 Environment

If you build a container from the [Container Library](#) or [Docker Hub](#), the environment will be included with the container at build time. You can also define new environment variables in your definition file as follows:

```
Bootstrap: library
From: default/alpine

%environment
  #First define the variables
  VARIABLE_PATH=/usr/local/bootstrap
  VARIABLE_VERSION=3.0
  #Then export them
  export VARIABLE_PATH VARIABLE_VERSION
```

You may need to add environment variables to your container during the `%post` section. For instance, maybe you will not know the appropriate value of a variable until you have installed some software. To add variables to the environment during `%post` you can use the `$$SINGULARITY_ENVIRONMENT` variable with the following syntax:

```
%post
  echo 'export VARIABLE_NAME=variable_value' >>$$SINGULARITY_ENVIRONMENT
```

Text in the `%environment` section will be appended to the file `/.singularity.d/env/90-environment.sh` while text redirected to `$$SINGULARITY_ENVIRONMENT` will appear in the file `/.singularity.d/env/91-environment.sh`. If nothing is redirected to `$$SINGULARITY_ENVIRONMENT` in the `%post` section, the file `/.singularity.d/env/91-environment.sh` will not exist.

Because files in `/.singularity.d/env` are sourced in alpha-numerical order, variables added using `$$SINGULARITY_ENVIRONMENT` take precedence over those added via the `%environment` section.

If you need to define a variable in the container at runtime, when you execute Singularity pass a variable prefixed with `SINGULARITYENV_`. These variables will be transposed automatically and the prefix will be stripped. For example, let's say we want to set the variable `HELLO` to have value `world`. We can do that as follows:

```
$ SINGULARITYENV_HELLO=world singularity exec centos7.img env | grep HELLO
HELLO=world
```

The `--cleanenv` option can be used to remove the host environment and execute a container with a minimal environment.



```
$ singularity exec --cleanenv centos7.img env
LD_LIBRARY_PATH=/usr/local/lib:/usr/local/lib64
SINGULARITY_NAME=test.img
PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin
PWD=/home/gmk/git/singularity
LANG=en_US.UTF-8
SHLVL=0
SINGULARITY_INIT=1
SINGULARITY_CONTAINER=test.img
```

Without the `--cleanenv` flag, the environment on the host system will be present within the container at run time.

If you need to change the `$PATH` of your container at run time there are a few special environmental variables you can use:

- `SINGULARITYENV_PREPEND_PATH=/good/stuff/at/beginning` to prepend directories to the beginning of the `$PATH`
- `SINGULARITYENV_APPEND_PATH=/good/stuff/at/end` to append directories to the end of the `$PATH`
- `SINGULARITYENV_PATH=/a/new/path` to override the `$PATH` within the container

## 13.3 Labels

Your container stores metadata about its build, along with Docker labels, and custom labels that you define during build in a `%labels` section.

For containers that are generated with Singularity version 3.0 and later, labels are represented using the [rc1 Label Schema](#). For example:

```
$ singularity inspect jupyter.sif
{
  "OWNER": "Joana",
  "org.label-schema.build-date": "Friday_21_December_2018_0:49:50_CET",
  "org.label-schema.schema-version": "1.0",
  "org.label-schema.usage": "/.singularity.d/runscript.help",
  "org.label-schema.usage.singularity.deffile.bootstrap": "library",
  "org.label-schema.usage.singularity.deffile.from": "debian:9",
  "org.label-schema.usage.singularity.runscript.help": "/.singularity.d/
↪runscript.help",
  "org.label-schema.usage.singularity.version": "3.0.1-236.g2453fdfe"
}
```

You will notice that the one label doesn't belong to the label schema, `OWNER`. This was a user provided label during bootstrap.

You can add custom labels to your container in a bootstrap file:

```
Bootstrap: docker
From: ubuntu: latest

%labels
  OWNER Joana
```

The `inspect` command is useful for viewing labels and other container meta-data. The next section will detail its various options.

## 13.4 The inspect command

The `inspect` command gives you the ability to print out the labels and/or other metadata that was added to your container using the definition file.

### 13.4.1 --labels

This flag corresponds to the default behavior of the `inspect` command. When you run a `singularity inspect <your-container.sif>` you will get output like this.

```
$ singularity inspect --labels jupyter.sif
{
  "org.label-schema.build-date": "Friday_21_December_2018_0:49:50_CET",
  "org.label-schema.schema-version": "1.0",
  "org.label-schema.usage": "/.singularity.d/runscript.help",
  "org.label-schema.usage.singularity.deffile.bootstrap": "library",
  "org.label-schema.usage.singularity.deffile.from": "debian:9",
  "org.label-schema.usage.singularity.runscript.help": "/.singularity.d/runscript.
↪help",
  "org.label-schema.usage.singularity.version": "3.0.1-236.g2453fdfe"
}
```

This is the same as running `singularity inspect jupyter.sif`.

### 13.4.2 --deffile

This flag gives you the def file(s) that was used to create the container.

```
$ singularity inspect --deffile jupyter.sif
```

And the output would look like:

```
Bootstrap: library
From: debian:9

%help
  Container with Anaconda 2 (Conda 4.5.11 Canary) and Jupyter Notebook 5.6.0 for_
↪Debian 9.x (Stretch).
  This installation is based on Python 2.7.15

%environment
  JUP_PORT=8888
  JUP_IPNAME=localhost
  export JUP_PORT JUP_IPNAME

%startscript
  PORT=""
  if [ -n "$JUP_PORT" ]; then
  PORT="--port=${JUP_PORT}"
  fi

  IPNAME=""
  if [ -n "$JUP_IPNAME" ]; then
```

(continues on next page)

(continued from previous page)

```

IPNAME="--ip=${JUP_IPNAME}"
fi

exec jupyter notebook --allow-root ${PORT} ${IPNAME}

%setup
#Create the .condarc file where the environments/channels from conda are
↳specified, these are pulled with preference to root
cd /
touch .condarc

%post
echo 'export RANDOM=123456' >>${SINGULARITY_ENVIRONMENT}
#Installing all dependencies
apt-get update && apt-get -y upgrade
apt-get -y install \
build-essential \
wget \
bzip2 \
ca-certificates \
libglib2.0-0 \
libxext6 \
libsm6 \
libxrender1 \
git
rm -rf /var/lib/apt/lists/*
apt-get clean
#Installing Anaconda 2 and Conda 4.5.11
wget -c https://repo.continuum.io/archive/Anaconda2-5.3.0-Linux-x86_64.sh
/bin/bash Anaconda2-5.3.0-Linux-x86_64.sh -bfp /usr/local
#Conda configuration of channels from .condarc file
conda config --file /.condarc --add channels defaults
conda config --file /.condarc --add channels conda-forge
conda update conda
#List installed environments
conda list

```

Which is a definition file for a `jupyter.sif` container.

### 13.4.3 --runscript

This flag shows the runscript for the image.

```
$ singularity inspect --runscript jupyter.sif
```

And the output would look like:

```

#!/bin/sh
OCI_ENTRYPOINT=""
OCI_CMD="bash"
# ENTRYPOINT only - run entrypoint plus args
if [ -z "$OCI_CMD" ] && [ -n "$OCI_ENTRYPOINT" ]; then
SINGULARITY_OCI_RUN="${OCI_ENTRYPOINT} ${@"
fi

```

(continues on next page)

(continued from previous page)

```

# CMD only - run CMD or override with args
if [ -n "$OCI_CMD" ] && [ -z "$OCI_ENTRYPOINT" ]; then
if [ $# -gt 0 ]; then
    SINGULARITY_OCI_RUN="$@"
else
    SINGULARITY_OCI_RUN="{OCI_CMD}"
fi
fi

# ENTRYPOINT and CMD - run ENTRYPOINT with CMD as default args
# override with user provided args
if [ $# -gt 0 ]; then
    SINGULARITY_OCI_RUN="{OCI_ENTRYPOINT} $@"
else
    SINGULARITY_OCI_RUN="{OCI_ENTRYPOINT} {OCI_CMD}"
fi

exec $SINGULARITY_OCI_RUN

```

### 13.4.4 --test

This flag shows the test script for the image.

```
$ singularity inspect --test jupyter.sif
```

This will output the corresponding %test section from the definition file.

### 13.4.5 --environment

This flag shows the environment settings for the image. The respective environment variables set in %environment section ( So the ones in 90-environment.sh ) and SINGULARITY\_ENV variables set at runtime (that are located in “91-environment.sh”) will be printed out.

```
$ singularity inspect --environment jupyter.sif
```

And the output would look like:

```

==90-environment.sh==
#!/bin/sh

JUP_PORT=8888
JUP_IPNAME=localhost
export JUP_PORT JUP_IPNAME

==91-environment.sh==
export RANDOM=123456

```

As you can see, the JUP\_PORT and JUP\_IPNAME were previously defined in the %environment section of the definition file, while the RANDOM variable shown regards to the use of SINGULARITYENV\_ variables, so in this case SINGULARITYENV\_RANDOM variable was set and exported at runtime.

### 13.4.6 --helpfile

This flag will show the container's description in the %help section of its definition file.

You can call it this way:

```
$ singularity inspect --helpfile jupyter.sif
```

And the output would look like:

```
Container with Anaconda 2 (Conda 4.5.11 Canary) and Jupyter Notebook 5.6.0 for Debian_
↪9.x (Stretch).
This installation is based on Python 2.7.15
```

### 13.4.7 --json

This flag gives you the possibility to output your labels in a JSON format.

You can call it this way:

```
$ singularity inspect --json jupyter.sif
```

And the output would look like:

```
{
  "attributes": {
    "labels": "{\n\t\"org.label-schema.build-date\": \"Friday_21_
↪December_2018_0:49:50_CET\", \n\t\"org.label-schema.schema-version\": \"1.0\", \n\t\
↪\"org.label-schema.usage\": \"/.singularity.d/runscript.help\", \n\t\"org.label-
↪schema.usage.singularity.deffile.bootstrap\": \"library\", \n\t\"org.label-sche-
↪ma.usage.singularity.deffile.from\": \"debian:9\", \n\t\"org.label-schema.usag-
↪e.singularity.runscript.help\": \"/.singularity.d/runscript.help\", \n\t\"org.label-
↪schema.usage.singularity.version\": \"3.0.1-236.g2453fdfe\"\n}"
  },
  "type": "container"
}
```

## 13.5 Container Metadata

Inside of the container, metadata is stored in the /.singularity.d directory. You probably shouldn't edit any of these files directly but it may be helpful to know where they are and what they do:

```
/.singularity.d/
├── actions
│   ├── exec
│   ├── run
│   ├── shell
│   ├── start
│   └── test
├── env
│   ├── 01-base.sh
│   ├── 10-docker2singularity.sh
│   └── 90-environment.sh
```

(continues on next page)

(continued from previous page)

```
| | | 91-environment.sh
| | | 94-appsbase.sh
| | | 95-apps.sh
| | | 99-base.sh
| | |
| | | labels.json
| | | libs
| | | runscript
| | | runscript.help
| | | Singularity
| | | startscript
```

- **actions:** This directory contains helper scripts to allow the container to carry out the action commands. (e.g. `exec`, `run` or `shell`) In later versions of Singularity, these files may be dynamically written at runtime.
- **env:** All `*.sh` files in this directory are sourced in alpha-numeric order when the container is initiated. For legacy purposes there is a symbolic link called `/environment` that points to `./singularity.d/env/90-environment.sh`.
- **labels.json:** The json file that stores a containers labels described above.
- **libs:** At runtime the user may request some host-system libraries to be mapped into the container (with the `--nv` option for example). If so, this is their destination.
- **runscript:** The commands in this file will be executed when the container is invoked with the `run` command or called as an executable. For legacy purposes there is a symbolic link called `/singularity` that points to this file.
- **runscript.help:** Contains the description that was added in the `%help` section.
- **Singularity:** This is the definition file that was used to generate the container. If more than 1 definition file was used to generate the container additional Singularity files will appear in numeric order in a sub-directory called `bootstrap_history`.
- **startscript:** The commands in this file will be executed when the container is invoked with the `instance start` command.

## OCI RUNTIME SUPPORT

### 14.1 Overview

OCI is an acronym for the [Open Containers Initiative](#) - an independent organization whose mandate is to develop open standards relating to containerization. To date, standardization efforts have focused on container formats and runtimes. Singularity's compliance with respect to the OCI Image Specification is considered in detail *elsewhere*. It is Singularity's compliance with the OCI Runtime Specification that is of concern here.

Briefly, compliance with respect to the OCI Runtime Specification is addressed in Singularity through the introduction of the `oci` command group. Although this command group can, in principle, be used to provide a runtime that supports end users, in this initial documentation effort, emphasis is placed upon interoperability with Kubernetes; more specifically, interoperability with Kubernetes via the [Singularity Container Runtime Interface](#).

Owing to this restricted focus, a subset of the Singularity `oci` command group receives attention here; specifically:

- Mounting and unmounting OCI filesystem bundles
- Creating OCI compliant container instances

Some context for integration with Kubernetes via the Singularity CRI is provided at the end of the section.

---

**Note:** All commands in the `oci` command group require `root` privileges.

---

### 14.2 Mounted OCI Filesystem Bundles

#### 14.2.1 Mounting an OCI Filesystem Bundle

`BusyBox` is used here for the purpose of illustration.

Suppose the Singularity Image Format (SIF) file `busybox_latest.sif` exists locally. (Recall:

```
$ singularity pull docker://busybox
INFO:   Starting build...
Getting image source signatures
Copying blob sha256:fc1a6b909f82ce4b72204198d49de3aaf757b3ab2bb823cb6e47c416b97c5985
 738.13 KiB / 738.13 KiB [=====] 0s
Copying config sha256:5fffaf1f2c1830a6a8cf90eb27c7a1a8476b8c49b4b6261a20d6257d031ce4f3
 575 B / 575 B [=====] 0s
Writing manifest to image destination
Storing signatures
```

(continues on next page)

(continued from previous page)

```
INFO:    Creating SIF file...
INFO:    Build complete: busybox_latest.sif
```

This is one way to bootstrap creation of this image in SIF that *retains* a local copy - i.e., a local copy of the SIF file *and* a cached copy of the OCI blobs. Additional approaches and details can be found in the section [Support for Docker and OCI](#).

For the purpose of bootstrapping the creation of an OCI compliant container, this SIF file can be mounted as follows:

```
$ sudo singularity oci mount ./busybox_latest.sif /var/tmp/busybox
```

By issuing the mount command, the root filesystem encapsulated in the SIF file `busybox_latest.sif` is mounted on `/var/tmp/busybox` as an overlay file system,

```
$ sudo df -k
Filesystem                1K-blocks    Used Available Use% Mounted on
udev                      475192         0   475192    0% /dev
tmpfs                     100916     1604    99312    2% /run
/dev/mapper/vagrant--vg-root 19519312 2620740 15883996 15% /
tmpfs                     504560         0   504560    0% /dev/shm
tmpfs                     5120          0     5120    0% /run/lock
tmpfs                     504560         0   504560    0% /sys/fs/cgroup
tmpfs                     100912         0   100912    0% /run/user/900
overlay                   19519312 2620740 15883996 15% /var/tmp/busybox/rootfs
```

with permissions as follows:

```
$ sudo ls -ld /var/tmp/busybox
drwx----- 4 root root 4096 Apr  4 14:30 /var/tmp/busybox
```

## 14.2.2 Content of an OCI Compliant Filesystem Bundle

The *expected* contents of the mounted filesystem are as follows:

```
$ sudo ls -la /var/tmp/busybox
total 28
drwx----- 4 root root 4096 Apr  4 14:30 .
drwxrwxrwt 4 root root 4096 Apr  4 14:30 ..
-rw-rw-rw- 1 root root 9879 Apr  4 14:30 config.json
drwx----- 4 root root 4096 Apr  4 14:30 overlay
drwx----- 1 root root 4096 Apr  4 14:30 rootfs
```

From the perspective of the [OCI runtime specification](#), this content is expected because it prescribes a

“... a format for encoding a container as a **filesystem bundle** - a set of files organized in a certain way, and containing all the necessary data and metadata for any compliant runtime to perform all standard operations against it.”

Critical to compliance with the specification is the presence of the following *mandatory* artifacts residing locally in a single directory:

1. The `config.json` file - a file of configuration data that must reside in the root of the bundle directory under this name
2. The container’s root filesystem - a referenced directory



**Note:** Because the directory itself, i.e., `/var/tmp/busybox` is *not* part of the bundle, the mount point can be chosen arbitrarily.

The `filtered` `config.json` file corresponding to the OCI mounted `busybox_latest.sif` container can be detailed as follows via `$ sudo cat /var/tmp/busybox/config.json | jq`:

```
{
  "ociVersion": "1.0.1-dev",
  "process": {
    "user": {
      "uid": 0,
      "gid": 0
    },
    "args": [
      "/.singularity.d/actions/run"
    ],
    "env": [
      "PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
      "TERM=xterm"
    ],
    "cwd": "/",
    "capabilities": {
      "bounding": [
        "CAP_CHOWN",
        "CAP_DAC_OVERRIDE",
        "CAP_FSETID",
        "CAP_FOWNER",
        "CAP_MKNOD",
        "CAP_NET_RAW",
        "CAP_SETGID",
        "CAP_SETUID",
        "CAP_SETFCAP",
        "CAP_SETPCAP",
        "CAP_NET_BIND_SERVICE",
        "CAP_SYS_CHROOT",
        "CAP_KILL",
        "CAP_AUDIT_WRITE"
      ],
      "effective": [
        "CAP_CHOWN",
        "CAP_DAC_OVERRIDE",
        "CAP_FSETID",
        "CAP_FOWNER",
        "CAP_MKNOD",
        "CAP_NET_RAW",
        "CAP_SETGID",
        "CAP_SETUID",
        "CAP_SETFCAP",
        "CAP_SETPCAP",
        "CAP_NET_BIND_SERVICE",
        "CAP_SYS_CHROOT",
        "CAP_KILL",
        "CAP_AUDIT_WRITE"
      ],
      "inheritable": [
        "CAP_CHOWN",

```

(continues on next page)

(continued from previous page)

```

    "CAP_DAC_OVERRIDE",
    "CAP_FSETID",
    "CAP_FOWNER",
    "CAP_MKNOD",
    "CAP_NET_RAW",
    "CAP_SETGID",
    "CAP_SETUID",
    "CAP_SETFCAP",
    "CAP_SETPCAP",
    "CAP_NET_BIND_SERVICE",
    "CAP_SYS_CHROOT",
    "CAP_KILL",
    "CAP_AUDIT_WRITE"
  ],
  "permitted": [
    "CAP_CHOWN",
    "CAP_DAC_OVERRIDE",
    "CAP_FSETID",
    "CAP_FOWNER",
    "CAP_MKNOD",
    "CAP_NET_RAW",
    "CAP_SETGID",
    "CAP_SETUID",
    "CAP_SETFCAP",
    "CAP_SETPCAP",
    "CAP_NET_BIND_SERVICE",
    "CAP_SYS_CHROOT",
    "CAP_KILL",
    "CAP_AUDIT_WRITE"
  ],
  "ambient": [
    "CAP_CHOWN",
    "CAP_DAC_OVERRIDE",
    "CAP_FSETID",
    "CAP_FOWNER",
    "CAP_MKNOD",
    "CAP_NET_RAW",
    "CAP_SETGID",
    "CAP_SETUID",
    "CAP_SETFCAP",
    "CAP_SETPCAP",
    "CAP_NET_BIND_SERVICE",
    "CAP_SYS_CHROOT",
    "CAP_KILL",
    "CAP_AUDIT_WRITE"
  ]
},
"rlimits": [
  {
    "type": "RLIMIT_NOFILE",
    "hard": 1024,
    "soft": 1024
  }
]
},
"root": {
  "path": "/var/tmp/busybox/rootfs"
}

```

(continues on next page)

(continued from previous page)

```
},
"hostname": "mrsdalloway",
"mounts": [
  {
    "destination": "/proc",
    "type": "proc",
    "source": "proc"
  },
  {
    "destination": "/dev",
    "type": "tmpfs",
    "source": "tmpfs",
    "options": [
      "nosuid",
      "strictatime",
      "mode=755",
      "size=65536k"
    ]
  },
  {
    "destination": "/dev/pts",
    "type": "devpts",
    "source": "devpts",
    "options": [
      "nosuid",
      "noexec",
      "newinstance",
      "ptmxmode=0666",
      "mode=0620",
      "gid=5"
    ]
  },
  {
    "destination": "/dev/shm",
    "type": "tmpfs",
    "source": "shm",
    "options": [
      "nosuid",
      "noexec",
      "nodev",
      "mode=1777",
      "size=65536k"
    ]
  },
  {
    "destination": "/dev/mqueue",
    "type": "mqueue",
    "source": "mqueue",
    "options": [
      "nosuid",
      "noexec",
      "nodev"
    ]
  },
  {
    "destination": "/sys",
    "type": "sysfs",
```

(continues on next page)

(continued from previous page)

```
"source": "sysfs",
"options": [
  "nosuid",
  "noexec",
  "nodev",
  "ro"
]
},
],
"linux": {
  "resources": {
    "devices": [
      {
        "allow": false,
        "access": "rwm"
      }
    ]
  },
},
"namespaces": [
  {
    "type": "pid"
  },
  {
    "type": "network"
  },
  {
    "type": "ipc"
  },
  {
    "type": "uts"
  },
  {
    "type": "mount"
  }
],
"seccomp": {
  "defaultAction": "SCMP_ACT_ERRNO",
  "architectures": [
    "SCMP_ARCH_X86_64",
    "SCMP_ARCH_X86",
    "SCMP_ARCH_X32"
  ],
  "syscalls": [
    {
      "names": [
        "accept",
        "accept4",
        "access",
        "alarm",
        "bind",
        "brk",
        "capget",
        "capset",
        "chdir",
        "chmod",
        "chown",
        "chown32",
```

(continues on next page)

(continued from previous page)

```
"clock_getres",
"clock_gettime",
"clock_nanosleep",
"close",
"connect",
"copy_file_range",
"creat",
"dup",
"dup2",
"dup3",
"epoll_create",
"epoll_create1",
"epoll_ctl",
"epoll_ctl_old",
"epoll_pwait",
"epoll_wait",
"epoll_wait_old",
"eventfd",
"eventfd2",
"execve",
"execveat",
"exit",
"exit_group",
"faccessat",
"fadvise64",
"fadvise64_64",
"fallocate",
"fanotify_mark",
"fchdir",
"fchmod",
"fchmodat",
"fchown",
"fchown32",
"fchownat",
"fcntl",
"fcntl64",
"fdatasync",
"fgetxattr",
"flistxattr",
"flock",
"fork",
"removexattr",
"fsetxattr",
"fstat",
"fstat64",
"fstatat64",
"fstatfs",
"fstatfs64",
"fsync",
"ftruncate",
"ftruncate64",
"futex",
"futimesat",
"getcpu",
"getcwd",
"getdents",
"getdents64",
```

(continues on next page)

(continued from previous page)

```
"getegid",
"getegid32",
"geteuid",
"geteuid32",
"getgid",
"getgid32",
"getgroups",
"getgroups32",
"getitimer",
"getpeername",
"getpgid",
"getpgrp",
"getpid",
"getppid",
"getpriority",
"getrandom",
"getresgid",
"getresgid32",
"getresuid",
"getresuid32",
"getrlimit",
"get_robust_list",
"getrusage",
"getsid",
"getsockname",
"getsockopt",
"get_thread_area",
"gettid",
"gettimeofday",
"getuid",
"getuid32",
"getxattr",
"inotify_add_watch",
"inotify_init",
"inotify_init1",
"inotify_rm_watch",
"io_cancel",
"ioctl",
"io_destroy",
"io_getevents",
"ioprio_get",
"ioprio_set",
"io_setup",
"io_submit",
"ipc",
"kill",
"lchown",
"lchown32",
"lgetxattr",
"link",
"linkat",
"listen",
"listxattr",
"llistxattr",
"_llseek",
"lremovexattr",
"lseek",
```

(continues on next page)

(continued from previous page)

```
"lsetxattr",
"lstat",
"lstat64",
"madvise",
"memfd_create",
"mincore",
"mkdir",
"mkdirat",
"mknod",
"mknodat",
"mlock",
"mlock2",
"mlockall",
"mmap",
"mmap2",
"mprotect",
"mq_getsetattr",
"mq_notify",
"mq_open",
"mq_timedreceive",
"mq_timedsend",
"mq_unlink",
"mremap",
"msgctl",
"msgget",
"msgrcv",
"msgsnd",
"msync",
"munlock",
"munlockall",
"munmap",
"nanosleep",
"newfstatat",
"_newselect",
"open",
"openat",
"pause",
"pipe",
"pipe2",
"poll",
"ppoll",
"prctl",
"pread64",
"preadv",
"prlimit64",
"pselect6",
"pwrite64",
"pwritev",
"read",
"readahead",
"readlink",
"readlinkat",
"readv",
"recv",
"recvfrom",
"recvmmsg",
"recvmsg",
```

(continues on next page)

(continued from previous page)

```
"remap_file_pages",
"removexattr",
"rename",
"renameat",
"renameat2",
"restart_syscall",
"rmdir",
"rt_sigaction",
"rt_sigpending",
"rt_sigprocmask",
"rt_sigqueueinfo",
"rt_sigreturn",
"rt_sigsuspend",
"rt_sigtimedwait",
"rt_tgsigqueueinfo",
"sched_getaffinity",
"sched_getattr",
"sched_getparam",
"sched_get_priority_max",
"sched_get_priority_min",
"sched_getscheduler",
"sched_rr_get_interval",
"sched_setaffinity",
"sched_setattr",
"sched_setparam",
"sched_setscheduler",
"sched_yield",
"seccomp",
"select",
"semctl",
"semget",
"semop",
"semtimedop",
"send",
"sendfile",
"sendfile64",
"sendmmsg",
"sendmsg",
"sendto",
"setfsgid",
"setfsgid32",
"setfsuid",
"setfsuid32",
"setgid",
"setgid32",
"setgroups",
"setgroups32",
"setitimer",
"setpgid",
"setpriority",
"setregid",
"setregid32",
"setresgid",
"setresgid32",
"setresuid",
"setresuid32",
"setreuid",
```

(continues on next page)



(continued from previous page)

```
"setreuid32",
"setrlimit",
"set_robust_list",
"setsid",
"setsockopt",
"set_thread_area",
"set_tid_address",
"setuid",
"setuid32",
"setxattr",
"shmat",
"shmctl",
"shmdt",
"shmget",
"shutdown",
"sigaltstack",
"signalfd",
"signalfd4",
"sigreturn",
"socket",
"socketcall",
"socketpair",
"splice",
"stat",
"stat64",
"statfs",
"statfs64",
"symlink",
"symlinkat",
"sync",
"sync_file_range",
"syncfs",
"sysinfo",
"syslog",
"tee",
"tgkill",
"time",
"timer_create",
"timer_delete",
"timerfd_create",
"timerfd_gettime",
"timerfd_settime",
"timer_getoverrun",
"timer_gettime",
"timer_settime",
"times",
"tkill",
"truncate",
"truncate64",
"ugetrlimit",
"umask",
"uname",
"unlink",
"unlinkat",
"utime",
"utimensat",
"utimes",
```

(continues on next page)

(continued from previous page)

```

        "vfork",
        "vmsplice",
        "wait4",
        "waitid",
        "waitpid",
        "write",
        "writev"
    ],
    "action": "SCMP_ACT_ALLOW"
},
{
    "names": [
        "personality"
    ],
    "action": "SCMP_ACT_ALLOW",
    "args": [
        {
            "index": 0,
            "value": 0,
            "op": "SCMP_CMP_EQ"
        },
        {
            "index": 0,
            "value": 8,
            "op": "SCMP_CMP_EQ"
        },
        {
            "index": 0,
            "value": 4294967295,
            "op": "SCMP_CMP_EQ"
        }
    ]
},
{
    "names": [
        "chroot"
    ],
    "action": "SCMP_ACT_ALLOW"
},
{
    "names": [
        "clone"
    ],
    "action": "SCMP_ACT_ALLOW",
    "args": [
        {
            "index": 0,
            "value": 2080505856,
            "op": "SCMP_CMP_MASKED_EQ"
        }
    ]
},
{
    "names": [
        "arch_prctl"
    ],
    "action": "SCMP_ACT_ALLOW"
}

```

(continues on next page)

(continued from previous page)

```

    },
    {
      "names": [
        "modify_ldt"
      ],
      "action": "SCMP_ACT_ALLOW"
    }
  ]
}
}
}
}
}

```

Furthermore, and through use of `$ sudo cat /var/tmp/busybox/config.json | jq [.root.path]`, the property

```

[
  "/var/tmp/busybox/rootfs"
]

```

identifies `/var/tmp/busybox/rootfs` as the container's root filesystem, as required by the standard; this filesystem has contents:

```

$ sudo ls /var/tmp/busybox/rootfs
bin dev environment etc home proc root singularity sys tmp usr var

```

---

**Note:** `environment` and `singularity` above are symbolic links to the `.singularity.d` directory.

---

Beyond `root.path`, the `config.json` file includes a multitude of additional properties - for example:

- `ociVersion` - a mandatory property that identifies the version of the OCI runtime specification that the bundle is compliant with
- `process` - an optional property that specifies the container process. When invoked via Singularity, subproperties such as `args` are populated by making use of the contents of the `.singularity.d` directory, e.g. via `$ sudo cat /var/tmp/busybox/config.json | jq [.process.args]`:

```

[
  [
    "/.singularity.d/actions/run"
  ]
]

```

where `run` equates to the *familiar runscript* for this container. If image creation is bootstrapped via a Docker or OCI agent, Singularity will make use of `ENTRYPOINT` or `CMD` (from the OCI image) to populate `args`; for additional discussion, please refer to *Directing Execution* in the section *Support for Docker and OCI*.

For a comprehensive discussion of all the `config.json` file properties, refer to the [implementation guide](#).

Technically, the `overlay` directory was *not* content expected of an OCI compliant filesystem bundle. As detailed in the section dedicated to [Persistent Overlays](#), these directories allow for the introduction of a writable file system on an otherwise immutable read-only container; thus they permit the illusion of read-write access.

---

**Note:** SIF is stated to be an extensible format; by encapsulating a filesystem bundle that conforms with the OCI

---

runtime specification, this extensibility is evident.

---

## 14.3 Creating OCI Compliant Container Instances

SIF files encapsulate the OCI runtime. By ‘OCI mounting’ a SIF file (see above), this encapsulated runtime is revealed; please refer to the note below for additional details. Once revealed, the filesystem bundle can be used to bootstrap the creation of an OCI compliant container instance as follows:

```
$ sudo singularity oci create -b /var/tmp/busybox busybox1
```

**Note:** Data for the `config.json` file exists within the SIF file as a descriptor for images pulled or built from Docker/OCI registries. For images sourced elsewhere, a default `config.json` file is created when the `singularity oci mount ...` command is issued.

Upon invocation, `singularity oci mount ...` also mounts the root filesystem stored in the SIF file on `/bundle/rootfs`, and establishes an overlay filesystem on the mount point `/bundle/overlay`.

---

In this example, the filesystem bundle is located in the directory `/var/tmp/busybox` - i.e., the mount point identified above with respect to ‘OCI mounting’. The `config.json` file, along with the `rootfs` and `overlay` filesystems, are all employed in the bootstrap process. The instance is named `busybox1` in this example.

---

**Note:** The outcome of this creation request is truly a container **instance**. Multiple instances of the same container can easily be created by simply changing the name of the instance upon subsequent invocation requests.

---

The state of the container instance can be determined via `$ sudo singularity oci state busybox1`:

```
{
"ociVersion": "1.0.1-dev",
"id": "busybox1",
"status": "created",
"pid": 6578,
"bundle": "/var/tmp/busybox",
"createdAt": 1554389921452964253,
"attachSocket": "/var/run/singularity/instances/root/busybox1/attach.sock",
"controlSocket": "/var/run/singularity/instances/root/busybox1/control.sock"
}
```

Container state, as conveyed via these properties, is in compliance with the OCI runtime specification as detailed [here](#).

The `create` command has a number of options available. Of these, real-time logging to a file is likely to be of particular value - e.g., in deployments where auditing requirements exist.

## 14.4 Unmounting OCI Filesystem Bundles

To unmount a mounted OCI filesystem bundle, the following command should be issued:

```
$ sudo singularity oci umount /var/tmp/busybox
```

---


**Note:** The argument provided to `oci mount` above is the name of the bundle path, `/var/tmp/busybox`, as opposed to the mount point for the overlay filesystem, `/var/tmp/busybox/rootfs`.

---

## 14.5 Kubernetes Integration

As noted at the [outset here](#), in documenting support for an OCI runtime in Singularity, the impetus is initially derived from the requirement to integrate with [Kubernetes](#). Simply stated, Kubernetes is an open-source system for orchestrating containers; developed originally at Google, Kubernetes was contributed as seed technology to the [Cloud Native Compute Foundation \(CNCF\)](#). At this point, Kubernetes is regarded as a Graduated Project by CNCF, and is being used widely in production deployments. Even though Kubernetes emphasizes an orientation around services, it is appealing to those seeking to orchestrate containers having compute-driven requirements. Furthermore, emerging classes of workload in AI for example, appear to have requirements that are best addressed by a combination of service and traditional HPC infrastructures. Thus there is ample existing, as well as emerging, interest in integrating Singularity containers with Kubernetes.

The connection with support for the OCI runtime documented here, within the context of a Singularity-Kubernetes integration, can be best established through an architectural schematic. Dating back to the introduction of a Container Runtime Interface (CRI) for Kubernetes in late 2016, the schematic below is a modified version of the original presented in a [Kubernetes blog post](#). The lower branch of this schematic is essentially a reproduction of the original; it does however, place emphasis on OCI compliance in terms of the CRI and containers (the runtime as well as their instances).



From this schematic it is evident that integrating Singularity containers with Kubernetes requires the following efforts:

1. Implementation of a CRI for Singularity
2. Implementation of an OCI runtime in Singularity

The implementation of a CRI for Singularity is the emphasis of a separate and distinct [open source project](#); the implementation of this CRI is documented here. For the rationale conveyed through the architectural schematic, Singularity CRI's dependence upon Singularity with OCI runtime support is made clear as an [installation prerequisite](#). User-facing documentation for Singularity CRI details usage in a Kubernetes context - usage, of course, that involves orchestration of a Singularity container obtained from the [Sylabs Cloud Container Library](#). Because the entire Kubernetes-based

deployment can exist within a single instance of a Singularity container, Singularity CRI can be easily evaluated via Sykube; inspired by [Minikube](#), [use of Sykube](#) is included in the documentation for Singularity CRI.

Documenting the implementation of an OCI-compliant runtime for Singularity has been the emphasis here. Although this standalone runtime can be used by end users independent of anything to do with Singularity and Kubernetes, the primary purpose here has been documenting it within this integrated context. In other words, by making use of the OCI runtime presented by Singularity, commands originating from Kubernetes (see, e.g., [Basic Usage](#) in the Singularity CRI documentation) have impact ultimately on Singularity containers via the CRI. Singularity CRI is implemented as a [gRPC](#) server - i.e., a persistent service available to [Kubelets](#) (node agents). Taken together, this integration allows Singularity containers to be manipulated directly from Kubernetes.

## KEY COMMANDS

Singularity 3.2 introduces the abilities to import, export and remove PGP keys following the OpenPGP standard via **GnuPGP (GPG)**. These commands only modify the local keyring and are not related to the cloud keystore.

### 15.1 Key import command

Singularity 3.2 allows you import keys reading either from binary or armored key format and automatically detect if it is a private or public key and add it to the correspondent local keystore.

To give a quick view on how it works, we will first consider the case in which a user wants to import a secret (private) key to the local keystore.

First we will check what's the status of the local keystore (which keys are stored by the moment before importing a new key).

```
$ singularity key list --secret
```

---

**Note:** Remember that using `--secret` flag or `-s` flag will return the secret or private local keyring as output.

---

The output will look as it follows:

```
Private key listing (/home/joana/.singularity/syppgp/pgp-secret):  
  
0) U: Johnny Cash (none) <cash@sylabs.io>  
C: 2019-04-11 22:22:28 +0200 CEST  
F: 47282BDC661F58FA4BEBEF47CA576CBD8EF1A2B4  
L: 3072  
-----  
1) U: John Green (none) <john@sylabs.io>  
C: 2019-04-11 13:08:45 +0200 CEST  
F: 5720799FE7B048CF36FAB8445EE1E2BD7B6342C5  
L: 1024  
-----
```

---

**Note:** Remember that running that same command but with `sudo` privilege, will give you a totally different list since it will be the correspondent keystore from user `root`

---

After this, you can simply import the key you need by adding the exact location to the file, let's say you own a gpg key file named `pinkie-pie.asc` which is a secret GPG key you want to import. Then you will just need to run the following command to import your key:

```
$ singularity key import $HOME/pinkie-pie.asc
```

**Note:** This location is considering your key was located on the `$HOME` directory. You can specify any location to the file.

---

Since you're importing a private (secret) key, you will need to specify the passphrase related to it and then a new passphrase to be added on your local keystore.

```
Enter your old password :
Enter a new password for this key :
Retype your passphrase :
Key with fingerprint 8C10B902F438E4D504C3ACF689FCFFAED5F34A77 successfully added to
↳the keyring
```

After this you can see if that key was correctly added to your local keystore by running `singularity key list -s` command:

```
Private key listing (/home/joana/.singularity/syppg/pgp-secret):

0) U: Johnny Cash (none) <cash@sylabs.io>
C: 2019-04-11 22:22:28 +0200 CEST
F: 47282BDC661F58FA4BEBEF47CA576CBD8EF1A2B4
L: 3072
-----
1) U: John Green (none) <john@sylabs.io>
C: 2019-04-11 13:08:45 +0200 CEST
F: 5720799FE7B048CF36FAB8445EE1E2BD7B6342C5
L: 1024
-----
3) U: Pinkie Pie (Eternal chaos comes with chocolate rain!) <balloons@sylabs.io>
C: 2019-04-26 12:07:07 +0200 CEST
F: 8C10B902F438E4D504C3ACF689FCFFAED5F34A77
L: 1024
-----
```

You will see the imported key at the bottom of the list. Remember you can also import an `ascii` armored key and this will be automatically detected by the `key import` command (no need to specify the format).

---

**Note:** In case you would like to import a public key the process remains the same, as the import command will automatically detect whether this key to be imported is either public or private.

---

## 15.2 Key export command

The key export command allows you to export a key that is on your local keystore. This key could be either private or public, and the key can be exported on ASCII armored format or on binary format. Of course to identify the keystore and the format the syntax varies from the `key import` command.

For example to export a public key in binary format you can run:

```
$ singularity key export 8C10B902F438E4D504C3ACF689FCFFAED5F34A77 $HOME/mykey.asc
```



This will export a public binary key named `mykey.asc` and will save it under the home folder. If you would like to export the same public key but in an ASCII armored format, you would need to run the following command:

```
$ singularity key export --armor 8C10B902F438E4D504C3ACF689FCFFAED5F34A77 $HOME/mykey.  
↪asc
```

And in the case in which you may need to export a secret key on ASCII armored format, you would need to specify from where to find the key, since the fingerprint is the same.

```
$ singularity key export --armor --secret 8C10B902F438E4D504C3ACF689FCFFAED5F34A77  
↪$HOME/mykey.asc
```

and on binary format instead:

```
$ singularity key export --secret 8C10B902F438E4D504C3ACF689FCFFAED5F34A77 $HOME/  
↪mykey.asc
```

---

**Note:** Exporting keys will not change the status of your local keystore or keyring. This will just obtain the content of the keys and save it on a local file on your host.

---

## 15.3 Key remove command

In case you would want to remove a public key from your public local keystore, you can do so by running the following command:

```
$ singularity key remove 8C10B902F438E4D504C3ACF689FCFFAED5F34A77
```

---

**Note:** Remember that this will only delete the public key and not the private one with the same matching fingerprint.

---



## SIGNING AND VERIFYING CONTAINERS

Singularity 3.0 introduces the abilities to create and manage PGP keys and use them to sign and verify containers. This provides a trusted method for Singularity users to share containers. It ensures a bit-for-bit reproduction of the original container as the author intended it.

### 16.1 Verifying containers from the Container Library

The `verify` command will allow you to verify that a container has been signed using a PGP key. To use this feature with images that you pull from the container library, you must first generate an access token to the Sylabs Cloud. If you don't already have a valid access token, follow these steps:

- 1) Go to : <https://cloud.sylabs.io/>
- 2) Click “Sign in to Sylabs” and follow the sign in steps.
- 3) Click on your login id (same and updated button as the Sign in one).
- 4) Select “Access Tokens” from the drop down menu.
- 5) Click the “Manage my API tokens” button from the “Account Management” page.
- 6) Click “Create”.
- 7) Click “Copy token to Clipboard” from the “New API Token” page.
- 8) Paste the token string into your `~/.singularity/sylabs-token` file.

Now you can verify containers that you pull from the library, ensuring they are bit-for-bit reproductions of the original image.

```
$ singularity pull library://alpine
$ singularity verify alpine_latest.sif
Verifying image: alpine_latest.sif
Data integrity checked, authentic and signed by:
  Sylabs Admin <support@sylabs.io>, KeyID 51BE5020C508C7E9
```

In this example you can see that **Sylabs Admin** has signed the container.

### 16.2 Signing your own containers

#### 16.2.1 Generating and managing PGP keys

To sign your own containers you first need to generate one or more keys.

If you attempt to sign a container before you have generated any keys, Singularity will guide you through the interactive process of creating a new key. Or you can use the `newpair` subcommand in the `key` command group like so:

```
$ singularity key newpair
Enter your name (e.g., John Doe) : Dave Godlove
Enter your email address (e.g., john.doe@example.com) : d@sylabs.io
Enter optional comment (e.g., development keys) : demo
Generating Entity and OpenPGP Key Pair... Done
Enter encryption passphrase :
```

The `list` subcommand will show you all of the keys you have created or saved locally.‘

```
$ singularity key list
Public key listing (/home/david/.singularity/sypgp/pgp-public):

0) U: Dave Godlove (demo) <d@sylabs.io>
   C: 2018-10-08 15:25:30 -0400 EDT
   F: 135E426D67D8416DE1D6AC7FFED5BBA38EE0DC4A
   L: 4096
   -----
```

In the output above, the letters stand for the following:

- U: User
- C: Creation date and time
- F: Fingerprint
- L: Key length

After generating your key you can optionally push it to the [Keystore](#) using the fingerprint like so:

```
$ singularity key push 135E426D67D8416DE1D6AC7FFED5BBA38EE0DC4A
public key `135E426D67D8416DE1D6AC7FFED5BBA38EE0DC4A` pushed to server successfully
```

This will allow others to verify images that you have signed.

If you delete your local public PGP key, you can always locate and download it again like so.

```
$ singularity key search Godlove
Search results for 'Godlove'

Type bits/keyID      Date          User ID
-----
pub  4096R/8EE0DC4A 2018-10-08 Dave Godlove (demo) <d@sylabs.io>
-----

$ singularity key pull 8EE0DC4A
1 key(s) fetched and stored in local cache /home/david/.singularity/sypgp/pgp-public
```

But note that this only restores the *public* key (used for verifying) to your local machine and does not restore the *private* key (used for signing).

### 16.2.2 Searching for keys

Singularity allows you to search the keystore for public keys. You can search for names, emails, and fingerprints (key IDs). When searching for a fingerprint, you need to use `0x` before the fingerprint, check the example:

```
# search for key ID:
$ singularity key search 0x8883491F4268F173C6E5DC49EDECE4F3F38D871E

# search for the sort ID:
$ singularity key search 0xF38D871E

# search for user:
$ singularity key search Godlove

# search for email:
$ singularity key search @gmail.com
```

### 16.2.3 Signing and validating your own containers

Now that you have a key generated, you can use it to sign images like so:

```
$ singularity sign my_container.sif
Signing image: my_container.sif
Enter key passphrase:
Signature created and applied to my_container.sif
```

Because your public PGP key is saved locally you can verify the image without needing to contact the Keystore.

```
$ singularity verify my_container.sif
Verifying image: my_container.sif
Data integrity checked, authentic and signed by:
    Dave Godlove (demo) <d@sylabs.io>, KeyID FED5BBA38EE0DC4A
```

If you've pushed your key to the Keystore you can also verify this image in the absence of a local key. To demonstrate this, first delete your local keys, and then try to use the `verify` command again.

```
$ rm ~/.singularity/syppg/*

$ singularity verify my_container.sif
Verifying image: my_container.sif
INFO:   key missing, searching key server for KeyID: FED5BBA38EE0DC4A...
INFO:   key retrieved successfully!
Store new public key 135E426D67D8416DE1D6AC7FFED5BBA38EE0DC4A? [Y/n] y
Data integrity checked, authentic and signed by:
    Dave Godlove (demo) <d@sylabs.io>, KeyID FED5BBA38EE0DC4A
```

Answering yes at the interactive prompt will store the Public key locally so you will not have to contact the Keystore again the next time you verify your container.



## 17.1 Overview

A Singularity plugin is a package that can be dynamically loaded by the Singularity runtime, augmenting Singularity with experimental, non-standard and/or vendor-specific functionality. Currently, plugins are able to add commands and flags to Singularity. In the future, plugins will also be able to interface with more complex subsystems of the Singularity runtime.

## 17.2 Using Plugins

The `list` command prints the currently installed plugins.

```
$ singularity plugin list
There are no plugins installed.
```

Plugins are packaged and distributed as binaries encoded with the versatile Singularity Image Format (SIF). However, plugin authors may also distribute the source code of their plugins. A plugin can be compiled from its source code with the `compile` command. A sample plugin `test-plugin` is included with the Singularity source code.

```
$ singularity plugin compile examples/plugins/test-plugin/
```

Upon successful compilation, a SIF file will appear in the directory of the plugin's source code.

```
$ ls examples/plugins/test-plugin/ | grep sif
test-plugin.sif
```

---

**Note:** Currently, **all** plugins must be compiled from the Singularity source code tree.

---

Every plugin encapsulates various information such as the plugin's author, the plugin's version, etc. To view this information about a plugin, use the `inspect` command.

```
$ singularity plugin inspect examples/plugins/test-plugin/test-plugin.sif
Name: sylabs.io/test-plugin
Description: This is a short test plugin for Singularity
Author: Michael Bauer
Version: 0.0.1
```

To install a plugin, use the `install` command. This operation requires root privilege.

```
$ sudo singularity plugin install examples/plugins/test-plugin/test-plugin.sif
$ singularity plugin list
ENABLED  NAME
      yes  sylabs.io/test-plugin
```

After successful installation, the plugin will automatically be enabled. Any plugin can be disabled with the `disable` command and re-enabled with the `enable` command. Both of these operations require root privilege.

```
$ sudo singularity plugin disable sylabs.io/test-plugin
$ singularity plugin list
ENABLED  NAME
      no  sylabs.io/test-plugin
$ sudo singularity plugin enable sylabs.io/test-plugin
$ singularity plugin list
ENABLED  NAME
      yes  sylabs.io/test-plugin
```

Finally, to uninstall a plugin, use the `uninstall` command. This operation requires root privilege.

```
$ sudo singularity plugin uninstall sylabs.io/test-plugin
Uninstalled plugin "sylabs.io/test-plugin".
$ singularity plugin list
There are no plugins installed.
```

## 17.3 Writing a Plugin

Developers interested in writing Singularity plugins can get started by reading the [Go documentation](#) for the plugin package. Furthermore, reading through the [source code](#) for the example test plugin will prove valuable. More detailed plugin development documentation is in the works and will be released at a future date.



## SECURITY OPTIONS

Singularity 3.0 introduces many new security related options to the container runtime. This document will describe the new methods users have for specifying the security scope and context when running Singularity containers.

### 18.1 Linux Capabilities

Singularity provides full support for granting and revoking Linux capabilities on a user or group basis. For example, let us suppose that an admin has decided to grant a user capabilities to open raw sockets so that they can use `ping` in a container where the binary is controlled via capabilities (i.e. a recent version of CentOS).

To do so, the admin would issue a command such as this:

```
$ sudo singularity capability add --user david CAP_NET_RAW
```

This means the user `david` has just been granted permissions (through Linux capabilities) to open raw sockets within Singularity containers.

The admin can check that this change is in effect with the `capability list` command.

```
$ sudo singularity capability list --user david
CAP_NET_RAW
```

To take advantage of this new capability, the user `david` must also request the capability when executing a container with the `--add-caps` flag like so:

```
$ singularity exec --add-caps CAP_NET_RAW library://centos ping -c 1 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=128 time=18.3 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 18.320/18.320/18.320/0.000 ms
```

If the admin decides that it is no longer necessary to allow the user `dave` to open raw sockets within Singularity containers, they can revoke the appropriate Linux capability like so:

```
$ sudo singularity capability drop --user david CAP_NET_RAW
```

The `capability add` and `drop` subcommands will also accept the case insensitive keyword `all` to grant or revoke all Linux capabilities to a user or group. Similarly, the `--add-caps` option will accept the `all` keyword. Of course appropriate caution should be exercised when using this keyword.

## 18.2 Security related action options

Singularity 3.0 introduces many new flags that can be passed to the action commands; `shell`, `exec`, and `run` allowing fine grained control of security.

### 18.2.1 `--add-caps`

As explained above, `--add-caps` will “activate” Linux capabilities when a container is initiated, providing those capabilities have been granted to the user by an administrator using the `capability add` command. This option will also accept the case insensitive keyword `all` to add every capability granted by the administrator.

### 18.2.2 `--allow-setuid`

The SetUID bit allows a program to be executed as the user that owns the binary. The most well-known SetUID binaries are owned by root and allow a user to execute a command with elevated privileges. But other SetUID binaries may allow a user to execute a command as a service account.

By default SetUID is disallowed within Singularity containers as a security precaution. But the root user can override this precaution and allow SetUID binaries to behave as expected within a Singularity container with the `--allow-setuid` option like so:

```
$ sudo singularity shell --allow-setuid some_container.sif
```

### 18.2.3 `--keep-privs`

It is possible for an admin to set a different set of default capabilities or to reduce the default capabilities to zero for the root user by setting the `root default capabilities` parameter in the `singularity.conf` file to `file` or `no` respectively. If this change is in effect, the root user can override the `singularity.conf` file and enter the container with full capabilities using the `--keep-privs` option.

```
$ sudo singularity exec --keep-privs library://centos ping -c 1 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=128 time=18.8 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 18.838/18.838/18.838/0.000 ms
```

### 18.2.4 `--drop-caps`

By default, the root user has a full set of capabilities when they enter the container. You may choose to drop specific capabilities when you initiate a container as root to enhance security.

For instance, to drop the ability for the root user to open a raw socket inside the container:

```
$ sudo singularity exec --drop-caps CAP_NET_RAW library://centos ping -c 1 8.8.8.8
ping: socket: Operation not permitted
```

The `drop-caps` option will also accept the case insensitive keyword `all` as an option to drop all capabilities when entering the container.

## 18.2.5 --security

The `--security` flag allows the root user to leverage security modules such as SELinux, AppArmor, and seccomp within your Singularity container. You can also change the UID and GID of the user within the container at runtime.

For instance:

```
$ sudo whoami
root

$ sudo singularity exec --security uid:1000 my_container.sif whoami
david
```

To use seccomp to blacklist a command follow this procedure. (It is actually preferable from a security standpoint to whitelist commands but this will suffice for a simple example.) Note that this example was run on Ubuntu and that Singularity was installed with the `libseccomp-dev` and `pkg-config` packages as dependencies.

First write a configuration file. An example configuration file is installed with Singularity, normally at `/usr/local/etc/singularity/seccomp-profiles/default.json`. For this example, we will use a much simpler configuration file to blacklist the `mkdir` command.

```
{
  "defaultAction": "SCMP_ACT_ALLOW",
  "archMap": [
    {
      "architecture": "SCMP_ARCH_X86_64",
      "subArchitectures": [
        "SCMP_ARCH_X86",
        "SCMP_ARCH_X32"
      ]
    }
  ],
  "syscalls": [
    {
      "names": [
        "mkdir"
      ],
      "action": "SCMP_ACT_KILL",
      "args": [],
      "comment": "",
      "includes": {},
      "excludes": {}
    }
  ]
}
```

We'll save the file at `/home/david/no_mkdir.json`. Then we can invoke the container like so:

```
$ sudo singularity shell --security seccomp:/home/david/no_mkdir.json my_container.sif

Singularity> mkdir /tmp/foo
Bad system call (core dumped)
```

Note that attempting to use the blacklisted `mkdir` command resulted in a core dump.

The full list of arguments accepted by the `--security` option are as follows:

```
--security="seccomp:/usr/local/etc/singularity/seccomp-profiles/default.json"  
--security="apparmor:/usr/bin/man"  
--security="selinux:context"  
--security="uid:1000"  
--security="gid:1000"  
--security="gid:1000:1:0" (multiple gids, first is always the primary group)
```

## NETWORK VIRTUALIZATION

Singularity 3.0 introduces full integration with `cni`, and several new features to make network virtualization easy.

A few new options have been added to the action commands (`exec`, `run`, and `shell`) to facilitate these features, and the `--net` option has been updated as well. These options can only be used by root.

### 19.1 `--dns`

The `--dns` option allows you to specify a comma separated list of DNS servers to add to the `/etc/resolv.conf` file.

```
$ nslookup sylabs.io | grep Server
Server:                127.0.0.53

$ sudo singularity exec --dns 8.8.8.8 ubuntu.sif nslookup sylabs.io | grep Server
Server:                8.8.8.8

$ sudo singularity exec --dns 8.8.8.8 ubuntu.sif cat /etc/resolv.conf
nameserver 8.8.8.8
```

### 19.2 `--hostname`

The `--hostname` option accepts a string argument to change the hostname within the container.

```
$ hostname
ubuntu-bionic

$ sudo singularity exec --hostname hal-9000 my_container.sif hostname
hal-9000
```

### 19.3 `--net`

Passing the `--net` flag will cause the container to join a new network namespace when it initiates. New in Singularity 3.0, a bridge interface will also be set up by default.

```
$ hostname -I
10.0.2.15
```

(continues on next page)

(continued from previous page)

```
$ sudo singularity exec --net my_container.sif hostname -I
10.22.0.4
```

## 19.4 --network

The `--network` option can only be invoked in combination with the `--net` flag. It accepts a comma delimited string of network types. Each entry will bring up a dedicated interface inside container.

```
$ hostname -I
172.16.107.251 10.22.0.1

$ sudo singularity exec --net --network ptp ubuntu.sif hostname -I
10.23.0.6

$ sudo singularity exec --net --network bridge,ptp ubuntu.sif hostname -I
10.22.0.14 10.23.0.7
```

When invoked, the `--network` option searches the singularity configuration directory (commonly `/usr/local/etc/singularity/network/`) for the cni configuration file corresponding to the requested network type(s). Several configuration files are installed with Singularity by default corresponding to the following network types:

- bridge
- ptp
- ipvlan
- macvlan
- none (must be used alone)

None is the only network option that can be used by non-privileged users. It isolates the container network from the host network with a loopback interface.

Administrators can also define custom network configurations and place them in the same directory for the benefit of users.

## 19.5 --network-args

The `--network-args` option provides a convenient way to specify arguments to pass directly to the cni plugins. It must be used in conjunction with the `--net` flag.

For instance, let's say you want to start an [NGINX](#) server on port 80 inside of the container, but you want to map it to port 8080 outside of the container:

```
$ sudo singularity instance start --writable-tmpfs \
  --net --network-args "portmap=8080:80/tcp" docker://nginx web2
```

The above command will start the Docker Hub official NGINX image running in a background instance called `web2`. The NGINX instance will need to be able to write to disk, so we've used the `--writable-tmpfs` argument to allocate some space in memory. The `--net` flag is necessary when using the `--network-args` option, and specifying the `portmap=8080:80/tcp` argument which will map port 80 inside of the container to 8080 on the host.

Now we can start NGINX inside of the container:

```
$ sudo singularity exec instance://web2 nginx
```

And the `curl` command can be used to verify that NGINX is running on the host port 8080 as expected.

```
$ curl localhost:8080
10.22.0.1 - - [16/Oct/2018:09:34:25 -0400] "GET / HTTP/1.1" 200 612 "-" "curl/7.58.0"
↪ "-"
<!DOCTYPE html>
<html>
<head>
<title>Welcome to nginx!</title>
<style>
  body {
    width: 35em;
    margin: 0 auto;
    font-family: Tahoma, Verdana, Arial, sans-serif;
  }
</style>
</head>
<body>
<h1>Welcome to nginx!</h1>
<p>If you see this page, the nginx web server is successfully installed and
working. Further configuration is required.</p>

<p>For online documentation and support please refer to
<a href="http://nginx.org/">nginx.org</a>.<br/>
Commercial support is available at
<a href="http://nginx.com/">nginx.com</a>.</p>

<p><em>Thank you for using nginx.</em></p>
</body>
</html>
```

For more information about cni, check the [cni specification](#).





## LIMITING CONTAINER RESOURCES WITH CGROUPS

Starting in Singularity 3.0, users have the ability to limit container resources using cgroups.

### 20.1 Overview

Singularity cgroups support can be configured and utilized via a TOML file. An example file is typically installed at `/usr/local/etc/singularity/cgroups/cgroups.toml`. You can copy and edit this file to suit your needs. Then when you need to limit your container resources, apply the settings in the TOML file by using the path as an argument to the `--apply-cgroups` option like so:

```
$ sudo singularity shell --apply-cgroups /path/to/cgroups.toml my_container.sif
```

The `--apply-cgroups` option can only be used with root privileges.

### 20.2 Examples

#### 20.2.1 Limiting memory

To limit the amount of memory that your container uses to 500MB (524288000 bytes), follow this example. First, create a `cgroups.toml` file like this and save it in your home directory.

```
[memory]
  limit = 524288000
```

Start your container like so:

```
$ sudo singularity instance start --apply-cgroups /home/$USER/cgroups.toml \
  my_container.sif instance1
```

After that, you can verify that the container is only using 500MB of memory. (This example assumes that `instance1` is the only running instance.)

```
$ cat /sys/fs/cgroup/memory/singularity/*/memory.limit_in_bytes
524288000
```

After you are finished with this example, be sure to cleanup your instance with the following command.

```
$ sudo singularity instance stop instance1
```

Similarly, the remaining examples can be tested by starting instances and examining the contents of the appropriate subdirectories of `/sys/fs/cgroup/`.

### 20.2.2 Limiting CPU

Limit CPU resources using one of the following strategies. The `cpu` section of the configuration file can limit memory with the following:

#### 20.2.2.1 shares

This corresponds to a ratio versus other `cgroups` with `cpu` shares. Usually the default value is `1024`. That means if you want to allow to use 50% of a single CPU, you will set `512` as value.

```
[cpu]
  shares = 512
```

A `cgroup` can get more than its share of CPU if there are enough idle CPU cycles available in the system, due to the work conserving nature of the scheduler, so a contained process can consume all CPU cycles even with a ratio of 50%. The ratio is only applied when two or more processes conflicts with their needs of CPU cycles.

#### 20.2.2.2 quota/period

You can enforce hard limits on the CPU cycles a `cgroup` can consume, so contained processes can't use more than the amount of CPU time set for the `cgroup`. `quota` allows you to configure the amount of CPU time that a `cgroup` can use per period. The default is 100ms (100000us). So if you want to limit amount of CPU time to 20ms during period of 100ms:

```
[cpu]
  period = 100000
  quota = 20000
```

#### 20.2.2.3 cpus/mems

You can also restrict access to specific CPUs and associated memory nodes by using `cpus/mems` fields:

```
[cpu]
  cpus = "0-1"
  mems = "0-1"
```

Where container has limited access to CPU 0 and CPU 1.

---

**Note:** It's important to set identical values for both `cpus` and `mems`.

---

For more information about limiting CPU with `cgroups`, see the following external links:

- [Red Hat resource management guide section 3.2 CPU](#)
- [Red Hat resource management guide section 3.4 CPuset](#)
- [Kernel scheduler documentation](#)

### 20.2.3 Limiting IO

You can limit and monitor access to I/O for block devices. Use the `[blockIO]` section of the configuration file to do this like so:

```
[blockIO]
  weight = 1000
  leafWeight = 1000
```

`weight` and `leafWeight` accept values between 10 and 1000.

`weight` is the default weight of the group on all the devices until and unless overridden by a per device rule.

`leafWeight` relates to `weight` for the purpose of deciding how heavily to weigh tasks in the given cgroup while competing with the cgroup's child cgroups.

To override `weight/leafWeight` for `/dev/loop0` and `/dev/loop1` block devices you would do something like this:

```
[blockIO]
  [[blockIO.weightDevice]]
    major = 7
    minor = 0
    weight = 100
    leafWeight = 50
  [[blockIO.weightDevice]]
    major = 7
    minor = 1
    weight = 100
    leafWeight = 50
```

You could limit the IO read/write rate to 16MB per second for the `/dev/loop0` block device with the following configuration. The rate is specified in bytes per second.

```
[blockIO]
  [[blockIO.throttleReadBpsDevice]]
    major = 7
    minor = 0
    rate = 16777216
  [[blockIO.throttleWriteBpsDevice]]
    major = 7
    minor = 0
    rate = 16777216
```

To limit the IO read/write rate to 1000 IO per second (IOPS) on `/dev/loop0` block device, you can do the following. The rate is specified in IOPS.

```
[blockIO]
  [[blockIO.throttleReadIOPSDevice]]
    major = 7
    minor = 0
    rate = 1000
  [[blockIO.throttleWriteIOPSDevice]]
    major = 7
    minor = 0
    rate = 1000
```

For more information about limiting IO, see the following external links:

- [Red Hat resource management guide section 3.1 blkio](#)
- [Kernel block IO controller documentation](#)
- [Kernel CFQ scheduler documentation](#)

### 20.2.3.1 Limiting device access

You can limit read, write, or creation of devices. In this example, a container is configured to only be able to read from or write to `/dev/null`.

```
[[devices]]
  access = "rwm"
  allow = false
[[devices]]
  access = "rw"
  allow = true
  major = 1
  minor = 3
  type = "c"
```

For more information on limiting access to devices the [Red Hat resource management guide section 3.5 DEVICES](#).

## 21.1 Singularity's environment variables

Singularity 3.0 comes with some environment variables you can set or modify depending on your needs. You can see them listed alphabetically below with their respective functionality.

### 21.1.1 A

1. **SINGULARITY\_ADD\_CAPS**: To specify a list (comma separated string) of capabilities to be added. Default is an empty string.
2. **SINGULARITY\_ALL**: List all the users and groups capabilities.
3. **SINGULARITY\_ALLOW\_SETUID**: To specify that setuid binaries should or not be allowed in the container. (root only) Default is set to false.
4. **SINGULARITY\_APP** and **SINGULARITY\_APPNAME**: Sets the name of an application to be run inside a container.
5. **SINGULARITY\_APPLY\_CGROUPS**: Used to apply cgroups from an input file for container processes. (it requires root privileges)

### 21.1.2 B

1. **SINGULARITY\_BINDPATH** and **SINGULARITY\_BIND**: Comma separated string `source:<dest>` list of paths to bind between the host and the container.
2. **SINGULARITY\_BOOT**: Set to false by default, considers if executing `/sbin/init` when container boots (root only).
3. **SINGULARITY\_BUILDER**: To specify the remote builder service URL. Defaults to our remote builder.

### 21.1.3 C

1. **SINGULARITY\_CACHEDIR**: Specifies the directory for image downloads to be cached in.
2. **SINGULARITY\_CLEANENV**: Specifies if the environment should be cleaned or not before running the container. Default is set to false.
3. **SINGULARITY\_CONTAIN**: To use minimal `/dev` and empty other directories (e.g. `/tmp` and `$HOME`) instead of sharing filesystems from your host. Default is set to false.

4. **SINGULARITY\_CONTAINALL**: To contain not only file systems, but also PID, IPC, and environment. Default is set to false.
5. **SINGULARITY\_CONTAINLIBS**: Used to specify a string of file names (comma separated string) to bind to the `/.singularity.d/libs` directory.

#### 21.1.4 D

1. **SINGULARITY\_DEFFILE**: Shows the Singularity recipe that was used to generate the image.
2. **SINGULARITY\_DESC**: Contains a description of the capabilities.
3. **SINGULARITY\_DETACHED**: To submit a build job and print the build ID (no real-time logs and also requires `--remote`). Default is set to false.
4. **SINGULARITY\_DNS**: A list of the DNS server addresses separated by commas to be added in `resolv.conf`.
5. **SINGULARITY\_DOCKER\_LOGIN**: To specify the interactive prompt for docker authentication.
6. **SINGULARITY\_DOCKER\_USERNAME**: To specify a username for docker authentication.
7. **SINGULARITY\_DOCKER\_PASSWORD**: To specify the password for docker authentication.
8. **SINGULARITY\_DROP\_CAPS**: To specify a list (comma separated string) of capabilities to be dropped. Default is an empty string.

#### 21.1.5 E

1. **SINGULARITY\_ENVIRONMENT**: Contains all the environment variables that have been exported in your container.
2. **SINGULARITYENV\_\***: Allows you to transpose variables into the container at runtime. You can see more in detail how to use this variable in our *environment and metadata section*.
3. **SINGULARITYENV\_APPEND\_PATH**: Used to append directories to the end of the `$PATH` environment variable. You can see more in detail on how to use this variable in our *environment and metadata section*.
4. **SINGULARITYENV\_PATH**: A specified path to override the `$PATH` environment variable within the container. You can see more in detail on how to use this variable in our *environment and metadata section*.
5. **SINGULARITYENV\_PREPEND\_PATH**: Used to prepend directories to the beginning of `$PATH` environment variable. You can see more in detail on how to use this variable in our *environment and metadata section*.

#### 21.1.6 F

1. **SINGULARITY\_FAKEROOT**: Set to false by default, considers running the container in a new user namespace as uid 0 (experimental).
2. **SINGULARITY\_FORCE**: Forces to kill the instance.

#### 21.1.7 G

1. **SINGULARITY\_GROUP**: Used to specify a string of capabilities for the given group.

### 21.1.8 H

1. **SINGULARITY\_HELPFILE**: Specifies the runscrip helpfile, if it exists.
2. **SINGULARITY\_HOME** : A home directory specification, it could be a source or destination path. The source path is the home directory outside the container and the destination overrides the home directory within the container.
3. **SINGULARITY\_HOSTNAME**: The container's hostname.

### 21.1.9 I

1. **SINGULARITY\_IMAGE**: Filename of the container.

### 21.1.10 J

1. **SINGULARITY\_JSON**: Specifies the structured json of the def file, every node as each section in the def file.

### 21.1.11 K

1. **SINGULARITY\_KEEP\_PRIVS**: To let root user keep privileges in the container. Default is set to false.

### 21.1.12 L

1. **SINGULARITY\_LABELS**: Specifies the labels associated with the image.
2. **SINGULARITY\_LIBRARY**: Specifies the library to pull from. Default is set to our Cloud Library.

### 21.1.13 N

1. **SINGULARITY\_NAME**: Specifies a custom image name.
2. **SINGULARITY\_NETWORK**: Used to specify a desired network. If more than one parameters is used, addresses should be separated by commas, where each network will bring up a dedicated interface inside the container.
3. **SINGULARITY\_NETWORK\_ARGS**: To specify the network arguments to pass to CNI plugins.
4. **SINGULARITY\_NOCLEANUP**: To not clean up the bundle after a failed build, this can be helpful for debugging. Default is set to false.
5. **SINGULARITY\_NOHTTPS**: Sets to either false or true to avoid using HTTPS for communicating with the local docker registry. Default is set to false.
6. **SINGULARITY\_NO\_HOME**: Considers not mounting users home directory if home is not the current working directory. Default is set to false.
7. **SINGULARITY\_NO\_INIT** and **SINGULARITY\_NOSHIMINIT**: Considers not starting the shim process with `--pid`.
8. **SINGULARITY\_NO\_NV**: Flag to disable Nvidia support. Opposite of **SINGULARITY\_NV**.
9. **SINGULARITY\_NO\_PRIVS**: To drop all the privileges from root user in the container. Default is set to false.
10. **SINGULARITY\_NV**: To enable experimental Nvidia support. Default is set to false.

### 21.1.14 o

1. **SINGULARITY\_OVERLAY** and **SINGULARITY\_OVERLAYIMAGE**: To indicate the use of an overlay file system image for persistent data storage or as read-only layer of container.

### 21.1.15 p

1. **SINGULARITY\_PWD** and **SINGULARITY\_TARGET\_PWD**: The initial working directory for payload process inside the container.

### 21.1.16 R

1. **SINGULARITY\_REMOTE**: To build an image remotely. (Does not require root) Default is set to false.
2. **SINGULARITY\_ROOTFS**: To reference the system file location.
3. **SINGULARITY\_RUNSCRIPT**: Specifies the runscript of the image.

### 21.1.17 s

1. **SINGULARITY\_SANDBOX**: To specify that the format of the image should be a sandbox. Default is set to false.
2. **SINGULARITY\_SCRATCH** and **SINGULARITY\_SCRATCHDIR**: Used to include a scratch directory within the container that is linked to a temporary directory. (use `-W` to force location)
3. **SINGULARITY\_SECTION**: To specify a comma separated string of all the sections to be run from the deffile (setup, post, files, environment, test, labels, none)
4. **SINGULARITY\_SECURITY**: Used to enable security features. (SELinux, Apparmor, Seccomp)
5. **SINGULARITY\_SECRET**: Lists all the private keys instead of the default which display the public ones.
6. **SINGULARITY\_SHELL**: The path to the program to be used as an interactive shell.
7. **SINGULARITY\_SIGNAL**: Specifies a signal sent to the instance.

### 21.1.18 T

1. **SINGULARITY\_TEST**: Specifies the test script for the image.
2. **SINGULARITY\_TMPDIR**: Used with the `build` command, to consider a temporary location for the build.

### 21.1.19 U

1. **SINGULARITY\_UNSHARE\_PID**: To specify that the container will run in a new PID namespace. Default is set to false.
2. **SINGULARITY\_UNSHARE\_IPC**: To specify that the container will run in a new IPC namespace. Default is set to false.
3. **SINGULARITY\_UNSHARE\_NET**: To specify that the container will run in a new network namespace (sets up a bridge network interface by default). Default is set to false.
4. **SINGULARITY\_UNSHARE\_UTS**: To specify that the container will run in a new UTS namespace. Default is set to false.



5. **SINGULARITY\_UPDATE**: To run the definition over an existing container (skips the header). Default is set to false.
6. **SINGULARITY\_URL**: Specifies the key server URL.
7. **SINGULARITY\_USER**: Used to specify a string of capabilities for the given user.
8. **SINGULARITY\_USERNS** and **SINGULARITY\_UNSHARE\_USERNS**: To specify that the container will run in a new user namespace, allowing Singularity to run completely unprivileged on recent kernels. This may not support every feature of Singularity. (Sandbox image only). Default is set to false.

### 21.1.20 w

1. **SINGULARITY\_WORKDIR**: The working directory to be used for /tmp, /var/tmp and \$HOME (if -c or --contain was also used)
2. **SINGULARITY\_WRITABLE**: By default, all Singularity containers are available as read only, this option makes the file system accessible as read/write. Default set to false.
3. **SINGULARITY\_WRITABLE\_TMPFS**: Makes the file system accessible as read-write with non-persistent data (with overlay support only). Default is set to false.

## 21.2 Build Modules

### 21.2.1 library bootstrap agent

#### 21.2.1.1 Overview

You can use an existing container on the Container Library as your “base,” and then add customization. This allows you to build multiple images from the same starting point. For example, you may want to build several containers with the same custom python installation, the same custom compiler toolchain, or the same base MPI installation. Instead of building these from scratch each time, you could create a base container on the Container Library and then build new containers from that existing base container adding customizations in %post, %environment, %runscript, etc.

#### 21.2.1.2 Keywords

```
Bootstrap: library
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
From: <entity>/<collection>/<container>:<tag>
```

The From keyword is mandatory. It specifies the container to use as a base. `entity` is optional and defaults to `library`. `collection` is optional and defaults to `default`. This is the correct namespace to use for some official containers (`alpine` for example). `tag` is also optional and will default to `latest`.

```
Library: http://custom/library
```

The Library keyword is optional. It will default to `https://library.sylabs.io`.

## 21.2.2 docker bootstrap agent

### 21.2.2.1 Overview

Docker images are comprised of layers that are assembled at runtime to create an image. You can use Docker layers to create a base image, and then add your own custom software. For example, you might use Docker's Ubuntu image layers to create an Ubuntu Singularity container. You could do the same with CentOS, Debian, Arch, Suse, Alpine, BusyBox, etc.

Or maybe you want a container that already has software installed. For instance, maybe you want to build a container that uses CUDA and cuDNN to leverage the GPU, but you don't want to install from scratch. You can start with one of the `nvidia/cuda` containers and install your software on top of that.

Or perhaps you have already invested in Docker and created your own Docker containers. If so, you can seamlessly convert them to Singularity with the `docker bootstrap` module.

### 21.2.2.2 Keywords

```
Bootstrap: docker
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
From: <registry>/<namespace>/<container>:<tag>@<digest>
```

The From keyword is mandatory. It specifies the container to use as a base. `registry` is optional and defaults to `index.docker.io`. `namespace` is optional and defaults to `library`. This is the correct namespace to use for some official containers (ubuntu for example). `tag` is also optional and will default to `latest`

See *Singularity and Docker* for more detailed info on using Docker registries.

```
Registry: http://custom_registry
```

The Registry keyword is optional. It will default to `index.docker.io`.

```
Namespace: namespace
```

The Namespace keyword is optional. It will default to `library`.

```
IncludeCmd: yes
```

The IncludeCmd keyword is optional. If included, and if a `%runscript` is not specified, a Docker CMD will take precedence over ENTRYPOINT and will be used as a runscript. Note that the IncludeCmd keyword is considered valid if it is not empty! This means that `IncludeCmd: yes` and `IncludeCmd: no` are identical. In both cases the IncludeCmd keyword is not empty, so the Docker CMD will take precedence over an ENTRYPOINT.

See *Singularity and Docker* for more info on order of operations for determining a runscript.

### 21.2.2.3 Notes

Docker containers are stored as a collection of tarballs called layers. When building from a Docker container the layers must be downloaded and then assembled in the proper order to produce a viable file system. Then the file system must be converted to Singularity Image File (sif) format.

Building from Docker Hub is not considered reproducible because if any of the layers of the image are changed, the container will change. If reproducibility is important to your workflow, consider hosting a base container on the Container Library and building from it instead.

For detailed information about setting your build environment see *Build Customization*.

## 21.2.3 shub bootstrap agent

### 21.2.3.1 Overview

You can use an existing container on Singularity Hub as your “base,” and then add customization. This allows you to build multiple images from the same starting point. For example, you may want to build several containers with the same custom python installation, the same custom compiler toolchain, or the same base MPI installation. Instead of building these from scratch each time, you could create a base container on Singularity Hub and then build new containers from that existing base container adding customizations in `%post`, `%environment`, `%runscript`, etc.

### 21.2.3.2 Keywords

```
Bootstrap: shub
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
From: shub://<registry>/<username>/<container-name>:<tag>@digest
```

The From keyword is mandatory. It specifies the container to use as a base. `registry` is optional and defaults to `singularity-hub.org`. `tag` and `digest` are also optional. `tag` defaults to `latest` and `digest` can be left blank if you want the latest build.

### 21.2.3.3 Notes

When bootstrapping from a Singularity Hub image, all previous definition files that led to the creation of the current image will be stored in a directory within the container called `/.singularity.d/bootstrap_history`. Singularity will also alert you if environment variables have been changed between the base image and the new image during bootstrap.

## 21.2.4 localimage bootstrap agent

This module allows you to build a container from an existing Singularity container on your host system. The name is somewhat misleading because your container can be in either image or directory format.

### 21.2.4.1 Overview

You can use an existing container image as your “base”, and then add customization. This allows you to build multiple images from the same starting point. For example, you may want to build several containers with the same custom python installation, the same custom compiler toolchain, or the same base MPI installation. Instead of building these from scratch each time, you could start with the appropriate local base container and then customize the new container in `%post`, `%environment`, `%runscript`, etc.

### 21.2.4.2 Keywords

```
Bootstrap: localimage
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
From: /path/to/container/file/or/directory
```

The From keyword is mandatory. It specifies the local container to use as a base.

### 21.2.4.3 Notes

When building from a local container, all previous definition files that led to the creation of the current container will be stored in a directory within the container called `/.singularity.d/bootstrap_history`. Singularity will also alert you if environment variables have been changed between the base image and the new image during bootstrap.

## 21.2.5 yum bootstrap agent

This module allows you to build a Red Hat/CentOS/Scientific Linux style container from a mirror URI.

### 21.2.5.1 Overview

Use the `yum` module to specify a base for a CentOS-like container. You must also specify the URI for the mirror you would like to use.

### 21.2.5.2 Keywords

```
Bootstrap: yum
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
OSVersion: 7
```

The OSVersion keyword is optional. It specifies the OS version you would like to use. It is only required if you have specified a `%{OSVERSION}` variable in the `MirrorURL` keyword.

```
MirrorURL: http://mirror.centos.org/centos-%{OSVERSION}/%{OSVERSION}/os/$basearch/
```

The MirrorURL keyword is mandatory. It specifies the URI to use as a mirror to download the OS. If you define the OSVersion keyword, then you can use it in the URI as in the example above.

```
Include: yum
```

The Include keyword is optional. It allows you to install additional packages into the core operating system. It is a best practice to supply only the bare essentials such that the `%post` section has what it needs to properly complete the build. One common package you may want to install when using the `yum` build module is YUM itself.

### 21.2.5.3 Notes

There is a major limitation with using YUM to bootstrap a container. The RPM database that exists within the container will be created using the RPM library and Berkeley DB implementation that exists on the host system. If the RPM implementation inside the container is not compatible with the RPM database that was used to create the container, RPM and YUM commands inside the container may fail. This issue can be easily demonstrated by bootstrapping an older RHEL compatible image by a newer one (e.g. bootstrap a Centos 5 or 6 container from a Centos 7 host).

In order to use the `yum` build module, you must have `yum` installed on your system. It may seem counter-intuitive to install YUM on a system that uses a different package manager, but you can do so. For instance, on Ubuntu you can install it like so:

```
$ sudo apt-get update && sudo apt-get install yum
```

## 21.2.6 debootstrap build agent

This module allows you to build a Debian/Ubuntu style container from a mirror URI.

### 21.2.6.1 Overview

Use the `debootstrap` module to specify a base for a Debian-like container. You must also specify the OS version and a URI for the mirror you would like to use.

### 21.2.6.2 Keywords

```
Bootstrap: debootstrap
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
OSVersion: xenial
```

The OSVersion keyword is mandatory. It specifies the OS version you would like to use. For Ubuntu you can use code words like `trusty` (14.04), `xenial` (16.04), and `yakkety` (17.04). For Debian you can use values like `stable`, `oldstable`, `testing`, and `unstable` or code words like `wheezy` (7), `jesse` (8), and `stretch` (9).

```
MirrorURL: http://us.archive.ubuntu.com/ubuntu/
```

The MirrorURL keyword is mandatory. It specifies a URI to use as a mirror when downloading the OS.

```
Include: somepackage
```

The Include keyword is optional. It allows you to install additional packages into the core operating system. It is a best practice to supply only the bare essentials such that the `%post` section has what it needs to properly complete the build.

### 21.2.6.3 Notes

In order to use the `debootstrap` build module, you must have `debootstrap` installed on your system. On Ubuntu you can install it like so:

```
$ sudo apt-get update && sudo apt-get install debootstrap
```

On CentOS you can install it from the epel repos like so:

```
$ sudo yum update && sudo yum install epel-release && sudo yum install debootstrap.  
↪noarch
```

## 21.2.7 arch bootstrap agent

This module allows you to build a Arch Linux based container.

### 21.2.7.1 Overview

Use the `arch` module to specify a base for an Arch Linux based container. Arch Linux uses the aptly named `pacman` package manager (all puns intended).

### 21.2.7.2 Keywords

```
Bootstrap: arch
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

The Arch Linux bootstrap module does not name any additional keywords at this time. By defining the `arch` module, you have essentially given all of the information necessary for that particular bootstrap module to build a core operating system.

### 21.2.7.3 Notes

Arch Linux is, by design, a very stripped down, light-weight OS. You may need to perform a significant amount of configuration to get a usable OS. Please refer to this [README.md](#) and the [Arch Linux example](#) for more info.

## 21.2.8 busybox bootstrap agent

This module allows you to build a container based on BusyBox.

### 21.2.8.1 Overview

Use the `busybox` module to specify a BusyBox base for container. You must also specify a URI for the mirror you would like to use.

### 21.2.8.2 Keywords

```
Bootstrap: busybox
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
MirrorURL: https://www.busybox.net/downloads/binaries/1.26.1-defconfig-multiarch/  
↳busybox-x86_64
```

The MirrorURL keyword is mandatory. It specifies a URI to use as a mirror when downloading the OS.

### 21.2.8.3 Notes

You can build a fully functional BusyBox container that only takes up ~600kB of disk space!

## 21.2.9 zypper bootstrap agent

This module allows you to build a Suse style container from a mirror URI.

### 21.2.9.1 Overview

Use the `zypper` module to specify a base for a Suse-like container. You must also specify a URI for the mirror you would like to use.

### 21.2.9.2 Keywords

```
Bootstrap: zypper
```

The Bootstrap keyword is always mandatory. It describes the bootstrap module to use.

```
OSVersion: 42.2
```

The OSVersion keyword is optional. It specifies the OS version you would like to use. It is only required if you have specified a `%{OSVERSION}` variable in the `MirrorURL` keyword.

```
Include: somepackage
```

The Include keyword is optional. It allows you to install additional packages into the core operating system. It is a best practice to supply only the bare essentials such that the `%post` section has what it needs to properly complete the build. One common package you may want to install when using the `zypper` build module is `zypper` itself.