

WenQuanYi Micro Hei [Scale=0.9]WenQuanYi Micro Hei Mono song-
WenQuanYi Micro Hei sfWenQuanYi Micro Hei "zh" = 0pt plus 1pt

**KaliArch's
awesome-kubernetes-notes**
Release 1.0.0

Dec 22, 2019

1	äÿĂ KuberneſesæċĈēř	9
1.1	1.1 āōzāŽīçijŮæŎšāũēāĚũ	9
1.2	1.2 kuberneſes	9
1.3	1.3 ċŎřāċĈæđũæđĎ	10
1.4	1.4 æđũæđĎāŠŇċzĎžzũ	12
2	ăžŇ æăÿăŧĈċzĎžzũ/ėŽĎžzũ	13
2.1	2.1 Controller	13
2.2	2.2 Service	14
2.3	2.3 ċ;ŠċzIJæĺăđŇ	14
2.4	2.4 kube-proxy	15
2.5	2.5 etcd	15
2.6	2.6 flanel	15
2.7	ċșċĕřĖăřŔċzȘ	16
3	ăÿL' éŽĖċ;đ'ėĈĺċ;ș	17
3.1	3.1 éĈĺċ;șăL'■ăĖĖăđ'Ė	17
3.2	3.2 éĈĺċ;ș Master	19
3.3	3.3 éĈĺċ;ș Node	22
4	ăžŽ āĚĕĕŮĺăŠ;ăžđ'	26
4.1	4.1 kubectl	26
4.2	4.2 run	28
4.3	4.3 expose	29
4.4	4.4 cp	29
4.5	4.5 port-forward	30
4.6	4.6 coredns	30
4.7	4.7 æĺăæŇș POD ĕċnăĹăėŽđ'	31
4.8	4.8 æĺăæŇș service ĕċnăĹăėŽđ'	31
4.9	4.9 labels	31
4.10	4.10 âĹăĂĂăĹ'ăōž	33
4.11	4.11 æžŽăĹăă■Ėċžġ	33
4.12	4.12 éŽĖċ;đ'ăđ'ŮĕŏĕéŮŏ	35

4.13	4.13 æŒšæšæŸæŁæŁŸ	35
4.14	4.14 èŁđăĚĚ POD áőžáŽÍ	35
5	äžŦ éĚ■ç;őæÿĖăŦăŦăçŦí	36
5.1	5.1 âŔréĚ■ç;őçŽĐâřžèşą	36
5.2	5.2 éĚ■ç;őæÿĖăŦăŦčžĐæĹŔ	36
5.3	5.3 èŬăŔŮŰæÿĖăŦăŦăŸőăĹ'	37
5.4	5.4 æÿĖăŦăŦăşžæĹĴăăĵăĵŔ	37
5.5	5.5 âĤŋăēŮèŬăŔŮŰæÿĖăŦăŦ	38
5.6	5.6 create âĹŽăžž	38
5.7	5.7 delete âĹăéŽď'	38
5.8	5.8 apply âĹŽăžžæĹŮæŽť æŮř	39
6	ăĚ■ POD éĚ■ç;őæÿĖăŦăŦ	40
6.1	6.1 pods.metadata PODăĚČæŦřăēő	40
6.2	6.2 pods.spec èğĐěŇČ	40
7	äÿČ æŐğăĹŮăŽíéĚ■ç;őæÿĖăŦăŦ	51
7.1	7.1 ReplicaSet æŐğăĹŮăŽÍ	51
7.2	7.2 DeploymentæŐğăĹŮăŽÍ	52
7.3	7.3 DaemonSetæŐğăĹŮăŽÍ	55
8	ăĚŋ Service éĚ■ç;őæÿĖăŦăŦ	58
8.1	8.1 Service âŮëăĴăăĵăŔ	58
8.2	8.2 Service çşžăđŇ	58
8.3	8.3 èťĐæžŔèőřăĴ	59
8.4	8.4 Service æÿĖăŦăŦ	59
8.5	8.5 service.spec èğĐěŇČ	59
8.6	8.6 ClusterIP çşžăđŇçŽĐ service	59
8.7	8.7 NodePort çşžăđŇçŽĐ service	60
8.8	8.8 loadBalancerIP çşžăđŇ	60
8.9	8.9 æŮăéŽĚçĴď' âĴřăĹĂçŽĐ Service	61
8.10	8.10 externalName çşžăđŇ	61
9	äži ingress æŐğăĹŮăŽÍ	62
9.1	9.1 ingress.spec èğĐěŇČ	63
9.2	9.2 ingress-nginx äžčçŘĚ	63
9.3	9.3 ingress-tomcat äžčçŘĚ	65
10	ă■A POD âŮăĆăĹăŮ	68
10.1	10.1 âŮüçŽĐçşžăđŇ	68
10.2	10.2 áőžăŽíăŇČèĴĲéĂĹéąż	69
10.3	10.3 èĹČçÇzăŮăĆí	69
10.4	10.4 çĴŞçžĴăŮăĆí	72
10.5	10.5 âĹĚăÿČăĵŔăŮăĆí	73
10.6	10.6 StorageClass Ceph RBD	77
11	ă■AäÿĂ éĚ■ç;őăĤæAŕăőžăŽíăŇŮ	81
11.1	11.1 POD èŬăŔŮŰçŐŕăćČăŔŸéĞŔ	82

11.2	11.2 configMap	82
11.3	11.3 secret	86
12	ā■AāžŃ StatefulSet æŌğĀLúāZÍ	90
12.1	12.1 æyĒā■TæāijāijR	91
12.2	12.2 āLŽāžž NFS PV	91
12.3	12.3 āLŽāžž statefulSet	93
12.4	12.4 æLl'āōžāŠŃā■Ğčžğ	94
13	ā■AäyL ċTlæLúèød'èrAçşçzçş	96
13.1	13.1 ċTlæLúçŽDçşzādŃ	96
13.2	13.2 PODæÇäi;TēfđæŌēéZĖç;d'	97
13.3	13.3 serviceaccount āřžèşā	97
13.4	13.4 kubectl éĚ■ç;őæŰĞäzű	98
13.5	13.5 æūzāLæèfAäžççTlæLúāLř config	99
13.6	13.6 āLŽāžžæŰř config æŰĞäzű	101
13.7	13.7 āşžžŌ token èød'èrA	101
14	ā■AāžŽ ċTlæLúæİČéŽRçşçzçş	103
14.1	14.1 æİČéŽŘāLŰèāĹ	104
14.2	14.2 āLŽāžž Role	104
14.3	14.3 āLŽāžž rolebinding	105
14.4	14.4 āLŽāžž clusterrole	105
14.5	14.5 āLŽāžž clusterrolebinding	107
14.6	14.6 rolebinding äyŌ clusterrole	107
14.7	14.7 RBACæŌLæİČ	108
15	ā■AäžT dashboard	110
15.1	15.1 éČíç;şætAçĹŃ	110
15.2	15.2 äj;ççTlāzd'çL'ŃçŽžāi;T	111
15.3	15.3 āLEçžğçőaçŘĖ	112
15.4	15.4 éĚ■ç;őæŰĞäzűèød'èrA	112
16	ā■AāĚ■ ç;ŚçzIJéĀŽāfa	114
16.1	16.1 éĀŽāfaæĹāādŃ	114
16.2	16.2 éĀŽāfaæĹāādŃāžTāsČ	114
16.3	16.3 K8S āŘ■çğřçl'žéŮř	115
16.4	16.4 K8Sç;ŚçzIJæŃŞæL'S	115
16.5	16.5 flannel	116
16.6	16.6 Calico	118
17	ā■AäyČ èřČāžçç■ŰçTě	123
17.1	17.1 PODāLŽāžžætAçĹŃ	123
17.2	17.2 ServiceāLŽāžžèfĞçĹŃ	124
17.3	17.3 èřDæžŘéŽŘāLúçzt'āžç	124
17.4	17.4 Scheduler èřČāžçèfĞçĹŃ	124
17.5	17.4 écĎéĀL'āŽāçř'ā	125
17.6	17.5 äijYéĀL'āĞjæTř	127
17.7	17.6 éĀL'æŃl'āĞjæTř	128

18	ā■AāĒn éñŸčžgèrČāžēēō;ç;ĩ	129
18.1	18.1 èĽĆçĆzéĀĽ'æŃĲāZĪ	129
18.2	18.2 áržèĽĆçĆžçŽDžšāŠŃæĀğ	130
18.3	18.3 árž POD çŽDžšāŠŃæĀğ	132
18.4	18.4 árž POD çŽDāR■āžšāŠŃæĀğ	134
18.5	18.5 node æśaçĆž	136
18.6	18.6 POD æśaçĆžāōžāf■	138
19	ā■AāZĪ āōžāZĪlètDæžŘéZŘāLŮ	141
19.1	19.1 ètDæžŘéZŘāLŮ	141
19.2	19.2 qos èt'léĜŘçōaçŘĚ	143
20	āžŃā■A HeapSterçŽŚæŌgüjĽāžšāijČäy■üjĽ	144
20.1	20.1 āōĽ'èčĚ influx DB	144
20.2	20.2 āōĽ'èčĚ HeapSter	146
20.3	20.3 āōĽ'èčĚ Grafana	148
21	āžŃā■AäyĀ æŮrāyĀžčçŽŚæŌgæđūædĎ	150
21.1	21.1 æäyāfČæŃĜæāĜætĀært'čžf	150
21.2	21.2çŽŚæŌgætĀært'čžf	150
21.3	21.3 āōĽ'èčĚ metrics-server	151
21.4	21.4 āōĽ'èčĚ prometheus	152
21.5	21.5 HPAāŚ;āzd'èāŃæŮžāijŘ	152
21.6	21.6 HPAæyĚā■T	153
22	āžŃā■AāžŃ K8SāŃĚçōaçŘĚāZĪ	155
22.1	22.1 āšžçāĀæçČāft	155
22.2	22.2 Helm āūēā;IJāŌšçŘĚ	156
22.3	22.3 éČĪç;š Helm	157
22.4	22.4 ChartæŮĜāžūçžĎçžĜ	159
22.5	22.5 ä;fçTĪ Helm + Ceph éČĪç;š EFK	159
22.6	22.6 Storage Class	160
22.7	22.7 Helm Elasticsearch	161
22.8	22.8 Helm fluentd-elasticsearch	163
22.9	22.9 Helm kibana	166
23	āžŃā■AäyĽ ETCDèrēègč	167
23.1	23.1 ETCDæçCèfř	167
23.2	23.2 ETCDæđūædĎāRĽègčædŘ	169
23.3	23.3 āžTçTĪāIJæŽř	170
23.4	23.4 āōĽ'èčĚéČĪç;š	174
23.5	23.5 çōĀā■Tä;fçTĪ	181
23.6	23.6 çd'žä;Ń	185
24	āžŃā■AāZŽ āZ;āžgāōžāZĪçōaçŘĚāžšāRřKubeSphereāōđæĽYæŌŠéTŽ	190
24.1	24.1 æyĚçŘĚéĀĀāĜçĽūæĀAçŽĎāōžāZĪ	190
24.2	24.2 æyĚçŘĚāijČäyāĽĽūécénél' séĀŘçŽĎ pod	190
24.3	24.3 Docker æTřæ■ōēfAçgž	192
24.4	24.4 kubosphere ç;ŚçzIJæŌŠéTŽ	192

24.5	24.5 kubescape ǎŹŤčŤlèŭčŤśǎijĆǎŷŷ	194
24.6	24.6 Jenkins çŽĎ Agent	195
24.7	24.7 Devops äŷ■ MailçŽĎǎŔŚéǎǎ	197
24.8	ǎŔĆèǎŤĆéŞ;æŎě	199
25	Indices and tables	201

Contents

- *äyǺ KuberñetesæġĆèĚř*
 - *1.1 âôžâZĹċijŮæŎŠâũěâĚũ*
 - *1.2 kuberñetes*
 - *1.3 ċŎřâċČæđũæđĎ*
 - *1.4 æđũæđĎâŠŇċžĎäzũ*
- *äžŇ æäyâĚČċžĎäzũ/éŽĎäzũ*
 - *2.1 Controller*
 - *2.2 Service*
 - *2.3 ċĴŠċžIJæĹqâđŇ*
 - *2.4 kube-proxy*
 - *2.5 etcd*
 - *2.6 flanel*
 - *ċšěërĒârĚċžŚ*
- *äyĹ éŽĒċžĎ' éČĹċž*
 - *3.1 éČĹċžšâĹ■âĢĒâđ'Ģ*
 - * *3.1.1 âĚšéŮ■ firewalld âŠŇ selinux*
 - * *3.1.2 âĹăèĴĴ ipvs âĒĚæäyæĹqâĪŮ*
 - * *3.1.3 äyŇèĴĴ Docker âŠŇ K8S*
 - * *3.1.4 èôĴċžôâĒĚæäyâĹĹ K8S âĹĆæĤř*
 - *3.2 éČĹċžš Master*
 - * *3.2.1 æĹŘâĹ■æŇĹâĹŮéĤIJâČĹ*
 - * *3.2.2 âĹĹăġŇâŇŮMaster*
 - *3.3 éČĹċžš Node*
 - * *3.3.1 âĹăâĚěéŽĒċžĎ'*
 - * *3.3.2 æšċIJŇèĚŽâžġ*
 - * *3.3.3 éĤIJâČĹäyŇèĴĴâđ'ĹæĚĆ*
- *âŽŽ âĚěéŮĹâŚĴäzđ'*
 - *4.1 kubectĹ*
 - *4.2 run*
 - *4.3 expose*

- 4.4 cp
- 4.5 port-forward
- 4.6 coredns
- 4.7 *aelqaeNš POD ècñâLăéŽd'*
- 4.8 *aelqaeNš service ècñâLăéŽd'*
- 4.9 labels
- 4.10 *âLlâĀAæLl'âôž*
- 4.11 *æžŽâLlâ■Ğçžğ*
- 4.12 *éŽEç;d'âd'ŮèôŁéŮô*
- 4.13 *æŮŠæšæŮëâŁŮ*
- 4.14 *èŁđâĚě POD âôžâŽl*
- *äžŤ éĚ■ç;őæÿĚâ■Ťä;ŁçŤl*
 - 5.1 *âRřéĚ■ç;őçŽĎáržèšq*
 - 5.2 *éĚ■ç;őæÿĚâ■ŤçžĎæĹŔ*
 - 5.3 *èŮâRŮŮæÿĚâ■ŤäÿôâLl'*
 - 5.4 *æÿĚâ■ŤâšžæIJñæäijâijŔ*
 - 5.5 *âŁnáæ■ûèŮâRŮŮæÿĚâ■Ť*
 - 5.6 *create âĹŽâžž*
 - 5.7 *delete âĹăéŽd'*
 - 5.8 *apply âĹŽâžžæĹŮæŽt'æŮř*
- *âĚ■ POD éĚ■ç;őæÿĚâ■Ť*
 - 6.1 *Pods.metadata PODâĚĚČæŤřæ■ő*
 - * 6.1.1 *labels æăĞç■;*
 - 6.2 *Pods.spec èğĎèŇČ*
 - * 6.2.1 *nodeName èŁŘèqŇèŁĆçĆž*
 - * 6.2.2 *nodeSelector èŁĆçĆžéĀLæŇl'*
 - * 6.2.3 *restartPolicy PODéĞ■âŘřç■ŮçŤě*
 - * 6.2.4 *hostNetwork äÿžæIJžç;šçžIJçl'žéŮt'*
 - * 6.2.5 *hostPID äÿžæIJžPIDçl'žéŮt'*
 - * 6.2.6 *containers éĚ■ç;ő*
- *äÿČ æŮğâĹûâŽléĚ■ç;őæÿĚâ■Ť*
 - 7.1 *ReplicaSet æŮğâĹûâŽl*

- * 7.1.1 *replicaset.spec* èĝĎěŇĚ
- * 7.1.2 æÿĚā■Ŧçd'žä;Ň
- 7.2 *Deployment*æŎĝāLúāZÍ
 - * 7.2.1 *replicaset.spec* âřžèšqèĝĎěŇĚ
 - * 7.2.2 æÿĚā■Ŧçd'žä;Ň
 - * 7.2.3 âĚšäžŎæŽt'æŮř
 - * 7.2.4 æÍqæŇšéĜSäÿIéŽĀâŔSâÿĚ
 - * 7.2.5 æŽt'æŮřç■ŮçŦě
 - * 7.2.6 âĚšäžŎâŽđæžŽ
- 7.3 *DaemonSet*æŎĝāLúāZÍ
 - * 7.3.1 *DaemonSet.spec*èĝĎěŇĚ
 - * 7.3.2 æÿĚā■Ŧçd'žä;Ň
 - * 7.3.3 âĚšäžŎæŽt'æŮř
- âĚŇ *Service* éĚ■ç;ŏæÿĚā■Ŧ
 - 8.1 *Service* âũěä;IJæÍqâijŔ
 - 8.2 *Service* çšžâđŇ
 - 8.3 èŦĎæžŔèŏřâ;Ŧ
 - 8.4 *Service* æÿĚā■Ŧ
 - 8.5 *service.spec* èĝĎěŇĚ
 - 8.6 *ClusterIP* çšžâđŇçŽĎ *service*
 - 8.7 *NodePort* çšžâđŇçŽĎ *service*
 - 8.8 *loadBalancerIP* çšžâđŇ
 - 8.9 æŮăéŽĚç; d'âIJřâĬçŽĎ *Service*
 - 8.10 *externalName* çšžâđŇ
- äžĬ *ingress* æŎĝāLúāZÍ
 - 9.1 *ingress.spec* èĝĎěŇĚ
 - 9.2 *ingress-nginx* äžčçŔĚ
 - 9.3 *ingress-tomcat* äžčçŔĚ
- â■A *POD* â■ŸâĆÍâ■ũ
 - 10.1 â■ũçŽĎçšžâđŇ
 - 10.2 âŏžâŽÍæŇĚç; j;éĂĽéqž
 - 10.3 èĽĆçĆžâ■ŸâĆÍ

- * 10.3.1 *hostpath*
- * 10.3.2 *gitRepo*
- * 10.3.3 *emptyDir*
- 10.4 *ç;ŚçzIJā■ŸāĆí■ũ*
 - * 10.4.1 *nfs*
- 10.5 *āŁĚäÿČäijŘā■ŸāĆí*
 - * 10.5.1 *PersistentVolume*
 - * 10.5.2. *PersistentVolumeClaim*
 - * 10.5.3 *StorageClass*
- 10.6 *StorageClass Ceph RBD*
 - * 10.6.1 *éĚ■ç;ř Ceph āĆíā■Ÿæšā*
 - * 10.6.2 *āōLèčĚ rbd-provisioner*
 - * 10.6.3 *ä;ŁçŤl StorageClass*
- *ā■AäÿĀ éĚ■ç;řāŁæAřāōžāŽlāŇŮ*
 - 11.1 *POD èŌūāRŮçŎřāćČāRŸéĜŘ*
 - 11.2 *configMap*
 - * 11.2.1 *æšlāĚĚ POD ENV*
 - * 11.2.2 *æŇČè;ĵäÿž POD ā■ũ*
 - 11.3 *secret*
 - * 11.3.1 *çğAæIJL'üzšāžšèōd'èřA1*
 - * 11.3.2 *çğAæIJL'üzšāžšèōd'èřA2*
 - * 11.3.3 *āŁŽāžž TLS èřAžę*
- *ā■AžŇ StatefulSet æŎğāŁūāŽl*
 - 12.1 *æÿĚā■ŤæäijäijŘ*
 - 12.2 *āŁŽāžž NFS PV*
 - 12.3 *āŁŽāžž statefulSet*
 - 12.4 *æL'l'āōžāšŇā■Ğçžğ*
- *ā■AäÿL çŤlæLūèōd'èřAçšžçžš*
 - 13.1 *çŤlæLūçŽĐçšžāđŇ*
 - 13.2 *PODāęĆü;ŤèŁđæŎěéŽEç;d'*
 - 13.3 *serviceaccount āřžèšq*
 - * 13.3.1 *āIJl POD äÿ■ü;ŁçŤl serviceaccount*

- 13.4 kubectľ éĚĲ;őæŮĜäzű
- 13.5 æűzâĹăërAäzëçŤĹæĹûâĹř config
 - * 13.5.1 âĹZâzžSSLèrAäzëçŤĹæĹû
 - * 13.5.2 æűzâĹăSSLèrAäzëçŤĹæĹûâĹřconfig
 - * 13.5.3 âĹZâzžâĹĜæĲcäyĹäyŇæŮĜ
- 13.6 âĹZâzžæŮř config æŮĜäzű
- 13.7 âšžäžŎ token èöd'èrA
 - * 13.7.1 âĹZâzž serviceaccount
 - * 13.7.2 çžŠâöŽéZĚç;đ'çőqçŘĚâŠŸèğŠèĹš
 - * 13.7.3 éĂŽèĚĜ serviceaccount â;ŮâĹř Token
- âĲAâZŽ çŤĹæĹûæĲĚŽŘçžçžš
 - 14.1 æĲĚŽŘâĹŮèqĲ
 - 14.2 âĹZâzž Role
 - 14.3 âĹZâzž rolebinding
 - 14.4 âĹZâzž clusterrole
 - 14.5 âĹZâzž clusterrolebinding
 - 14.6 rolebinding äyŎ clusterrole
 - 14.7 RBACæŎĹæĲĚ
- âĲAäžŤ dashboard
 - 15.1 éĲĲ;šæĲAçĲŇ
 - 15.2 ä;ĚçŤĹäzd'çĹŇçŽzâ;Ť
 - 15.3 âĹĚçžğçőqçŘĚ
 - 15.4 éĚĲ;őæŮĜäzűèöd'èrA
- âĲAâĚĲ;šçžIJéĂŽăĚq
 - 16.1 éĂŽăĚqæĲqâđŇ
 - 16.2 éĂŽăĚqæĲqâđŇăžŤâšĲ
 - 16.3 K8S âŘĲğřçĲ'žéŮř
 - 16.4 K8Sç;šçžIJæŇšæĹŠ
 - 16.5 flannel
 - * 16.5.1 flannel âűěä;IJæĲqâijŘ
 - * 16.5.2 VXLAN éĂŽăĚqèĚĜçĲŇ
 - * 16.5.3 flannel éĲĲ;šæŮžâijŘ

- * 16.5.4 flannel éĚ■çjőæÚĜäzű
- * 16.5.5 äŁőæŤžâũěü;IJæĺqâijŘ
- 16.6 Calico
 - * 16.6.1 âóL'èčĚ canal
 - * 16.6.2 æyĚâ■ŤâóŽäzL'
 - * 16.6.3 policyTypes
- â■AäyČ èřČäžęç■ŮçTě
 - 17.1 PODâĹZâžžæŤAçĹŇ
 - 17.2 ServiceâĹZâžžèŁĜçĹŇ
 - 17.3 èŤDæžŘéŽŘâĹúçžt'âžę
 - 17.4 Scheduler èřČäžęèŁĜçĹŇ
 - 17.4 éčĎéĀLâŽăçt'ă
 - 17.5 äijYéĀLâĜjæŤř
 - 17.6 éĀLæŇl'âĜjæŤř
- â■AâĚŇ éŋYçžğèřČäžęèőçjő
 - 18.1 èŁĆçĆžéĀLæŇl'âŽĹ
 - 18.2 âřžèŁĆçĆžçŽĎäžšâŠŇæĀğ
 - 18.3 âřž POD çŽĎäžšâŠŇæĀğ
 - 18.4 âřž POD çŽĎâR■äžšâŠŇæĀğ
 - 18.5 node æsqçĆž
 - 18.6 POD æsqçĆžâőžâŁ■
- â■AäzĹ âőžâŽĹèŤDæžŘéŽŘâĹú
 - 19.1 èŤDæžŘéŽŘâĹú
 - 19.2 qos èŤléĜŘçőqçŘĚ
- äžŇâ■A HeapSterçŽŚæŎğüijĹäžšâijČäy■üijL
 - 20.1 âóL'èčĚ influx DB
 - 20.2 âóL'èčĚ HeapSter
 - 20.3 âóL'èčĚ Grafana
- äžŇâ■AäyĀ æŮřäyĀüzççŽŚæŎğæđúæđĎ
 - 21.1 æăyâŁČæŇĜæăĜæŤAæřt'çžŁ
 - 21.2 çŽŚæŎğæŤAæřt'çžŁ
 - 21.3 âóL'èčĚ metrics-server

- 21.4 *prometheus*
- 21.5 *HPA*
- 21.6 *HPA*
- *äžÑā■AäžÑ K8SāÑĚçõqçŘĚāŽÍ*
 - 22.1 *āšžçqĀæĈā£†*
 - 22.2 *Helm āũěä;IJāÕšçŘĚ*
 - 22.3 *éČĺç;š Helm*
 - 22.4 *ChartæŮĜäžŮčžĎčžĜ*
 - 22.5 *ä;£çŤĺ Helm + Ceph éČĺç;š EFK*
 - 22.6 *Storage Class*
 - 22.7 *Helm Elasticsearch*
 - 22.8 *Helm fluentd-elasticsearch*
 - 22.9 *Helm kibana*
- *äžÑā■AäyL ETCDèréèğč*
 - 23.1 *ETCDæĈè£ř*
 - * 23.1.1 *ETCDçõĀäžÑ*
 - * 23.1.2 *āŘŚāšŤāÕĚāŘš*
 - * 23.1.3 *ETCDçL'žçĆž*
 - * 23.1.4 *æĈā£†æIJřè■*
 - * 23.1.5 *çŽyāĚšāÕšçŘĚ*
 - 23.2 *ETCDæđŮæđĎāŘĽèğčæđŘ*
 - * 23.2.1 *æđŮæđĎāŽ;*
 - * 23.2.2 *æđŮæđĎèğčæđŘ*
 - 23.3 *āžŤçŤĺāIJžæŽř*
 - * 23.3.1 *æIJ■āŁæşĺāĚÑāŘŚçÕř*
 - * 23.3.2 *æŮĽæAřāŘŚāyČäyÕèócéŸĚ*
 - * 23.3.3 *èt'sè;jāĬĜèqq*
 - * 23.3.4 *āĽĚäyČäijRéĀŽçšäyŌā■RéřČ*
 - * 23.3.5 *āĽĚäyČäijRéŤĀ*
 - * 23.3.6 *āĽĚäyČäijRéŸšāĽŮ*
 - * 23.3.7 *éŽĚç;d'āŘĽçĽsä;jäyŌLeaderéĀĽäy;*
 - 23.4 *āōĽèčĚéČĺç;š*

- * 23.4.1 *â■TæIJžâôL'èčĚ*
- * 23.4.2 *éŽEç;d' éČlç;š*
- 23.5 *çóĀâ■Tä;£çTl*
 - * 23.5.1 *âćđâĽă*
 - * 23.5.2 *âĽăéŽd'*
 - * 23.5.3 *æŽt'æŮř*
 - * 23.5.4 *æ§ěèřć*
 - * 23.5.5 *watch*
 - * 23.5.6 *âd' Ğăzj*
 - * 23.5.7 *member*
- 23.6 *çd'žă;Ň*
- *ăžŇâ■AăŽŽ âŽj;ăžğăôžâŽlçóqçŘĚăžšăŘřKubeSphereăôđæĽŸæŎŠéTŽ*
 - 24.1 *æÿĚçŘĚéĀĀăĜžçĽŭæĀçŽĎăôžâŽl*
 - 24.2 *æÿĚçŘĚăijĆăÿÿæĽŮèćnél'séĀŘçŽĎ pod*
 - 24.3 *Docker æTřæ■óè£Açğž*
 - 24.4 *kubesphere ç;ŚçzIJæŎŠéTŽ*
 - 24.5 *kubesphere âžTçTl'èŭřçTšăijĆăÿÿ*
 - 24.6 *Jenkins çŽĎ Agent*
 - 24.7 *Devops äÿ■ MailçŽĎăŘŚéĀA*
 - *âŘĆèĀČéç;æŎě*
 - * *èğĚéćŚ*
 - * *èřŭæĽŚăŮlăŠŮăTqăŸTřÿŘ*

[TOC]

CHAPTER 1

äyÄ KubernetesæçĈèĖř

1.1 1.1 áóžáŽÍçijŮæŮŠăũěăĚů

- docker áóŸæŮžçijŮæŮŠăũěăĚů

```
docker compose      # á■ŤæIJžçijŮæŮŠăũěăĚů
docker swarm        # ářĚăd'ŽăŘř docker_
→æŘŘăžŽçŽĎěőaçőŮèťĎæžŘæŤt'ăŘĹçŽĎăőěăŘčĭijŇěŽŘăŘő docker compose_
→çijŮæŮŠçŽĎăŮăăŽăŘĹéIJĂèèĂéÍcăŘSèŁŽăŸĹæŤt'ăŘĹçŽĎăőěăŘčèŁŽèăŇçijŮæŮŠăřsèăŇĭijŇă
docker mechine      # ářĚăŸĂăŸĹăŸzæIJžăĹĹăġŇăŇŮăŸzăŸĂăŸĹèĈ;ăd'SăĹăăĚě_
→docker swarm éŽĚççd'ăŸ■çŽĎéćĎç;őçÍŇăžŘ
```

- mesos IDC æŞ■ă;IJçşžçzŞĭijŇěĈ;ăřĚăŸĂăŸĹ IDC_

```
IDC æŞ■ă;IJçşžçzŞĭijŇěĈ;ăřĚăŸĂăŸĹ IDC_
→æŘŘăžŽçŽĎçăñăzűèťĎæžŘĭijŇçzSăŸĂeřĈăžèăŞŇăĹĚéĖ■ĭijŇăőĈăŘĹăŸŕăŸĂăŸĹèťĎæžŘăĹĚéĖ■ă
```

- kubernetes

```
çŽőăĹ'■æIJĂæŤĂèăŇçŽĎăóžăŽÍçijŮæŮŠăũěăĚů,ăŸĈăIJžă■ăæIJĹ'çőĖæIJĂénŸ
```

1.2 1.2 kubernetes

kubernetes æŸŕăŸŇěĚĹèŕ■ĭijŇçŁžèŕSèŁĖăĹææŸŕĭijŽèĹŤæĹŇçŽĎăĎŔæĂĭĭijŇăőĈçŽĎăőŠăĎŇæŸŕèũ.

Borg éŽĚççd'çőaçĤĚçşžçzŞĭijŇăŔŕăžèèŕt'æŸŕéŽĚççzŞăžĚ Borg

èőĹèőăæĂĬæĈşçŽĎçşçă■ŮĭijŇăžűăŸŤăŔŸæŤűăžĚ Borg çşççzşşŸ■çŽĎçzŔéĹŇăŞŇæŤŽèő■ăĂĈ

ăőĈçŽĎçŽőăăĖă■ăžĚăžĚæŸŕăŸĂăŸĹçijŮæŮŠçşççzşĭijŇěĂŇæŸŕăŘăžŽăŸĂăŸĹèġĎěŇĈĭijŇăŔŕăžèè

kubernetes aIJl 2014 aZt' aRSayCazEcnayAaylclLaeIjNijNcZoal' aijAezRazual' YcoajlI
Github aylAaC

<https://github.com/Kubernetes>

czOal' iijNAWSaAAeYfeGNazSaAAa;oe;razSijNcZoal' auaszRaOsctSaTraNA
K8S iijNcZoal' auaszRaRrazeeol' ctlaLuZt' aeOeClc; sazSaOscTsctZdaIJaaAaC

- aIJL' azAazLaijYaL

```
- aSzazO Borg
→ cSzczSijNèozèoaelRcESijNaijAazRaAAayTè; zèGRczgiijNcOaAaTaySaAaAaOzaYScREèg
-
→ aIaaiUaNiijNaRraRSaNiijNaTraNAeSl' aRiijNaRrazzaDRczdaRliijNä;NaeciijZc;SczI
→ flanneliijNaaYacIaRSazuiijZ
-
→ aTfEeZIJARScOriijLaYat' zaAgaeOceSlaiijL' aSNèGlalSafoad' eC; aLZiijLaL' raiJnaTreGR
-
→ aRraL' l' asTczZDetDazReGlalèrCazaeIJzaLuaiijLad' Zczt' azeqZdaert' azseGlalalL' l' aoz
```

1.3 1.3 cOraCcaeduaedD

Kubernetes aYrayAayleZEcl' d' iijNaTt' aRLad' ZaRreoaqoUaeIJZcZDeoaqoUeC; aLZiijNaocaeYrayAag
K8S eZEcl' ayayzaeIJzaLeayzaed' cggegSèLsijZ

```
MasteriijZeZEcl' d' cZDcoaqREelCcCziijNaIJL' ayAaylaLUeAeayAcZDeLcCcziijNayAeln
→ 3 ayleusad' SaZeAaC
nodesiijJaerRa; ZeoaqoUetDazRczZDeLcCcziijNaarsayreRèaNaOzaZlczZDeLcCcziijNaRrazael'
```

aocaelucnralZazaRraLaOzaZlczZDeruaesCazd' czZ Master iijN-
Master ayLaIJL' ayAaylerCazeaZlaocceC; aLeaedRaRD nodes
elCcCzayLczZDetDazRcluaAAiijNaL; ayAaylaeIAeAcARLeRèaNgTlaLuaozaZlczZDeLcCcziijNazuaIJ
Docker aRraLeLZaylaOzaZliijNnode elCcCczZD Docker
aIJlaRraLaOzaZlaUuaAZaijZeUaELacAesaeIJnaIJraIJL' aSaIJL' eTIJaCRiijNaecadIJaSaIJL' arsaZOaz
pull cDuarROeRèaNaAaC

eCcaZLazSaZSaRrazeeRèaNaYzaOzaZliijNaL' AazeazSaRrazael' YcoajlI Kubernetes
azNayliijNaEuaoD Kubernetes aRrazael' YcoaeGlezniijNaGlaL' YcoajlI

- ApiServer

kubernetes aeOeaeTucTlaLuLaZazaOzaZlcl' eruaesCZdaYr Kubernetes Clus-
teriijNeCcaZLaocCarzad' UaeRRa; ZaIJaaLacZdaOeaeRcaarsaeYrayAayl API aeOeaeRc
iijNeLZaylaeOeaeRceIAeAeAciijUclNaieoeleUoiijNaLUeAeAeAZefGcijUaEzae; cZdaocaelucnrlNaZRaeleoe
Master ayLaIJL' ayAaylczZazuaarsaeYr ApiServeriijNaieaeOeaeTuaocnreuaesCiiijNegcaedRaocaelucnreuaesC
serveraeL' eC; aSaa; IetcdiijNaEuazUalaaUaeCsèeAeOuaRUaeTraeaeIAeAeAZefGapi
serveraeRRa; ZcZdaOeaeRcefZèaNgZyaEsaeTraeaeSaaiIJ

- Scheduler

scheduler watch apiserver iij Næ Ōēā RŪçşzçzş æ L Ūç T l æ L ū ē r ū æ s C æ Y r è f R è a N i i j N æ Ç ä ; T è e A è f R è a N ä y A Master ä i j Z ä ; Ç ç T l è r C ä z e ä Z l i i j L scheduler i i j L æ ä z æ ■ ð è r ū æ s C æ l è a l E è E ■ ä y Ä ä y l è Ç ; ä d' ş è f R è a N æ ö z ä Z l ç Z D nodes è L Ç ç C z i i j N ä ; N æ Ç i i j Z æ ä z æ ■ ö ç T l æ L ū a r z è t D æ z R è e A æ s C i i j N C P U a Ä A æ E ä ■ Y ä Ä A æ l è e r D ä i j r ä ş l ä y l nodes æ I J Ä ä R L é Ä Ç è f R è a N ä Ä Ç

ä d' g æ Ç ç Z D è f G ç l N æ r s æ Y r i i j Z è e Ū ä E L æ Y r é c D é Ä l i i j N ä z Ō nodes ä y ■ æ N S é Ä l' ä G z ç n e a R L ç T l æ L ū a ö z ä Z l è f R è a N è e A æ s C ç Z D i i j N ç D ū ä R Ō ä I J l è f Z ä z Z é c D é Ä l' ç z ş æ d I J ä y ■ è f Z æ node ä Ä Ç

- Controller i i j L æ Ō g ä L ū ä Z l i i j L

æ Ç æ d I J è f R è a N æ ö z ä Z l ç Z D è L Ç ç C z ä o T æ I J z æ L Ū è Ä E ä ö z ä Z l æ I J n è z n è f R è a N æ G z ç Ō r é Ū ö é c Y i i j N k u b e r n e s è Ç ; ä d' ş ä I J ä E ū ä z Ū è L Ç ç C z ä E ■ ä R f ä L ä y Ä ä y l ä y Ä æ l ä y Ä æ ä u ç Z D ä ö z ä Z l i i j N è f Z ä r s æ Y r K u - b e r n e s æ R R ä ; Z ç Z D è G l æ D L è Ç ; ä L Z ä Ä Ç

æ Ō g ä L ū ä Z l ä r s ä o d ç Ō r ä z E ç Z S æ Ō g ä o C æ L' Ä è t' ş è t' ç ç Z D æ r R ä y Ä ä y l ä o z ä Z l ç Z D ä A è ä z ū ç L ū æ Ä A i i j N ä y Ä æ Master ä R S é Ä A è r ū æ s C i i j N Master ä i j Z ä E ■ æ n a ç T s è r C ä z e ä Z l æ N S é Ä l' ä G z ä R L é Ä Ç ç Z D è L Ç ç C z ä E ■ æ n a ç è f R è a

ä o C è Ç ; æ N A ç z ■ æ Ä g æ Ō c æ t N æ L' Ä ç o a ç R E ç Z D ä ö z ä Z l i i j N ä y Ä æ Ū e ä y ■ ä A è ä z ū i i j N æ L Ū ä y ■ ç n e a R L ç T l æ L è Ä N K u b e r n e s æ T r æ N A ä i j Ū ä d' Z ç Z D æ Ō g ä L ū ä Z l i i j N æ T r æ N A ä ö z ä Z l ä A è ä z ū ç Z D æ Ō g ä L ū ä Z l ä R l æ Y r ä E

- ControllerManager i i j L ä L ū ä Z l ç o a ç R E ä Z l i i j L

ä I J l Master ä E E ç ; ö ç z D ä z ū ä y ■ æ I J L' ä y Ä ä y l æ Ō g ä L ū ä Z l ç o a ç R E ä Z l i i j N ä o C è t' ş è t' ç ç Z S è g E ç I Ä æ r R ä y Ä ä y l æ Master æ I J L' ä d' Z ä y l i i j N æ L' Ä ä z è ä E ū æ I J L' ä E Ū ä ; Z æ Ä g ä Ä Ç

- Pod i i j L ä Ō ş ä ■ R è r C ä z e ä ■ T ä E Ç i i j N æ Y r ä ö z ä Z l ç Z D ä r A è ç E i i j L'

ä I J l K u b e r n e s ä y L è r C ä z e ç Z D ä Ō ş ä ■ R ä ■ T ä E Ç i i j N K u b e r n e s ä y ■ ç Z t' æ Ō è e r C ä z e ä ö z ä Z l i i j N è Ä N æ Y r Pod i i j N Pod ä R r ä z è ç R E è g ç ä y z ä ö z ä Z l ç Z D ä z N æ n a ä r A è ç E i i j N ä R f ä z è ç T

ä R N ä y Ä ä y l POD é G N ç Z D ä ö z ä Z l i i j N è f Y è Ç ; ä E s ä z n ä R N ä y Ä ä y l ä ■ Y ä C l ä ■ i i j N ä ■ Y ä C l ä ■ ū ä R r ä z è ä s d ä z Ō POD ä Ä Ç

ä y Ä è L n ä y Ä ä y l POD ä R l è f R è a N ä y Ä ä y l ä ö z ä Z l i i j N æ Ç æ d I J è I J ä e e A ä I J l POD æ T l' ä d' Z ä y l ä ö z ä Z l i i j N è C ç ä z

- Node i i j L ä ū è ä ; I J è L Ç ç C z i i j L

æ R R ä ; Z è o a ç Ō Ū è t D æ z R ç Z D è L Ç ç C z i i j N æ r s æ Y r è f R è a N Pod ç Z D ä y z æ I J z i i j N K u b e r n e s Cluster ç z ş ä y Ä ç o a ç R E æ L' Ä æ I J L' ç Z D node è L Ç ç C z ç Z D è o a ç Ō Ū è t D æ z R i i j N ä ; ş ç T l æ L ū ē r ū æ s C ä L Z ä z z è t D æ z R ç Z D æ Ū ū ä Z i i j N ä R f ä z è æ ç Ä æ ş è ç Z o ä L' ■ é Ç

- Label i i j L æ ä G ç ■ ; i i j L

ä y Ä ä y l ç T s key = value ç z D æ L R ç Z D æ ä G ç ■ ; i i j N ä R f ä z è ä y z POD æ L' ş ä y L ä y Ä ä y l æ ä G ç ■ ; ä Ä Ç

- Selecter i i j L æ ä G ç ■ ; é Ä l' æ N l' ä Z l i i j L

é Z E ç ; d' ä y ■ è f R è a N ç Z D ä i j Ū ä d' Z POD i i j N ä L' ■ é l c æ R R ä l R ä y Ä ä y l æ Ō g ä L ū ä Z l ä R f ä z è ç o a ç R E è N è ä z ä y l POD i i j N è C ç ä z L æ Ō g ä L ū ä Z l ä e C ä ; T ä z Ō é Z E ç ; d' ä y ■ è f R è a N ç Z D æ L' Ä æ I J L' POD ä y ■ æ N S é Ä l' ä G z æ l è e G l ä ū s é I J ä e e A ç o a ç R E ç Z D POD ä S é ?

ä I J l ä L Z ä z z ä y Ä ä y l POD ç Z D æ Ū ū ä Z ä y z POD æ L' ş ä y L ä y Ä ä y l æ ä G ç ■ ; i i j N è o l' ç l N ä z R ä R f ä z è e Ä Z è f G è f pod i i j N ä R f ä z è ç z Z æ r R ä y l pod ä L ä ä y Ä ä y l K / V ç s z ä d N ç Z D æ ä G ç ■ ; ä e C i i j Z app = nginx i i j N ä R f ä e æ L' ; ä G z è f Z ä Z

nginx pod iijÑéCčázLælažúáŕsæYŕæázæ■ó æNěæIJL' key äyž app çŽDpod ážúäyT value äyž
nginx ælææNŠaGžèfŽçžD PODãÄČ

æāGç■;äy■æYŕ POD āTŕäyĀāEūæIJL'çŽDæIJzāLūiijNāEūázŪçŽDçžDāžúāRÑæāūāRŕäzæIJL' æāGç■;

1.4 1.4 ædūædDāŠNçžDäzū

- Etcd

çTlāžŌ Kubernetes çŽDāRŌçŕŕæTŕæ■ōāYāCí,
→æL' ĀæIJL' éžEçç;d' æTŕæ■ōéČ; ā■YāCíāIJlæ■d' ād' D

- Master èŁČÇCžèŕ' šèŕ' ççžt' æLd' éžEçç;d' çŽDçŽōæāGçLúæĀAiiijNäyLéIcèfRèaŇçŽDäyžæŌgçžDäzūæI

```
kube-apiserver # áŕžād' Ūæžt' éIJšāžE Kubernetes
→APIiijNāōČæYŕçŽD Kubernetes āL'■çŕŕæŌgāLūāsCiiijNāŕlæIJL' API
→Server äijžäyŌ etcd éĀžāfaiijNāEūāōČælaāiŪéČ; āfĒéažéĀžèfĠ API
→Server èōféŪōéžEçç;d' çLúæĀA
kube-controller-manager #
→ād' DçŕEéžEçç;d' äy■āyÿègDäzžāLaiijNāōČæYŕā■TçNŇçŽDèfŽçlŇiijNāEĒéČlāNĒāŕNād' žäyſæ
→POD æTŕéĠŔ
kube-scheduler # çŽSègEæŪŕāLžāžžçŽD Pod
→äyžæŪŕāLžāžžçŽD POD āLEéĒ■āŕLÉĀČçŽD node èŁČÇCž
```

- Node èŁČÇCžāōdéZĒèŕ' šèŕ' čāōdæŪ;iiijNāžšāŕsæYŕèfRèaŇ POD
çŽDèŁČÇCžiiijNäyLéIcèfRèaŇçŽDçžDäzūæIJL'

```
kubelet #
→èŁČÇCžèĠlæšlāEŇāŠNèŁČÇCžçLúæĀAæžt' æŪŕiiijNāōČçŽSætNāūsçžŕāLEéĒ■çžžèĠlāūsçŽD
→PodiiijNäyž POD āĠEād' Gā■ūiiijNäyNè;; POD æL' ĀéIJĀçŽD
→SecretiiijNäyNè;; éTIIJāČŕāžūèfRèaŇiiijNèfžèaŇçTšāS; āSīæIJSæŌcætNiiijNäyžæLē
→POD āŠNèŁČÇCžçLúæĀA
kube-proxy #
→éĀžèfĠççžt' æLd' äyžæIJžäyLçžŽDç; ŠçžIJègDāLžāžūæL' ġèaŇèfðæŌèè; nāŕSiiijNāŕE
→Kubernetes
→æŕŕä;žçŽDç; ŠçžIJæIJ■āLāāžççŕEāLŕæŕŕäyſèŁČÇCžäyLiiijNāōdçŌŕāžEĠKubernetesæIJ■āLāæ
docker # çTlāžŌèfRèaŇāōžāžl
```

- æŕŠäzū

æŕŠäzūæYŕācđaijžéžEçç;d' āLŠèČ; çŽD Pod āŠN Service,
→æŕŠäzūāŕžésæaIJnèžŇæYŕāŕŪāS; āŕ■çl' žéŪt' éžŕāLūçŽD, ècŇāLžāžžāžŌ kube-
→system āS; āŕ■çl' žéŪt' .

- DNS

èž; çDūāEūázŪæŕŠäzūāžūäy■æYŕāfĒéIJĀçŽD, ä; EæL' ĀæIJL' Kubernetes
→éžEçç;d' éČ; āžTèŕéāEūæIJL' Cluster DNS, èōyād' ŽāžTçTlā; iètŪāžŌāōČ, äyž
→Kubernetes æIJ■āLāæŕŕä;žDNSèōŕā; T, āōžāžlāŕŕāLlèŕéāŕŌäijžèĠlāLlāŕE
→DNS æIJ■āLāāžlāNĒāŕNāIJl resolv.conf äy■.

CHAPTER 2

äžŇ æăÿǻ£ČčžĎäžů/éŽĎäžů

2.1 2.1 Controller

[illegible]

- RelicationController

æŒǵáLúáRÑäYĂçsz POD áržèsaçŽĐàL'raeIJñæTřéĠŘijNăođčŎřcÍNăžRçŽĐæzŽăLíaZt' æŨrijNæLŮè
 alIjæzŽăLíaZt' æŨræUúăĂžijNăĚAèóyăyt' æŮüèùĚăGžegĐăďŽçŽĐàL'raeIJñæTřéĠŘijN

- RelicaSet

Deployment

- Deployment

ãöČăŕlêÇjçõäqřĚæUáčŁúæĂAçŽĐžTčŤlíijNěšZäyłæÖğáLúaŻlíijNæŦŕæÑAäžNçžgæÖğáLúaŻlíijNă
Pod AutoscaleríijNærťázş POD êĞlâŁlâijÿcijl' æÖğáLúaŻlíijLíijNă;Şet' şè;enŸçŽĐæUűăĂŻíijNěĞlâŁlâŔŕă
PODăĂĆ

- StatefulSet

çõaçŘĚæIJL'çŁúæĂAçŽĎăžŤçÍ

- DaemonSet

æĈæđIēIJǼēAāIJǣrŘäyÄäy! node äyLèŁRëaNäyÄäy!aL'ráIJniiJÑeĂŃäy■æÝřěŽRæĐŘèŁŘëaN

- Job

ɛ̃ʁRɛ̃aŋNäyÄæñæÄgä;IJäyŽiijNæŮúÉŮr'äy■ǎŽžǎŮŽçŽĐæŠ■ä;IJiijNä;NäeĆiijŽǎđ'Gäz;ǎÄAæyĚçŘEiij
 POD ælěɛ̃ʁŽèaŇǎđ'Gäz;çŽĐžžǎŁaiijNěšRɛ̃aŇǎŮNæĹŔǎřščzŠæİšǎžEǎĀĆ

æL'ÄäzëijÑæÖëÄëäd'ÚéCíèðféUöæUüäÄZüijÑëruæśCéçÚäĚLäLrè; node
 ç;ŚçzIJüijÑçDüäRÖ node ç;ŚçzIJäzççRĚèGš service ç;ŚçzIJüijÑservice æāzæ■ō iptables/ipvs
 èğDāLZäëë;ñāRŚāLr pod ç;ŚçzIJäy■çZD pod äyLāĀĆ ~~~ NODE ç;ŚçzIJ -> SVC ç;ŚçzIJ
 -> POD ç;ŚçzIJ ~~~

k8s æIJL'äyL'çg■éÄZäfaüijZ

- āRÑäyÄäyġ POD āĒĚçZDäd'ZäyġäözāZġéUť çZDēÄZäfaüijÑāRrāzëéÄZēfG lo
 éÄZäfaçZť æÖëéÄZëðfāĀĆ
- POD äyÖ POD éÄZäfaüijÑāçCædIJä;fçTġ flannel æL'ÄæIJL' POD
 éÇ;äd'DäzÖäyÄäyġç;ŚçzIJüijÑāRrāzëëüġ node äyÖāRĚäd'ÚçZD POD
 çZť æÖëéÄZäfaüijÑāZäyZä;fçTġläzĒāRāāLäç;ŚçzIJāĀĆ
- POD äyÖ Service éÄZäfaāĀĆ

2.4 2.4 kube-proxy

āIJġ node èLCçCzäyLèfRëaÑçZDäyÄäyġäöLæLd'èfZçġNüijÑāöÇet'šet'céZRæUüäyÖ
 apiserver èfZëaÑéÄZäfaüijÑāZäyZæfRäyġ pod āRŚçTšāRŸāNŪāRÖéIJÄëeAäfiā■YāIJġ
 apiserver äy■üijÑëÄÑ apiserver āRŚçTšæTzāRŸāRÖäijZçTšæLŸäyÄäyġéÄZçšëazNāzūüijÑëfZäyġazNāzūāR
 kube-proxy äyÄæUëāRŚçÖræšRäyġ service āRÖçnrçZD pod
 āIJŕāġāRŚçTšæTzāRŸüijÑëCçāzLāršçTš kube-proxy èť'šet'cāIJġæIJñāIJŕāŕĒāIJŕāġāĒEZāĒë
 iptables æLŪèÄĒ ipvs èğDāLZäy■āĀĆ

æL'Ääzë service çZDçöaçRĚæYŕéġä kube-proxy æġäöðçÖrçZDüijÑā;šā;āāLZāzzäyÄäyġ
 service üijÑëCçāzLāršéġä kube-proxy āIJġæfRäyġèLCçCzäyLāLZāzzäyZ iptables æLŪèÄĒ ipvs
 èğDāLZüijÑæfRäyġ service çZDāRŸāLäzšéIJÄëeA kube-proxy āR■āzTāLŕèğDāLZäyLāĀĆ

apiserver éIJÄëeAäfiā■YāRĎäyġ node äfaæAŕüijÑāöÇéIJÄëeAäfiā■YāIJġ etcd äy■āĀĆ

2.5 2.5 etcd

æYŕäyÄäyġéTōāAijā■YāCġçZDçszçzšüijÑäyÖ redis āġLāČŕüijÑā;ĒæYŕ etcd
 èfYæIJL'äyÄäzZā■RërČāLšëÇ;æYŕ redis æL'Ääy■āĒüäd'GçZDüijÑāöÇefYæIJL'èLCçCzéĀL'äyç■L'āLšëÇ
 etcd æZť āČŕ zookeepērāĀĆ

çTšāzÖæTť äyġéZĒç; d'çZDæL'ÄæIJL'äfaæAŕéÇ;äfiā■YāIJġ etcdüijÑæL'Ääzë
 etcd āçCædIJāöTæIJüijÑëCçāzLæTť äyġéZĒç; d'āršæNČāzĒüijÑāZäëÄÑ etcd
 éIJÄëeAāZëñYāŕŕçTġāĀĆ

2.6 2.6 flannel

æL'YçöaäyZ k8s çZDēZDäzüèfRëaÑ, āIJġ k8s äy■æIJL'āġLād'ZāĒüäzŪçZDāijÄæzŔç;ŚçzIJæŔŚāzūüijÑā
 calico äyL'āsČç;ŚçzIJæŔŚāzū,æĀğëÇ;āġLāë;üijÑæTŕæNĀèðféUöæÖğāLŪ

node ç;ŚçzIJüijZçL'ŕçRĒāŔDèLCçCzāzNéUť èfZëaÑéÄZäfa

POD ęŚçŻIİijŽæL'ĂæIJL' nodeäyŁçŽĐ POD āijæ■d'āzNéŮt'éĂŽèŁĠāRāāLāriijNæLŮèĂĚçŽt'æŌèèŭr
service ęŚçŻIİijŽçŤś kube-proxy èt'šet'čçōæŌġāŠNçŤśæLŘ

2.7 çšëèÉāŕŘçžŚ

- Master

```
kube-scheduler          # èřČāžę pod
kuber-controller-manager # çŏaçŘĚ pod
kube-apiserver          # æŌěæŤűèŕŭæśĆ
etcd                    # └
→ éŽĚççd'çŁúæĀĀā■ŸāĆÍiijNéŽĚççd'æL'ĂæIJL'çŽĐçžĐäžűçŽĐçŁúæĀĀéČ;ăĚĪā■ŸāIJĪéŁŽéĠN
```

- node

```
kubelet                # èŁĆçĆž/podçŏaçŘĚ
kube-proxy             # watch apiserverçŏaçŘĚservice
docker                 # āŌžāŽĪèŁŘèāNæŮű
```

CHAPTER 3

äyL' éZEÇçŁd' éČÍçĳš

äyžçõÄā■TäyLæL'Nä;ŞeiNāLŞèČĳĳĳNāRfäzēāĒLāL'çTĳkubeadmāōL'ècĒætNērTĳĳNçTšāžğçŎřácČāžž
æYř K8S āōYæŮzæRRä;ZçŽDāŁnéĀŞéČÍçĳšāũēāĒũĳĳNāōČæRRä;ZāžĒ kubeadm
init äžēāRĒ kubeadm join èŁŽäy'd'äyŁāŚĳāzd'äĳĳäyžāŁnéĀŞāLZāžž kubernetes
éZEÇçŁd'çŽDæĪĀäĳšāōðēũĳĳNæĪñçnäèLČērt'æYŎāžĒäĳŁçTĳ kubeadm ælēéČÍçĳš K8S
éZEÇçŁd'çŽDèŁĞçĳNāĀČ

- éZEÇçŁd'çŽDçzĞçzŞæđD

éaqçŽŎ	èrt'æYŎ
éZEÇçŁd'èğĐæĪā	MasterāĀĀnode1āĀĀnode2
çşžçžş	CentOS 7.3
çĳŚçzĪJèğĐāLŠ	PODĳĳŽ10.244.0.0/16āĀĀServiceĳĳŽ10.96.0.0/12

3.1 3.1 éČÍçĳšāL'■āĞĒāđ'Ğ

æĪñāřRèLČçŽDæL'ĀæĪJL'çŽDæŞ■āĳĳĳNāĪĪæL'ĀæĪJL'çŽDèLČçČzäyLèŁŽēāŃ

3.1.1 3.1.1 āĒŞēŮ■ firewalld āŠŃ selinux

```
setenforce 0
sed -i '/^SELINUX=/cSELINUX=disabled' /etc/selinux/config

systemctl stop firewalld
systemctl disable firewalld
```


æĈædIæIaäzŭäy■āĖAëöyāRřāzēāRCèĀĈ: <https://blog.csdn.net/jinguangliu/article/details/82792617> æĬëğĉāEşéTĬJāĈRĕŬőécŸāĀĈ

- éĖ■ĉjō Docker æNL'āRŬéTĬJāĈRæŬŭāĀŽĉŽDžĉĉRĖāIŖāIĀiijŃvim /usr/lib/systemd/system/docker.serviceāĀĈ

```
[Service]
Environment="HTTPS_PROXY=127.0.0.1:9666"
Environment="NO_PROXY=127.0.0.0/8,172.16.0.0/16"
```

- æRŘāL■æNL'āRŬāĬiāğNāŃŬĖIJĀèēAĉŽDĖTĬJāĈR

```
kubeadm config images pull
```

- äjĕĉTĬāĖŭäzŬæžRĖTĬJāĈR

```
docker pull mirrorgooglecontainers/kube-apiserver:v1.14.2
docker pull mirrorgooglecontainers/kube-controller-manager:v1.14.2
docker pull mirrorgooglecontainers/kube-scheduler:v1.14.2
docker pull mirrorgooglecontainers/kube-proxy:v1.14.2
docker pull mirrorgooglecontainers/pause:3.1
docker pull mirrorgooglecontainers/etcd:3.3.10
docker pull coredns/coredns:1.3.1

āĬl'ĉTĬ`kubeadm config images list` æSĖĉIJŃĖIJĀèēAĉŽDdocker image_
↪name

k8s.gcr.io/kube-apiserver:v1.14.2
k8s.gcr.io/kube-controller-manager:v1.14.2
k8s.gcr.io/kube-scheduler:v1.14.2
k8s.gcr.io/kube-proxy:v1.14.2
k8s.gcr.io/pause:3.1
k8s.gcr.io/etcd:3.3.10
k8s.gcr.io/coredns:1.3.1

# äfōæTžtag

docker tag docker.io/mirrorgooglecontainers/kube-apiserver:v1.14.2_
↪k8s.gcr.io/kube-apiserver:v1.14.2
docker tag docker.io/mirrorgooglecontainers/kube-scheduler:v1.14.2_
↪k8s.gcr.io/kube-scheduler:v1.14.2
docker tag docker.io/mirrorgooglecontainers/kube-proxy:v1.14.2 k8s.
↪gcr.io/kube-proxy:v1.14.2
docker tag docker.io/mirrorgooglecontainers/kube-controller-
↪manager:v1.14.2 k8s.gcr.io/kube-controller-manager:v1.14.2
docker tag docker.io/mirrorgooglecontainers/etcd:3.3.10 k8s.gcr.io/
↪etcd:3.3.10
docker tag docker.io/mirrorgooglecontainers/pause:3.1 k8s.gcr.io/
↪pause:3.1
docker tag docker.io/coredns/coredns:1.3.1 k8s.gcr.io/coredns:1.3.1
```

(continues on next page)

(continued from previous page)

```
docker rmi `docker images |grep docker.io/ |awk '{print $1":"$2}'`
```

3.2.2 3.2.2 aLlāgNāŃŮMaster

- ä;£çŤÍ kubeadm āLlāgNāŃŮ k8s éŽEç;d'

```
kubeadm init --kubernetes-version=v1.14.0 --pod-network-cidr=10.244.
↪0.0/16 --service-cidr=10.96.0.0/12 --ignore-preflight-errors=Swap
```

- āēČædIJæIJL'æLēéŤŽä;£çŤlāyNéIcāŚ;äzd'æšēçIJN

```
journalctl -xeu kubelet
```

- āēČædIJāLlāgNāŃŮēfGčlNēcñäy■æŮ■āRfäzēä;£çŤlāyNéIcāŚ;äzd'ælēæAčād'■

```
kubeadm reset
```

- äyNéIcāŸræIJāĀŔŌæL'gèaŃæLŔāLšæŸçd'žçŽDçzšædIJijNéIJĀēçAäflā■ŸēfŽäyĽæL'gèaŃçzšædI.
node èLČçCžāLāāĒēēZEç;d'

Your Kubernetes control-plane has initialized successfully!

To start using your cluster, you need to run the following as a a
↪regular user:

```
mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
```

You should now deploy a pod network to the cluster.

Run "kubectl apply -f [podnetwork].yaml" with one of the options a
↪listed at:

<https://kubernetes.io/docs/concepts/cluster-administration/addons/>

Then you can join any number of worker nodes by running the a
↪following on each as root:

```
kubeadm join 172.16.100.9:6443 --token 2dyd69.hrfsjkkxs4stim7n \
--discovery-token-ca-cert-hash a
↪sha256:4e30c1f41aefb177b708a404ccb7e818e31647c7dbdd2d42f6c5c9894b6f41e7
```

- æIJĀäç;äzēæŽōéĀŽçŤlāLūçŽDèžnāz;ēfRèaŃäyNéIcçŽDāŚ;äzd'

```
# āIJlā;šāL'■çŤlāLūāōūçŽōā;ŤäyNāLžāžž.
↪kubēçŽōā;ŤāžūēĒ■ç;ōēōēūōēēZEç;d'çŽDconfig æŮGžāžū
mkdir -p $HOME/.kube
```

(continues on next page)

(continued from previous page)

```
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
```

- éČĺčjš flannel çİŚçzIJæŘŠázŭ

```
kubectl apply -f https://raw.githubusercontent.com/coreos/flannel/
↳master/Documentation/kube-flannel.yml
```

- æšěçIJN kube-system åŚ;åŘ■çl'žéŮt'äy■ēfŘèqNçŽĎ pods

```
kubectl get pods -n kube-system
```

- æšěçIJN k8s éŽEçİd'çzĎžžçŽĎçŁŭæĀA

```
kubectl get ComponentStatus
```

- éĚ■çİōāŚİäzd'èačāĚÍ

```
yum install -y bash-completion
source /usr/share/bash-completion/bash_completion
source <(kubectl completion bash)
echo "source <(kubectl completion bash)" >> ~/.bashrc
```

3.3 3.3 éČĺčjš Node

æIJnårRèŁĆçŽĎæL'ĀæIJLçŽĎæŞ■äİIijNåRİåIJÍ Node èŁĆçĆžäyŁèfŽèqNāĀĆ

3.3.1 3.3.1 åŁāāĚēēŽEçİd'

- åŁāāĚēēŽEçİd'ijNæşİæĎŘåIJÍåŚİäzd'årçéČÍåŁäyŁ --ignore-
preflight-errors=Swap ijNāžēāfİçTě k8s åřžäyæIJŽ swap
çŽĎæčĀæšēijLk8sāyžāžEæĀğèÇİæL'ĀäžēēAæśĆèfŽāŁŭ swap ijL'

```
kubeadm join 172.16.100.9:6443 --token 2dyd69.hrfsjkkxs4stim7n \
--discovery-token-ca-cert-hash
↳sha256:4e30c1f41aefb177b708a404ccb7e818e31647c7dbdd2d42f6c5c9894b6f41e7
↳--ignore-preflight-errors=Swap
```

- èĚŤāŽđçzŞædIJijNēāİçd'žåŁāāĚēēŽEçİd'æĹŘāŁŞ

```
This node has joined the cluster:
* Certificate signing request was sent to apiserver and a response
↳was received.
* The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the control-plane to see this node join
↳the cluster.
```

(continues on next page)

(continued from previous page)

3.3.2 3.3.2 æšëçIJNè£Žāžę

ā;Š node èŁĈĉĆzāŁāāĔē K8S éŽĖçŁd'äy■āRŌiijŃMaster äijŽērĈāžęāŁr Node
èŁĈĉĆzāyŁāyĀāžŽçzĎžūiijŃçTlāžŌād'ĎçRĖéŽĖçŁd'āžNāŁāiijNè£ŽāžŽçzĎžūāšęæIJL'äyNè;āōNāēLŘā
Node èŁĈĉĆzāIJléŽĖçŁd'äy■ē£YæYræIJlāšçzŁçLūāĀĀ

- āIJl node æLğēāNāyNélcāŚ;āzd' iijNāRrāžęæšëçIJNéTlJāČRçŽĎäyNè;ç£ŽāžęiijNāyNélcæYræIJĀçz

```
$ docker image ls
REPOSITORY          TAG                 IMAGE ID            6
→CREATED            SIZE
k8s.gcr.io/kube-proxy v1.14.0            5cd54e388aba        6
→weeks ago          82.1MB
quay.io/coreos/flannel v0.11.0-amd64      ff281650a721        3
→months ago         52.6MB
k8s.gcr.io/pause     3.1                da86e6ba6ca1        16
→months ago         742kB
```

- āRrāžęāIJl Master äyŁā;ççTlāyNélcāŚ;āzd' ælēæšëçIJNæŪrāŁāāĔēçŽĎēŁĈĉĆzçLūāĀĀ

```
$ kubectl get nodes
NAME      STATUS    ROLES    AGE   VERSION
master    Ready     master   3d21h v1.14.1
node1     Ready     <none>   3d21h v1.14.1
node2     Ready     <none>   3d21h v1.14.1
```

- æšëçIJNéŽĖçŁd'çLūāĀĀ

```
[root@master ~]# kubectl cluster-info
Kubernetes master is running at https://10.234.2.204:6443
KubeDNS is running at https://10.234.2.204:6443/api/v1/namespaces/
→kube-system/services/kube-dns:dns/proxy
Metrics-server is running at https://10.234.2.204:6443/api/v1/
→namespaces/kube-system/services/https:metrics-server:/proxy

To further debug and diagnose cluster problems, use 'kubectl
→cluster-info dump'.
[root@master ~]# kubectl get componentstatuses
NAME                STATUS    MESSAGE                               ERROR
controller-manager  Healthy   ok
scheduler           Healthy   ok
etcd-0              Healthy   {"health": "true"}
```

āęĆādIJānŃç;ŚçzIJpulléTlJāČRæĔcāRrāžęāIJlāyĀāRrāyLélcārĖçTlJāČRæL'ŠāNēāRŚéĀĀęGšāĔūāžŪ

- | | | | | | | |
|--|---------------|------------|---------|-----------|--|--|
| \$ kubectl get pods -n kube-system -o wide | | | | | | |
| NAME | | READY | STATUS | RESTARTS | | |
| ↪AGE | IP | NOMINATED | NODE | READINESS | | |
| ↪GATES | | | | | | |
| coredns-fb8b8dccf-cp24r | | 1/1 | Running | 0 | | |
| ↪26m | 10.244.0.2 | i-xeahpl98 | <none> | <none> | | |
| coredns-fb8b8dccf-ljswp | | 1/1 | Running | 0 | | |
| ↪26m | 10.244.0.3 | i-xeahpl98 | <none> | <none> | | |
| etcd-i-xeahpl98 | | 1/1 | Running | 0 | | |
| ↪25m | 172.16.100.9 | i-xeahpl98 | <none> | <none> | | |
| kube-apiserver-i-xeahpl98 | | 1/1 | Running | 0 | | |
| ↪25m | 172.16.100.9 | i-xeahpl98 | <none> | <none> | | |
| kube-controller-manager-i-xeahpl98 | | 1/1 | Running | 0 | | |
| ↪25m | 172.16.100.9 | i-xeahpl98 | <none> | <none> | | |
| kube-flannel-ds-amd64-crft8 | | 1/1 | Running | 3 | | |
| ↪16m | 172.16.100.6 | i-me87b6gw | <none> | <none> | | |
| kube-flannel-ds-amd64-nckw4 | | 1/1 | Running | 0 | | |
| ↪6m41s | 172.16.100.10 | i-qhcc2owe | <none> | <none> | | |
| kube-flannel-ds-amd64-zb7sg | | 1/1 | Running | 0 | | |
| ↪23m | 172.16.100.9 | i-xeahpl98 | <none> | <none> | | |
| kube-proxy-7kjkf | | 1/1 | Running | 0 | | |
| ↪6m41s | 172.16.100.10 | i-qhcc2owe | <none> | <none> | | |
| kube-proxy-c5xs2 | | 1/1 | Running | 2 | | |
| ↪16m | 172.16.100.6 | i-me87b6gw | <none> | <none> | | |
| kube-proxy-rdzq2 | | 1/1 | Running | 0 | | |
| ↪26m | 172.16.100.9 | i-xeahpl98 | <none> | <none> | | |
| kube-scheduler-i-xeahpl98 | | 1/1 | Running | 0 | | |
| ↪25m | 172.16.100.9 | i-xeahpl98 | <none> | <none> | | |

node èŁĆçĆzéIJAèçAçŁzäçZäyNè||éTIIJäČŘäď'łæĚćiiJŇázžèőőä;ŁçTl docker
éTIIJäČŘčŽĎärijaĚĚärijaĠzäŁšèÇ; äĚŁäřĚmasterçŽĎäyL'äyłéTIIJäČŘæLŠäŇĚäŘSéĀAŁĹn-
odeèŁĆçĆziiJŇloadäŘŌäE■jion

-
- | | |
|----------------------------|-----------|
| 3.3. 3.3 éČícš Node | 24 |
|----------------------------|-----------|

```
docker image save -o /tmp/kube-proxy.tar k8s.gcr.io/kube-proxy
docker image save -o /tmp/flannel.tar quay.io/coreos/flannel
docker image save -o /tmp/pause.tar k8s.gcr.io/pause
```

- řířřřř

```
docker image load -i /tmp/kube-proxy.tar
docker image load -i /tmp/pause.tar
docker image load -i /tmp/flannel.tar
```

CHAPTER 4

4.1 4.1 kubecti

4.1 4.1 kubecti

kubecti æŸŕ apiserver ǵŽĐăôćæĹŭćŕćĹŃăžŔiijNèŁŽăŸĹăôćæĹŭćŕćĹŃăžŔæŸŕéĂŽèŁĜèđæŌě
master èĹĆćCžăŸĹćŽĐ apiserver iijNăôđćŌŕăŔĐćĝ■ k8s
ăržèśăǵŽĐăćđăĹăæŤžăšćē■ĹăšžæIJŋă\$■ă;IJiijNăIJĹ k8s ărŕèćŋćôăǵŔĖćŽĐăržèśăæIJĹăĹăđ'ŽăŸĹ

```
ăšžæIJŋăš;ăzd' (ăĹĹćžĝ) :  
  create      äžŌæŮĜăžŭæĹŬăĜăĜĖè;šăĖăăĹžăžžèťĐăžŔ  
  expose      èŌŭăŔŮăŸĂăŸĹăđ'■ăĹŭăŌĝăĹŭăžĹ, æIJ■ăĹă, _  
→éćĹć;šăĹŬèĂĖăžť'éIJšăŸĂăŸĹ POD ărĖăĖŭă;IJăŸžæŮŕćŽĐ Kubernetes _  
→æIJ■ăĹăăĖăăĹăijĂ  
  run         âĹžăžžăžŭèĹŕăăŇćĹ'žăôžćŽĐéŤIJăĈŔ, âĹžăžžă;ĹćŤĹ _  
→deployment æĹŬ job ǵôăǵŔĖćŽĐăôžăžĹ  
  set         èŌ;ć;ŏăŕžèśăǵŽĐćĹ'žăôžăĹšèć; , ä;NăēĆăŔšăŸć, _  
→ærŔăŋăăŌžset äŸ■ćŤĹćŽĐimage tag  
  
ăšžæIJŋăš;ăzd' (ăŸ■ćžĝ) :  
  explain     æŮĜăăćæĹŬèĂĖèťĐăžŔ, _  
→ăŔŕăžèćŤĹăĹăšēćIJŇèťĐăžŔăŸĖăăŤăĖžăşŤ  
  get         æŸ;ćđ'žăŸĂăŸĹăĹŬăđ'ŽăŸĹèťĐăžŔ  
  edit        ǵijŮè;šæIJ■ăĹăăžĹăŸĹćŽĐèťĐăžŔ  
  delete      æŇĹ'æŮĜăžŭăŔ■, æăĜăĜĖè;šăĖă, _  
→èťĐăžŔăšŇăŔ■ĝŕăĹŬèťĐăžŔăšŇăăĜć■;éĂĹ'æŇĹ'ăžĹăĹăéžđ'èťĐăžŔ  
  
éćĹć;šăš;ăzd' :  
  rollout     ǵôăǵŔĖèťĐăžŔćŽĐéćĹć;š  
  scale       äŸžéćĹć;šèŌ;ć;ŏăŮŕăđ'ĝăŕŔ, ReplicaSet, Replication _  
→Controller, Job
```

(continues on next page)

(continued from previous page)

autoscale	èĠlâĤlâL' l' âšTäyÄäyġéĈlċ;š, ReplicaSet, æĻŪèÄĔ
→ ReplicationController	
çċd' éŽEçõaçŘĔäš; äzd' :	
certificate	äĤöæTžèrAäžèĤĎæžŘ
cluster-info	æŸċd' žçċd' éŽEäfaæAř
top	æŸċd' žèĤĎæžŘ (CPU / äĔĔäŸ / äŸäĈl) ä; ĤçTġäĈĔäĔĤ,
→ éIJÄèæAäŌL' èĈĔmetrics-server	
cordon	ârĔèĤĈĈzæäĠèŏräyžäy■ârĤèĤĈäžĕ
uncordon	ârĔèĤĈĈzæäĠèŏräyžârĤèĤĈäžĕ
drain	èŏ; äŏž node èĤžäĔĕçzt' æĤd' æġaaijŘ
taint	æžt' æŮräyÄäyġæĻŪäd' žäyġèĤĈĈzäyĤçžĎæsaçĈz
æŤĔéŽIJæŌšéžd' äšNèĤĈérĤäš; äzd' :	
describe	æŸċd' žçĤL' žäŏžèĤĎæžŘæĻŪèĤĎæžŘçžĎçžĎèĤçzĔäfaæAř
logs	âIJġäŏžäžġläy■æL' šä■räŏžäžġġçžĎæŮèäĤŮ
attach	éžĎäĤäâġræ■ĈâIJġèĤĤæaŤçžĎäŏžäžġ
exec	âIJġäŏžäžġläy■æL' ġæaŤäš; äzd'
port-forward	ârĔäyÄäyġæĻŪäd' žäyġæIJŋâIJřçŋrârĈè; ŋâršäġr pod
proxy	èĤĤæaŤäžĈçŘĔäġr Kubernetes API æIJ■äĤaäžġ
cp	ârĔæŮĠäzŮäšNçžŏä; Ĥäd' ■äġŮäġrâŏžäžġ,
→ äšNäzŏäŏžäžġäd' ■äġŮ, èŮġäŏžäžġäd' ■äġŮæŮĠäzŮ	
auth	æĈÄæšæŏĤæĤĈ
énŸçžġäš; äzd' :	
diff	éšġäržârĔèæAäžĤçTġçžĎçĤL' ġæIJŋçžĎ Diff
→ äŏðæŮŮçĤL' ġæIJŋ	
apply	
→ éÄžèĤĠæŮĠäzŮäŤ■æĻŪæäĠäĠĔè; šäĔèæĤĔéĔ■ç; ŏäžĤçTġäžŏèĤĎæžŘ	
patch	ä; ĤçTġç■ŮçĤĔæġġäžŮèæäyAæžt' æŮĤèĤĎæžŘçžĎä■ŮæŏĤ
replace	çTġæŮĠäzŮäŤ■æĻŪæäĠäĠĔè; šäĔèæžĤæ■ĈèĤĎæžŘ
wait	äŏðéġNéŸŮæŏĤäš; äzd' :
→ âIJġäyÄäyġæĻŪäd' žäyġèĤĎæžŘäyĤç■æL' ä; ĔçĤL' žäŏžæġaäžŮ,	
→ äŏžäžĤL' äyÄäyġèġæâršäžġ	
convert	âIJġäy■ârĤçžĎAPIçĤL' ġæIJŋäžNéŮĤ' è; ŋæ■ĈéĔ■ç; ŏæŮĠäzŮ
kustomize	äžŏçžŏä; ĤæĻŪèĤIJçġġN URL æĎĎäžž kustomization çžŏæäĠ
èŏçç; ŏäš; äzd' :	
label	æžt' æŮĤèĤĎæžŘäyĤçžĎæäĠç■ç
annotate	æžt' æŮĤèĤĎæžŘäyĤçžĎæšġéĠĤ
completion	äš; äzd' èæäĔġçžyäĔšäĤSèĈ;
äĔŮäžŮäš; äzd' :	
api-resources	âIJġæIJ■äĤaäžġläyġæL' šä■räĤræŤäçžĎAPIèĤĎæžŘ
api-versions	äžĔ "group/version"
→ çžĎä; çaijŘâIJġæIJ■äĤaäžġläyġæL' šä■räĤræŤäçžĎAPIçĤL' ġæIJŋ	
config	äĤöæTž kubeconfig æŮĠäzŮ
plugin	æĤĤä; žäyŏæĤšäžŮäžd' äžšçžĎäŏðçTġçġġäžĤ
version	æL' šä■räŏçæġŮçŋrâršNæIJ■äĤaäžġĤL' ġæIJŋäĤæAř

4.2 4.2 run

- `kubectl run nginx --image=nginx:latest`

`kubectl run nginx --image=nginx:latest`

```
kubectl run nginx --image=nginx:latest
```

- `kubectl run nginx --image=nginx --replicas=5`

```
kubectl run nginx --image=nginx --replicas=5 # 5 replicas
```

- `kubectl run nginx --image=nginx --command -- <cmd> <arg1> ... <argN>`

```
kubectl run nginx --image=nginx --command -- <cmd> <arg1> ... <argN>
```

- `kubectl run pi --schedule="0/5 * * * ?" --image=perl --`

```
kubectl run pi --schedule="0/5 * * * ?" --image=perl --
  restart=OnFailure -- perl -Mbignum=bpi -wle 'print bpi(2000)'
```

- `kubectl run nginx --image=nginx --replicas=1 --dry-run`

`kubectl run nginx --image=nginx --replicas=1 --dry-run`

```
kubectl run nginx --image=nginx --port=80 --replicas=1 --dry-run=true
```

- `kubectl get deployment`

```
kubectl get deployment
```

- `kubectl get pod -o wide`

```
kubectl get pod -o wide
```

- `curl 10.244.2.2`

```
curl 10.244.2.2
```

- `kubectl delete pods nginx-deploy-5c9b546997-jsmk6`

```
kubectl delete pods nginx-deploy-5c9b546997-jsmk6
```

- `kubectl delete pods nginx-deploy-5c9b546997-jsmk6`

```
kubectl get pod -o wide
```

4.3 4.3 expose

service ælæäzcçŘE POD ælæäLŽäzzäyÄäyŁaŽžāōŽčŽĐçnrčCžāĀC

- āLŽäzzäyÄäyŁ service ælæäŽt' éIJšäyÄäyŁæIJ■āŁa

āIJæŌgāLūāŽÍ nginx-deploy äyŁāLŽäzzāR■ā■Uäyž nginx çŽĐ service , āōČauēā;IJçnrāRčäyž 80, äzcçŘEçŽĐāRŌçnrāōžāŽÍçnrāRč 80, ā■Rèōōäyž TCP

```
kubectl expose deployment nginx-deploy --name=nginx --port=80 --
→target-port=80 --protocol=TCP
```

- āRřäzēcIJNāLřāLŽāLŽāLŽäzzçŽĐāR■ā■Uäyž nginx çŽĐ service
iijNčŌrāIJlārsāRřäzēāIJléZEç; d'āEĚçTÍ service çŽĐāIJrāIÄælēēōēUōāžE,
āēČædIJād'ŪēČlēōēUōāRřäzēā;ēçTÍ NodePort ælqāijR

```
kubectl get service
```

- āLäēŽd'äyÄäyŁäzzāŁa

```
kubectl delete deployment nginx-deploy
```

4.4 4.4 cp

- æNùet'iaōēäyžæIJæŪGäzūæLŪçŽōā;TāLřpodäy■iijNāŽäyřRēçAæśČtarāžNēŁZāLūæŪGäzūāũšçzRā■

```
kubectl cp /tmp/foo_dir <some-pod>:/tmp/bar_dir

[root@master ~]# kubectl cp flannel.tar nginx-58cd4d4f44-8pwb7:/
→usr/share/nginx/html
[root@master ~]# kubectl cp mainfile/ nginx-58cd4d4f44-8pwb7:/usr/
→share/nginx/html
[root@master ~]# kubectl exec -it nginx-58cd4d4f44-8pwb7 -- /bin/
→bash
root@nginx-58cd4d4f44-8pwb7:/# ls -l /usr/share/nginx/html/
total 54108
-rw-r--r-- 1 root root      537 Jul 11  2017 50x.html
-rw-r--r-- 1 root root      355 May 27  06:47 dashboard-adminuser.
→yaml
-rw----- 1 root root 55390720 May 27  01:49 flannel.tar
-rw-r--r-- 1 root root      612 Jul 11  2017 index.html
drwxr-xr-x 4 root root       51 Aug 17 14:16 mainfile
```


4.7 4.7 ælaæNş POD ecnåLæÉZd'

- çŖåIJLæLSäznåLæÉZd' service åŖŖçŖçŽĐ POD iijNåL'raeIJnæŖgålŭåZlâijŽeGlaLlâLZâzzæŨŖçŽĐ PODiijNèĀN service åLZâijŽeGlaLlæNĠaŖSæŨŖåLZâzzçŽĐ POD

```
kubectl delete pods nginx-deploy-5c9b546997-4w24n
```

- æşççIJNçŤsål'raeIJnæŖgålŭåZlâijŽeGlaLlâLZâzzçŽĐ POD

```
kubectl get pods
```

- åIJ busybox èfZâyłåõzåZlây■erûæśĆ nginx èfZâyłåşşåŖ■çŽĐ service iijNèŖéŨŖæşæIJL'åŖŨåLŖå;şå■

```
wget -O - -q http://nginx:80/
```

4.8 4.8 ælaæNş service ecnåLæÉZd'

- å;şæLŖSäznåLæÉZd' service åZŭäYŤeĠ■æŨŖåzzçñNäYĀäYł service åE■æñæşççIJN service çŽĐåIJŖåĀåŭşçzŖåŖSçŤşåŖYåNŨäZĒ

```
kubectl delete service nginx
```

```
kubectl expose deployment nginx-deploy --name=nginx --port=80 --  
→target-port=80 --protocol=TCP
```

```
kubectl get service
```

- åIJ busybox èfZâyłåõzåZlây■erûæśĆ nginx èfZâyłåşşåŖ■çŽĐ service iijNèŖéŨŖæşæIJL'åz■çĐŭæşæIJL'åŖŨåLŖå;şå■

```
wget -O - -q http://nginx:80/
```

4.9 4.9 labels

äyžāzĀāzĹ Pod ecnåLæÉZd' åŖŖiijNservice åz■çĐŭeĈ;åd' şæ■ççæŖçŽĐeŖĈāzēåLŖæŨŖçŽĐ POD äyŁiijNèfZârşæŸŖ k8s çŽĐ labels èfZâyłæIJZålŭæIēæfIēŖAçŽĐāĀĈ

èĈ;åd' şä;ŖçŤlæāĠç■;æIJZålŭäy■æ■cæIJL podāĀAåIJ k8s
äy■å;Łād' ZâržèşæĈ;åŖŖäzēä;ŖçŤlæāĠç■;iijNä;NāçĈiijŽnodeāĀAservice

- æşççIJN service çŽĐeŖççzEæfæAŖiijNäijZârSçŖŖæāĠç■;éĀL'æNŖ'åZĹ

```
kubectl describe service nginx
```

```

Name:          nginx
Namespace:     default
Labels:        run=nginx-deploy
Annotations:    <none>
Selector:      run=nginx-deploy #
→ ēĹŽäyĹēĀĹ'æŃĹ'āžĹäijžēĠāĹĹēĀĹ'äy run æāĠççĹiijNäyŤāĀijäyž nginx-
→ deploy çŽĎ POD
Type:          ClusterIP
IP:            10.101.149.4
Port:          <unset> 80/TCP
TargetPort:    80/TCP
Endpoints:     10.244.2.4:80 # ā;Ŗ service
→ çŽĎāŖŌçŋŕiijŇā;Ŗ POD āŖŖçŤŖāŖŤāĹĹāĹŽçŋŇāŖšäijžæžt'æŮŖ
Session Affinity: None
Events:        <none>

```

- æšçĹJŇ POD çŽĎæāĠççĹiijNäijŽçĹJŇāĹŖæNēæĹĹ run=nginx-deploy æāĠççĹçŽĎāōžāŽĹiijNēĀŇāžžäyžāĹæŽd'äyĀäyĹ POD āŖŌiijŇāĹŖæĹJŇæŌġāĹāŽĹāĹŽāžžçŽĎāĹŖæĹJŇäyĹçŽĎæāĠççĹäyāijŽāŖŤāŇŮiijŇæĹĀžžæāĠççĹ service āĖŖāŤāĀ

```
kubectl get pods --show-labels
```

NAME	READY	STATUS	RESTARTS	AGE
→ LABELS				
client	1/1	Running	0	21m
→ run=client				
nginx-deploy-5c9b546997-kh88w	1/1	Running	0	8m37s
→ pod-template-hash=5c9b546997,run=nginx-deploy				

- æšçĹJŇ POD çŽĎēŖççŽĖāŖæĀŕiijŇāžŖāŖŖāžžæšçĹJŇāĹŖ POD çŽĎēŖççŽĖāŖæĀŖ

```
kubectl describe deployment nginx-deploy
```

- æāžæŌæāĠççĹēĹĠæžd'iijŇä;ŖçŤĹĹæĹæŇĠāōŽæāĠççĹāŖçġŖæĹŮāŖŇæŮŮēĹĠæžd'āĖŮāij

```
kubectl get pods --show-labels -l run=nginx-deploy
```

- æāĠççĹēĀĹ'æŃĹ'āžĹēŽĖäyēĹŖççŮ

```

āĖŖççžžäyŌ: KEY,KEY KEY=VALUE2,KEY=VALUE2 # -l run,app
çĹ'āĀijāĖŖççž:KEY = VALUE KEY != VALUE # -l run=nginx-
→ deploy,app!=myapp
ēŽĖāŖĹāĖŖççž:KEY in|not in (VALUE1,VALUE2) # -l "release in
→ (canary,bata,alpha) "

```

- æŸçd'žæŇĠāōŽçŽĎæāĠççĹçŽĎāijiijNäyŇēĹæŸçd'žāŖäyĹäyĹæāĠççĹ

```
kubectl get pods --show-labels -L run,pod-template-hash
```

- ```
kubectl label pods client release=canary
```

- ```
kubectl label pods client release=stable --overwrite
```

- ```
kubectl label nodes node2 disktype-
```

- ```
matchLabels: çẗt'æŎëçżŻăôŽéĤōăĂij
matchExpressions:
- ⤴ăŞşăžăŎçżŻăôŽçŽĐēālē;ǃăijRăīēăôŽăžL'ă;ŁçŦĺăăĞç█ēĂL'ăNl'ăžĺ: {key:
- ⤴"KEY", operator:"OPERATOR", value:[VAL1, VAL2, ...]}
    ä;ŁçŦĺ key äÿŎ value ēŁžēaN̄ operator ēēŘčôŮ,
- ⤴ăd'█ăŘĹăĪăäžúçŽĐăL'█ēcńēĂL'ăNl'
    æŞ█ă;IJçņę:
        In, NotIn: âĚű value âĻŰēāĺâfĔēązæIjL'ăĂij
        Exists, NotExists: âĚű value âfĔēązăžŸçł'ž
```

- ```
kubectl label nodes node2 disktype=ssd
```

```
kubectl get nodes --show-labels
```

- ```
kubectl scale --replicas=5 deployment nginx-deploy
```

```
kubectl set image deployment nginx-deploy nginx-deploy=ikubernetes/
↳myapp:v2
```

- æšćIJNæŽt' æŮřčŽDěfGčfNřijNčŽt' āLř 5 äyłāōžāZlāy■ēfRēāNčŽDēTĪJāČRāĖÉčČlæŽt' æŮřāōŇ

```
kubectl rollout status deployment nginx-deploy
```

```
[root@node1 ~]# kubectl rollout status deployment nginx-deploy
Waiting for deployment "nginx-deploy" rollout to finish: 3 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 3 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 3 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 3 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 3 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 4 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 4 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 4 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 4 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 4 out of 5
↳new replicas have been updated...
Waiting for deployment "nginx-deploy" rollout to finish: 2 old
↳replicas are pending termination...
Waiting for deployment "nginx-deploy" rollout to finish: 2 old
↳replicas are pending termination...
Waiting for deployment "nginx-deploy" rollout to finish: 2 old
↳replicas are pending termination...
Waiting for deployment "nginx-deploy" rollout to finish: 1 old
↳replicas are pending termination...
Waiting for deployment "nginx-deploy" rollout to finish: 1 old
↳replicas are pending termination...
Waiting for deployment "nginx-deploy" rollout to finish: 1 old
↳replicas are pending termination...
Waiting for deployment "nginx-deploy" rollout to finish: 4 of 5
↳updated replicas are available...
deployment "nginx-deploy" successfully rolled out
```

- āZdæzŽæš■ā;IJijNäy■æNĜāōZāzzā;TčŽDēTĪJāČRāLŽāyžāyLāyÄäyłčL'LæIJNčŽDēTĪJāČR

```
kubectl rollout undo deployment nginx-deploy
```

āēČædIJēYšæ■cæŽt' æŮřēfGčfNäy■ēcnērČāžēřijNēČčāzLāřséIJĀēēAā■ēāzāāřščzłæĀgæčĀætĪNæL'■ēČ

4.12 4.12 éŽĚçŁd'äd'ŮèóĚŮó

- äŁŉæŤž service çŽĎçŁŚçzIJçśžädNäyž NodePort

```
kubectl edit service nginx
```

```
type: ClusterIP -> type: NodePort
```

- æšĚçIJN service çŽĎäŁæAřijNäRŚçŮřäd'ŽäžEäyÄäyŁ 30982 çnräRč

```
kubectl get service
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	
↪ AGE					
kubernetes	ClusterIP	10.96.0.1	<none>	443/TCP	↪
↪ 15h					
nginx	NodePort	10.105.27.11	<none>	80:30982/TCP	↪
↪ 42m					

- äIJĚŽĚçŁd'äd'ŮèĬäĭçŁŤläžžæĎRçŽĎ node IP äIJřäĬ + çnräRčæĬèèóĚŮó

```
http://172.16.100.101:30982/
```

4.13 4.13 æŮŠæšĚæŮěåŁŮ

- æšĚçIJNäyÄäyŁ pod çŽĎæšŘäyŁäŉžåŽĬçŽĎèŁŘèäNäŮěåŁŮ

```
kubectl logs pod-demo busybox
```

4.14 4.14 èŁďäĚĚ POD äŉžåŽĬ

```
kubectl exec -it pod-demo -c myapp -- /bin/sh
```

CHAPTER 5

äžŤ éĚꞑǫæÿĚåꞤä|£çŦí

apiserver äẏĖăŎěăŦű json æäijäijRčŽDēḐdæẑŘăőŻăzL'ījN̄yāml
æäijäijRăōŻăzL'ăRŘă;ŹčŽDÉē■;őæÿĚă■TiiJNapiserwer âRrêGlâLlârEăĔūē;ñă■câyž
json æäijäijRīijNèĀÑăRŌăE■ēfZëąNăL'gëąNăĂĆ

5.1 5.1 ĄŔréĚ■çjőçŽĐǎřžèśą

- ǎRǝTlǝTǾæžRæÿĖǎ■TǝĖ■çjǝçŽǾǎrzǝsǎ

workload	Deployment	StatefulSet	DaemonSet	Job	CronJob
Service	Ingress				
Volume	CSI				
ConfigMap	Secret				
DownwardAPI					
Namespace	None	Role	ClusterRole	RoleBinding	ClusterRoleBinding
HPA	PodTemplate	LimitRange			

5.2 5.2 éĚ■çjőæÿĚǎ■ȚçžĎæŁŘ

- éĚ■ç;őăÿĚă■ȚçȚDăĹRěĆlăLEjijNăd'gěĆlăLEçȚDăzŘă;ĚçȚlěĚ■ç;őăÿĚă■ȚăŮzaijŘălěăĹZăzž

```
apiVersion
  # äzë "group/version" ä;çâijRæŃGæYÖiijŃèfZäyłårzésaąsďäžŌašłäył_
  ↪API çzĎiijŁçL'łæIJŃiijL'
```

(continues on next page)

(continues on next page)

(continued from previous page)

```

kind:
  # èťĎæžŘčśzǎłńiijŇæăĜèőřǎĹžǎžžăžĂăžĹčśzǎđŇčŽĎètĎæžŘ
metadata:
  # âĚČæŤřæ■ôâĚĚéČĹæŸřâtŇăěŮčŽĎă■Ůæôt
  # _
  →ăőŽăžĹ'ăžĚètĎæžŘǎřzèsaçŽĎăŘ■çğřăĂăăŚ;ăŘ■çĹ'žéŮt'iiĵĹk8sçžğǎĹŋçŽĎăŸ■æŸřçşzçzŞçŽĎ
spec:
  # _
  →èĝĎèŇČăőŽăžĹ'ètĎæžŘăžŤèřěăŇěăIJĹ'ăžĂăžĹæăŮčŽĎçĹ'žæĂĝiijŇă;ĪéĪăăŌĝǎĹŮăžĹçăőăĪç
  # âőČæŸřçŤĹæĹŮăőŽăžĹ'çŽĎæĹ'ĂăIJSæIJŽăžĚètĎæžŘçĹŮăĂă
status:
  # æŸççd'žètĎæžŘçŽĎă;ŞăĹ'■çĹŮăĂăiijŇk8s_
  →ăřsæŸřçăőăĪă;ŞăĹ'■çĹŮăĂăăŘŞçŽôăăĜçĹŮăĂăăŮăéŽŘéĪăèĹŚăžŌèĂŇæzaëŮşçŤĹæĹŮăIJSæIJŽ
  # âőČæŸřăŘĹèřzçŽĎiijŇăžçèăĹăžĚètĎæžŘă;ŞăĹ'■çĹŮăĂă

```

- èŌŮăŖŮăĚĹéČĹçŽĎ api çĹĹæIJŇ

```
kubectl api-versions
```

- èŌŮăŖŮăĚĹéČĹçŽĎ api èťĎæžŘǎřzèsa

ăžŌăĚĚăőzǎŖřăžèçIJŇăĹřăŸĂăžŽçijĹ'ăĚŽiijŇæŮžă;£ăĹŚăžŇæŮëăŸŸăŚ;ăžd'ăŘŌçôĂăĚŽ

```
kubectl api-resources
```

```

kubectl get po          # æŞěçIJŇpod
kubectl get deploy      # æŞěçIJŇdeployment
kubectl get svc         # æŞěçIJŇservice
kubectl get cm          # æŞěçIJŇ configmap
...

```

5.3 5.3 èŌŮăŖŮăæŸĚă■ŤăŸôăĹ'

- æŞěçIJŇ k8s æŞŖăŸĹăĚĚçĹǎřzèsaçŽĎĚĚ■çĹăŸĚă■ŤăăijăijŖiijŇăžŤèřěăŇěăŖŇăŇăŞĹăžŽă■ŮæôtĹiijŇă;£ăĹæŸççd'žă■ŮæôtçŽĎăăijăijŖăŸôăĹ'ăĹæăĹ'

```

kubectl explain pods
kubectl explain pods.metadata

```

5.4 5.4 æŸĚă■ŤăŞžæĹJŇæăijăijŖ

- âőŽăžĹăŸĂăŸĹètĎæžŘæŸĚă■Ť

- ä;£çTÍ -o ǎRĆæTṛæIēāNǧǎoŽǎrfzēsæTṛæ■ōçŽDè;ŠǎGžæāijāijRīijNǎ;£çTÍ –dry-run
æIēætNērTǎĀğæLğēāNǎyĀǎylæNǧǎzd'tijNǎōČǎyd'ǎylçzŠǎRĻètūǎIēīijNǎrsǎRfǎzēēĀŽē£ĠǎS'ǎzd'ǎĻ
yaml æāijāijRÉĒ■ç;ōæŪĠǎžūǎžE -o yaml –dry-run

- $\text{ál} \text{Z} \text{áz} \text{z} \text{è} \text{t} \text{D} \text{æ} \text{z} \text{R} \text{æ} \text{y} \text{Ě} \text{ǎ} \text{■} \text{T} \text{ä} \text{y} \text{■} \text{ç} \text{Z} \text{Ď} \text{è} \text{t} \text{D} \text{æ} \text{z} \text{R} \text{i} \text{i} \text{j} \text{N} \text{è} \text{f} \text{Z} \text{æ} \text{u} \text{ǎ} \text{ł} \text{Z} \text{áz} \text{z} \text{ç} \text{Z} \text{D} \text{ä} \text{y} \text{z} \text{è} \text{c} \text{y}$ POD
 $\text{i} \text{i} \text{j} \text{N} \text{æ} \text{s} \text{ǣ} \text{I} \text{J} \text{L} \text{'} \text{æ} \text{Ō} \text{ǵ} \text{ǎ} \text{L} \text{ú} \text{Ǻ} \text{Z} \text{í} \text{ç} \text{ō} \text{a} \text{ç} \text{R} \text{E} \text{i} \text{i} \text{j} \text{N} \text{æ} \text{L} \text{'} \text{Ä} \text{ä} \text{z} \text{é} \text{Ǻ} \text{Ł} \text{æ} \text{Z} \text{'d} \text{'} \text{ǎ} \text{R} \text{Ō} \text{ä} \text{y} \text{■} \text{ä} \text{i} \text{j} \text{Z} \text{è} \text{G} \text{ł} \text{Ǻ} \text{Ł} \text{é} \text{G} \text{■} \text{Ǻ} \text{z} \text{z} \text{i} \text{i} \text{j} \text{N} \text{æ} \text{L} \text{R} \text{ä} \text{y} \text{z} \text{è} \text{G} \text{ł} \text{ä} \text{y} \text{z} \text{Ǻ} \text{i} \text{R}$
 POD

```
kubectl delete -f pod-demo.yaml
```

5.8 5.8 apply

apply

```
kubectl apply -f pod-demo.yaml
```

CHAPTER 6

Pod

6.1 pods.metadata

6.1.1 labels

- labels are key-value pairs that are attached to a pod

```
labels:  
  app: myapp  
  tier: frontend
```

6.2 pods.spec

6.2.1 nodeName

- nodeName is the name of the node that the pod is scheduled on

```
apiVersion: v1  
kind: Pod  
metadata:  
  name: pod-deme  
  namespace: default  
  labels:  
    app: myapp  
    tier: frontend
```

(continues on next page)

6.2.3 restartPolicy

Always: Always restarts the container if it fails. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

OnFailure: Restarts the container if it fails. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

Never: Never restarts the container if it fails. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

```

# Always restart policy
apiVersion: v1
kind: Pod
metadata:
  name: always-restart
spec:
  restartPolicy: Always
  containers:
  - name: always-restart
    image: busybox
    command: ['sh', '-c', 'sleep 30']

```

```

# OnFailure restart policy
apiVersion: v1
kind: Pod
metadata:
  name: on-failure-restart
spec:
  restartPolicy: OnFailure
  containers:
  - name: on-failure-restart
    image: busybox
    command: ['sh', '-c', 'sleep 30']

```

6.2.4 hostNetwork

hostNetwork: true: The container shares the host's network namespace. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

6.2.5 hostPID

hostPID: true: The container shares the host's PID namespace. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

6.2.6 containers

kubectl explain pods.spec.containers

containers: List of containers that will run on this pod. Each container is a struct of type Container. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

- containers: List of containers that will run on this pod. Each container is a struct of type Container. The container is restarted after it has terminated. The container is restarted after it has terminated. The container is restarted after it has terminated.

args	
command	
env	
envFrom	
image	
imagePullPolicy	
lifecycle	
livenessProbe	
name	
ports	
readinessProbe	
resources	
securityContext	
stdin	
stdinOnce	
terminationMessagePath	
terminationMessagePolicy	
tty	
volumeDevices	
volumeMounts	
workingDir	

- `podSpec`

```

apiVersion: v1
kind: Pod
metadata:
  name: pod-deme                                # pod spec
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
    - name: myapp                                # myapp
      image: ikubernetes/myapp:v1               # ikubernetes/myapp:v1
      imagePullPolicy: IfNotPresent             # IfNotPresent
      ports:                                     # ports
        - name: http
          containerPort: 80
    - name: busybox
      image: busybox:latest
      command:
        - "/bin/sh"
        - "-c"
        - "sleep 10"

```

6.2.6.1 imagePullPolicyäÿÑè;ç■ÜçTě

- imagePullPolicy éŤIJăČŘěŎůăŔŮçŽĎç■ÜçTěiijÑěřěğĀiijŽkubect1 explain pods.spec.containers

```
Always          # æĀzæŸřăžŎăžšăžšăÿÑè; ;
Never           #
→ăžŎăÿ■ăÿÑè; ;iijÑăIJňăIJřăIJL'ăřšçŤlíijÑăšăæIJL'ăřšăd'set'ě
IfNotPresent    #
→ăęČăđIJăIJňăIJřă■ŸăIJlăřšçŽt'ăŎěă;ŁçŤlíijÑăęČăđIJăÿ■ă■ŸăIJlăřšăÿÑè; ;
```

ăęČăđIJăăĜç■ăŸř latest éČčăžŁăĹŽăĜŇçžŁăžŎăžšăžšăÿÑè; ;

6.2.6.2 ports çñŕăŔčăŁæAŕ

- ports âŏŽăžŁ'ăŏžăŽlăŁăŤŕéIJšçŽĎiijÑěřěğĀiijŽkubect1 explain pods.spec.containers.ports

ăIJlă■d'ăd'DăŽŕéIJšçŽĎçñŕăŔčăŔŕăÿžçšçžšăŔŔăŁăæIJL'ăĚšăŏžăŽlčŽĎç;ŠçzIJěŁăŎěçŽĎăŁăæAŕfi
0.0.0.0 âIJŕăĹăçŽĎçñŕăŔčăĈ;ăŔŕăžăžŎč;ŠçzIJěŎŁéŮŏ

```
ports:          # âŏŽăžŁ'ăÿd'ăÿlçñŕăŔčăŔŕăžšăăÿĀăÿl http
→ăÿĀăÿl https
- name: http     #
→âŏŽăžŁ'ăĚšăÿlçñŕăŔčăçŽĎăŔ■çğŕiijÑăŮžă;ŁăĹŇçžĎăŔŕăžšăăŔŮăijŤçŤl
  containerPort: 80      # çñŕăŔčăŔŮ
- name: https    # æŮžă;ŁăijŤçŤlçŽĎăŔ■çğŕ
  containerPort: 443     #
→ăĚšăÿlçñŕăŔčăŔŮăžĚăžĚăŸŕètŮăĹŕăŁăæAŕçŽĎă;IJçŤlíijÑăŮžă;ŁăšĚçIJŇăŠŇă;ŁçŤlăŔ■çğŕ
```

6.2.6.3 env äijăéĂšçŎŕăčČăŔŸéĜŔ

ăIJlăăŏžăŽlăÿ■ěŎůăŔŮ POD çŽĎăŁăæAŕ

```
ăŔŕăžăžă;ŁçŤlçŎŕăčČăŔŸéĜŔ
ăŔŕăžăžă;ŁçŤl downwardAPI
https://kubernetes.io/zh/docs/tasks/inject-data-application/
→downward-api-volume-expose-pod-information/
```

6.2.6.4 command ENTRYPOINT

- command âŏŽăžŁ'ăŏžăŽlăŁăŤŕăŇçŽĎçĹăžŔiijÑěřěğĀiijŽ

ăÿĀăÿl entrypoint array ěĂŇ command âŔŕăĹlçŽĎçĹăžŔăŸŕăÿ■ăijŽěŁŔăăŇăIJl Shell
ăÿ■çŽĎiijÑăęČăđIJăČšĚăĀĚŔăăŇăIJl Shell äÿ■éIJăĚĀĚĜlăŮšăăŇăĚŽiijÑăęČăđIJăšăæIJL'ăŔŔă;ŽĚŁŽăÿ
docker éŤIJăČŔăÿ■çŽĎ ENTRYPOINTăĂČ

6.2.6.5 args CMD

- args aŘŠ command äijäéĀŠāRCæTřčŽĎ

æĈæđIJäjäšæI JL'ăōŽăZL' args èĀNéTĪJăĈRăy■ăRĹă■YăIJĪ ENTRY-
POINT æŇĠăzd'ăŠŇ CMD æŇĠăzd'rijNéĈčăZĹéTĪJăĈRèĠăũščŽĎ CMD
ăŖĒä;IJăyžăRĈæTřäijäéĀŠčžŽ ENTRYPOINTăĀĈæĈæđIJăL'NăĹăæŇĠăōŽăžĒ args
éĈčăZĹéTĪJăĈRăy■čŽĎ CMD â■Ůăōtăy■ăĒ■ă;IJăyžăRĈæTřèĤŽăqNăijäéĀŠăĀĈ

æĈæđIJăIJĪ args äy■ăijTčTĹăžĒăRŸéĠRĭijNăĹŽéIJăĸăĀ;ĤčTĪ \$(VAR_NAME)
æĹăäijTčTĹăyĀăyĹăRŸéĠRĭijNăĸæĈæđIJăy■ăĈšăIJĹéĤŽéĠNéĤŽăqNăŠ;ăzd' æŽĒă■ĸrijNéĈčăZĹăRăžă
\$\$(VAR_NAME)rijNé;ăZL'ăŖŌăIJăăōZăŽĹăĒă;ĤčTĹăĀĈ

6.2.6.6 annotations æšlèġčăĚăĀăĀŖ

annotations äyŎ label äy■ăŖNčŽĎăIJăŤzăŎrijNăōĈăy■ăĸ;čTĹăžŎăŇSéĀL'èĤăžŖăŖžèšqijNăžĒ

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-deme
  namespace: default
  labels:
    app: myapp
    tier: frontend
  annotations:
    æšlèġčăĚăĀăĀŖ
    kaliarch/created-by: "xuel"
    æŭžăĹăăĒăĹăĀijăŖžčŽĎèĤăžŖăæšlèġč
spec:
  containers:
    - name: myapp
      image: ikubernetes/myapp:v1
      imagePullPolicy: IfNotPresent
```

6.2.6.7 POD çTšăŠjăŚĹæIJš

- äyĀèĹŋçĹŭăĀĀ

PendingiijžăũščžŖăĹŽăžžă;ĒăŸŖăšæI JL' éĀĈăŖĹèĤŖăqNăōĈčŽĎéĹĈčĈziijNăũščžŖèŖĈăžèiij
RunningiijžèĤŖăqNčĹŭăĀĀ
FailediijžăŖăăĹăđ'sèt'ě
SucceediijžăĹŖăĹSiiijNéĤžăyĹçĹŭăĀĀăĹĈS■
UnkowniijžăI JLčSèčŽĎçĹŭăĀĀiijNăĸæĈăđIJ Apiserver äyŎ kubelet
æĀžăĹăăđ'sèt'ăĹŽăiijžăđ'ăžŖăŖžèšăyĹçĹŭăĀĀ

- ăĹŽăžž POD éŸŭăōt


```

    çTlæLûçŽDāLZāzžèrûæšCæRŘāžd'çžZ apiserver iijNèĀŇ apiserver
    äijŽārEèrûæšCçŽDçŽōæāGçLûæĀAäflāYāIJl etcd äy■iijNèĀŇāRŌ apiserver
    äijŽèrûæšC schedule èfZèqNèrČāžèiijNāzūāyTæLĒrČāžèçŽDçzŠædIJæŽt'æŪrāIJl
    etcd çŽD pod çLûæĀAäy■iijNèŽRāRŌāyĀæŪèäflāYāIJl etcd äy■iijNāzūāōNæLŔ
    schedule æŽt'æŪrāRŌçŽōæāGèLÇçCžçŽD kubelet āršäijŽāzŌ etcd
    çŽDçLûæĀAāRŸāNŪā; ŪçšèæIJL'æŪrāzžāLāçzŽèGĥāūsiiijNæLĀāžèæ■d'æŪūäijŽæNĕāLŕçTlæLûæLĀāyN.
    PODiijNāèCædIJāLZāzžæLŔāLšæLŪèĀĒād'sèt'èiijNāLZārEçzŠædIJāRŠāŽdçžZ apiserver
    iijNapiserver āE■æñāäflāYāIJl etcd äy■āĀC
  
```

6.2.6.8 livenessProbe āYæt'zæĀgæŌcætN

èrèççEèġAiiijZkubectl explain pods.spec.containers.livenessProbe

- livenessProbe / readinessProbe æYr k8s äyd'äylçTšāS;āŚlæIJšiiijNèfZāyd'äylçTšāS;āŚlæIJšéČ;āRfrazè

```

livenessProbe #
→æNĠçd'žāōžāZlæYrāRĕæ■čāIJlèfRĕāNāĀCāèCædIJā■Yæt'zæŌcætNād'sèt'èiijNāLZā;Īæ■Ō
→restartPolicy ç■ŪçTĕæĪèèfZĕāNéG■āRr
readinessProbe #
→æNĠçd'žāōžāZlæYrāRĕāGĒād'Gāè;æIJ■āLæèrûæšCāĀCāèCædIJāRšçzLæŌcætNād'sèt'èçnrçCž
→Pod āNžéĒ■çŽDæL'ĀæIJL' Service çŽDçnrçCžäy■āLæéZd'èrĕ Pod çŽD IP
→āIJrāĪĀ
  
```

- livenessProbe / readinessProbe āRŕçTlçŽDæŌcéŚLāŠNæŌcéŚLçL'zæĀgiiijNæŌcéŚLāRĪèČ;āōžāzL'äy

```

exec #
→āIJlāōžāZlāEĒæL'ġèāNæNĠāōžāS;āzd'āĀCāèCædIJāS;āzd'ēĀĀāGžæŪūèfTāždçāĀäyž
→0 āLZèōd'äyžèrLæŪ■āLŔāLŠāĀC
tcpSocket # āRzæNĠāōžçnrāRçäyLçŽDāōžāZlçŽD IP āIJrāĪĀèfZĕāN TCP
→æČĀæSĕāĀCāèCædIJçnrāRçæL'ŠāijĀiijNāLZèrLæŪ■ècnèōd'äyžæYrāLŔāLŠçŽDāĀC
httpGet # HTTP GET
→èrûæšCāNĠāōžçnrāRçāŠNèurā;DäyLçŽDāōžāZlāĀCāèCædIJāS■āžTçāĀad'ġāžŌç■L'āžŌ200
→äyTārRāžŌ 400iijNāLZèrLæŪ■ècnèōd'äyžæYrāLŔāLŠçŽDāĀC
  
```

```

failureThreshold #
→æŌcætNāGāñæL'■āLd'āōžäyžæŌcætNād'sèt'èiijNéžYèōd'äyž 3 æñāāĀC
periodSeconds # ærRæñææŌcætNāŚlæIJSçŽDèŪt'èŽTæŪūèTĒāĀC
timeoutSeconds #
→æRŔæñææŌcætNāRŠāGžāRŌç■L'ā;EçzšædIJçŽDèŪĒæŪūæŪūèŪt'iijNéžYèōd'äyž
→1 çġšāĀC
initalDelaySeconds #
→āIJlāōžāZlāRŕāLlāRŌāzūèfSād'ŽāžĒāŌžèfZĕāNæŌcætNiiijNéžYèōd'äyžāRŕāLlāōžāZlāRŌçN
  
```

- ä;ççTl exec æŌcéŚLiiijNāōdēfNçzšædIJāžTèrēäyž 39 çġšāRŌ POD æYçd'ž ERROR
iijNā;Eäy■èGĥāLlĕG■āRr POD

```

apiVersion: v1
kind: Pod
metadata:
  
```

(continues on next page)

(continued from previous page)

```

name: execlive
namespace: default
labels:
  app: myapp
  tier: frontend
spec:
  containers:
  - name: busybox
    image: busybox
    command:
      - "/bin/sh"
      - "-c"
      - "touch /tmp/healthy; sleep 30; rm -rf /tmp/healthy; sleep_
→3600"      # ăĹŽăžžăÿĂăÿĹăŮĜăžŭĈ■Ĺ'ăĭĚ 30_
→ĉġšĭiĵNèĽŽăÿĹăŮŭéŮt'æŌcéšĹăžTèřěæŸræĹŖăĹSĉŽĎiĵN30_
→ĉġšăŖŌăĹŽăd'sèt'ě
      livenessProbe: #_
→ăŏžăžĹĭĈŽĎă■Ÿæt'zæĀġæĈĂætNĭĵNăĉĆăđIJăd'sèt'ěăĹŽăNĹ'ĉĚġ_
→restartPolicy ĉ■ŮĈTĕăĭééĜ■ăŖŕ POD
      exec: # exec_
→ĉśžăđNăŌcéšĹiĵNèĽŽăĚăŏžăžĹăĹ'ġèaNăÿĂăĭaăś;ăžd'
      command: ["test", "-e" ,"/tmp/healthy"] #_
→æĹ'ġèaNĉŽĎăś;ăžd'ăÿžætNèrTæŮĜăžŭă■ŸăIJĹăĀġ
      initialDelaySeconds: 2 #_
→ăŏžăžĹăŖăĹăĹăŖŌăžŭèĽSăd'ŽăžĚèĽŽèaNăŌćætN
      periodSeconds: 3 #_
→æŖŖăŋăæŌćætNăŖăĹăIJŖĉŽĎéŮt'éŽĬăŮŭéTĽăÿž 3 ĉġš
      failureThreshold: 3 # 3_
→æŋăăd'sèt'ěăŖŌăĹŽăĹd'ăŏžăÿžăŏžăžĹăŌćætNă■Ÿæt'zæĀġăd'sèt'ě
      restartPolicy: Never #_
→ăĭşæŌćætNăĹŖăŏžăžĹăd'sèt'ěæŸŖăŖĕéĜ■ăŖŕ POD

```

- äĭĤĉTĬ httpGet æŌcéšĹiĵNăŏđéĹNĉzşæđIJăžTèřăd'ġĉžĕ
40 ĉġšăŖŌăŌćætNă■Ÿæt'zæĀġăd'sèt'ëĭĵNèĜĹăĹéĜ■ăŖŕ
POĎĭĵNĉŋăÿĂăŋăĕĜ■ăŖŕăĭĴĉŋNă■şĕĽŽèaNĭĵNéŽŖăŖŌăŸŖ 30
ĉġšĉŽĎăă■ĈŽŖăĹŖ 300 ĉġšăĀĈ

```

apiVersion: v1
kind: Pod
metadata:
  name: httpgetlive
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
  - name: nginx
    image: ikubernetes/myapp:v1

```

(continues on next page)

(continued from previous page)

```

ports:
  - name: http
    containerPort: 80
  - name: https
    containerPort: 443
livenessProbe:
  httpGet:
    path: /error.html
    port: http
    httpHeaders:
      - name: X-Custom-Header
        value: Awesome
    initialDelaySeconds: 15
    timeoutSeconds: 1
    restartPolicy: Always

```

6.2.6.9 readinessProbe

Tomcat service

```

apiVersion: v1
kind: Service
metadata:
  name: tomcat
spec:
  selector:
    app: tomcat
  ports:
    - port: 80
      targetPort: 8080
    - port: 443
      targetPort: 8443
  type: ClusterIP

```

- livenessProbe / readinessProbe

```

livenessProbe:
  httpGet:
    path: /
    port: 80
  initialDelaySeconds: 15
  timeoutSeconds: 5
  successThreshold: 1
  failureThreshold: 3
  restartPolicy: Always
readinessProbe:
  httpGet:
    path: /
    port: 80
  initialDelaySeconds: 15
  timeoutSeconds: 5
  successThreshold: 1
  failureThreshold: 3
  restartPolicy: Always

```

- httpGet
 - path: /
 - port: 80
 - initialDelaySeconds: 15
 - timeoutSeconds: 5
 - successThreshold: 1
 - failureThreshold: 3
 - restartPolicy: Always

```

apiVersion: v1
kind: Pod
metadata:
  name: httpgetread
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
    - name: nginx
      image: ikubernetes/myapp:v1
      ports:
        - name: http
          containerPort: 80
        - name: https
          containerPort: 443
      livenessProbe:
        httpGet:
          path: /error.html
          port: http
          httpHeaders:
            - name: X-Custom-Header
              value: Awesome
          initialDelaySeconds: 15
          timeoutSeconds: 1
          restartPolicy: Always

```

- `index.html` äzëëgëäRŠäřšçzłæÄgæŒcéŠŁčŽDæčÄætŇ

```

kubectl exec -it httpgetread -- /bin/sh
$ rm -f /usr/share/nginx/html/index.html

```

- çŞædIJeŹäyŁ POD çŽD READY çŁúæÄÄäüšçzłæÄgæŒcéŠŁčŽDæčÄætŇ service äy■äijŽäE■ërČäžëäŁřëŹäyŁëŁČçČzäžE

```

[root@node1 ~]# kubectl get pods -w
NAME                                READY    STATUS    RESTARTS   AGE
httpgetread                         0/1      Running   0           2m50s

```

- äIJläöžäŽÍäEĖäE■äŁZäžäyÄäyŁæŮGäžüijNäžëëgëäRŠäřšçzłæÄgæŒcéŠŁčŽDæčÄætŇ

```
kubectl exec -it httpgetread -- /bin/sh
$ echo "hello world" >>/usr/share/nginx/html/index.html
```

- `çŽšædIJeŁŻäyŁ POD çŽĐçŽĐ READY çŁúæĀĀũščžŘcijŮčlŇăřščžłăžEĭijŇæ■d' æŮŮ`
service äijŽerČăžęĀŁrëŁŻäyŁèŁĆçĆzăžE

```
[root@node1 ~]# kubectl get pods -w
NAME                                READY    STATUS    RESTARTS   AGE
httpgetread                         1/1      Running   0           8m15s
```

6.2.6.10 lifecycle çŤšăŚ;ăŚÍæIJséŠ'ă■Ř

èřęèġAĭijŽkubectĭ explain pods.spec.containers.lifecycle

```
postStart                          #
→âIjĬăôžăŹĬăŘrăĬăŖŎčŇŇă■şæL' ġëaŇçŽĐăŚ;ăzd' iijŇăęĆædIJeŁŻäyŁæş■ă; IJăd' sét' ěăžEĭi
→restartPolicy æĬëăEşăôžăŸrăŘęéĢăŘr
preStop                            # âIjĬăôžăŹĬčžĬă■căL'■čŇŇă■şæL' ġëaŇçŽĐăŚ;ăzd'
```

- `postStart / preStop çŽĐăşžæIjňă;ŁçŤĬ`

```
apiVersion: v1
kind: Pod
metadata:
  name: lifecycle-demo
spec:
  containers:
  - name: lifecycle-demo-container
    image: nginx

    lifecycle:
      postStart:
        exec:
          command: ["/bin/sh", "-c", "echo Hello from the postStart_
→handler > /usr/share/message"]
      preStop:
        exec:
          command: ["/usr/sbin/nginx", "-s", "quit"]
```

PODæŎġăĬăŹĬ

æŎġăĬăŹĬçőăçŘEçŽĐ POD âŖrăžęăőđçŎřijŇëĢĬăĬčzt' æŁd' POD
ăĬŖăIjňæŤŕëĢŖĭijŇăőČëČ;ăőđçŎř POD çŽĐæĬĬăőžăŇŇcijĬăőžĭijŇă;EæŸŕăy■ëČ;ăőđçŎřæžŽçŽĐëČčăyĬă

7.1 ReplicaSet

7.1 ReplicaSet

ðeðgAiiZkubectl explain replicaset

- æyĖaTègDèÑČ

```
apiVersion <string>      # api  çL'ÍæIJñåRûii jÑäyĖèĹñäyž apps/v1
kind <string>             # _
→ ètDæžRçsžåĹñii jÑæăĖèõřåĹžåžžăžĂăžĹçsžăđNçŽDètDæžR
metadata <Object>        # POD  åĖČæTřæ■ō
spec <Object>            # åĖČæTřæ■ō
```

7.1.1 replicaset.spec ègDèÑČ

1. replicas åĹræIJñæTřéGRii jÑæÑĖåõŽäyĖäyĖæTřå■Ů
2. selector æăĖç■éĀĹæÑĹ'åŽĹii jÑåRřäžă;çTĹ matchLabelsāĀmatchExpressions
äyð'çĖ■çsžăđNçŽDèĀĹæÑĹ'åŽĹæĹéĀĹäy■çŽōæăĖ POD

```
matchLabelsii jŽçŽt' æŌëçžŽăõŽéTōāĀij
matchExpressionsii jŽăŞžăžŌçžžăõŽçŽDèaĹèç;ăijRæĹæăõŽăžĹ' ä;çTĹæăĖç■éĀĹæÑĹ'åŽĹii
→ {key: "KEY", operator: "OPERATOR", value: [VAL1, VAL2, ...]}
   ä;çTĹ key äyŌ value èĹžæāÑ operator _
→ èĹRçŌŮii jÑăđ' ■ăRĹæĹăžăžüçŽDæĹ' ■éçnéĀĹæÑĹ'
   æŞ■ă;IJçñëii jŽ
   InāĀNotInii jŽăĖŮ value åĹŮèaĹăĹĖĖæazæIJĹ'ăĀij
   ExistsāĀNotExistsii jŽăĖŮ value åĹĖĖæazăyžçĹ'ž
```

3. template POD
 ârzèšajijNèfZäylârzèšâRlâNĖâRñäžE pod.metadata âŠN pod.spec äyd' éCĭâLEāĂĈ

7.1.2 7.1.2 æyĖā■Tçd'žăĭN

```
apiVersion: apps/v1
kind: ReplicaSet
metadata:
  name: myrs
  namespace: default
spec:
  replicas: 2
  selector:
    matchLabels:
      app: myapp
      release: canary
  template:
    metadata:
      name: myapp-pod      # èfZäyĭâEûâôðæšaçTĭiijNâžâäyžâĬžâžžçŽĎ
    labels:
      app: myapp           # æăĜç■çäyĖâôžèèAçñèâĬ replicaset
      release: canary
    spec:
      containers:
        - name: myapp-containers
          image: ikubernetes/myapp:v1
          ports:
            - name: http
              containerPort: 80
```

7.2 7.2 DeploymentæŒğâĬúâŽĬ

Deployment éĂžèfĜæŒğâĬú ReplicaSet æĭèâôðçŒrâĬšèĈ;ĭijNéŽd' äžEæTŕæNĀ Repli-
 caSet çŽĎæĬĬ' çijĬ' âôžæĎRâĎ' ŪĭijNèfY æTŕæNĀæžŽâĬæŽt' æŪrâŠNâŽðæžŽç■ĬĭijNèfY æRĬäĭŽâžEâçræY

Deployment âĬĬæžŽâĬæŽt' æŪræŪûâĂžĭijNéĂžèfĜæŒğâĬúâĎ' ŽäyĬ ReplicaSet æĭèâôðçŒrĭijNReplicaSet âĬĬæŒğâĬúâĎ' ŽäyĬ PODĭijNâĎ' ŽäyĬ ReplicaSet çŽyâ;ŠâžŒâĎ' ŽäyĭâžTçTĭçŽĎçĬĬæĬĬâĂĈ

- æyĖā■TèğĎèNĈĭijNèrèèğAĭijŽkubectĭ explain deployment

```
apiVersion <string>      # apps/v1
kind <string>             #
    æĭèâôðçŒrâĬšèĈ;ĭijNâžâäyžâĬžâžžçŽĎ
```

(continues on next page)


```
kubectl rollout status deployment myapp-deploy
```

Waiting **for** deployment "myapp-deploy" rollout to finish: 1 out of 5 **new** replicas have been updated...

ç■L'â¿ĚćĹ;š"myapp-deploy"éĆĹ;šăŃæĹŔ: **5**äÿłæŮŕăL'ŕæIJňăÿçŽĐlăÿłăŭšæŽt'æŮŕ...

- æĈæđĲĖĜŚăÿĲéŽĂæšæĲĲ'ėŮŃėćŸĲĲŃéĈăžĹçžğçž■ăŔŕăžěă;ŕçŤĹçžğçž■æŽt'æŮŕçŽĐăŚ;ăzd'

```
kubectl rollout resume deployment myapp-deploy
```

7.2.5 7.2.5 æŽt'æŮŕç■ŮçŤĕ

- æĲĂăđ'găÿ■ăŔŕçŤĹăÿž 0 ĲĲŃæŽt'æŮŕæŮŭăĂŽăŔŕăžěăÿt'æŮŭėŭĚăĜžĹăÿĲ

```
kubectl patch deployment myapp-deploy -p '{"spec":{"strategy":{"rollingUpdate":{"maxSurge":1,"maxUnavailable":0}}}}'
```

7.2.6 7.2.6 âĖşăžŎăŽđæžŽ

1. rollout undo æŸŕăŽđæžŽçŽĐăŚ;ăzd'ĲĲŃéžŸėŃđ'æžŽăŽđăÿĹăÿĂçĹĹæĲĲ

```
kubectl rollout undo deployment myapp-deploy
```

2. æšĕçĲĲŃăŔŕăžěăŽđæžŽçŽĐçĹĹæĲĲ

```
kubectl rollout history deployment myapp-deploy
```

2. rollout undo æŃĜăŃŽăŽđæžŽçŽĐçĹĹæĲĲ

```
kubectl rollout undo deployment myapp-deploy --to-revision=2
```

3. æšĕçĲĲŃă;ŖăĹ■çŽĐăŭěă;ĲçĹĹæĲĲ

```
kubectl get rs -o wide
```

7.3 7.3 DaemonSetæŎğăĹŭăŽĲ

- æÿĚă■ŤĕğĐĕŃĈĲĲŃĕŕĕğğA kubectl explain daemonset

```
apiVersion <string> # apps/v1
```

```
kind <string> #
```

→ĕŤĐăžŔçśžăĹŲĲĲĲŃăĜĕŃŕăĹŽăžžăžĂăžĹçśžăđŃçžĐĕŤĐăžŔ

(continues on next page)

(continued from previous page)

```

name: filebeat-daemonset
namespace: default
spec:
  selector:
    matchLabels:
      app: filebeat
      release: stalbe
  template:
    metadata:
      labels:
        app: filebeat
        release: stalbe
    spec:
      containers:
        - name: filebeat
          image: ikubernetes/filebeat:5.6.5-alpine
          env:
            - name: REDIS_HOST
              value: redis.default.svc.cluster.local
            - name: REDIS_LOG_LEVEL
              value: info

```

7.3.3 7.3.3 æŽt'æŰř

- æŽt'æŰř filebeat-daemonset æŹäyŧ daemonset æŎğåŁŭåŹŭyŇçŽĎ filebeat
åőžåŹŭçŽĎéŧŭĀčŔ

```

kubectl set image daemonsets filebeat-daemonset_
→filebeat=ikubernetes/filebeat:5.6.6-alpine

```

CHAPTER 8

åĖñ Service éĖ■ç|őæÿĖå■Ţ

Service äyž POD æŎǵǎĹúǎŽĭæŎǵǎĹúčŽĎ POD éŽĚç; d' æŘŘä; ŽäyÄäyĽǎŽžǎŎŽčŽĎēŎéŮŏçnrčČzĭjŇS
čŽĎäüēä;IJēŸä;ĭetŮǎŹŎ K8s äy■čŽĎäyÄäyĽēŽĎäzŮĭjŇǎřsæŸf CoreDNS ĭjŇǎŏČǎřĚ Ser-
vice ħIJřǎĬÄæŘŘä; ŽäyÄäyĽǎ§šǎŘ■ēğčæđŘǎĀĆ

8.1 8.1 Service åũěä;JæíaåijR

1. userspace: 1.1 āžŅāĻ■çĻĹæIJñ
2. iptables: 1.10 āžŅāĻ■çĻĹæIJñ
3. ipvsiižĹ1.11 āžŅāŖŌçĻĹæIJñ

8.2 8.2 Service část

čšžďň	ä JčTí
ClusterIP	ézYèòď'āĀijījNāLēēĚ■äyÄäyŁ Service ç ŚçzIjçŽĐāIJřāĀĀijNāzĚçTlāžŎéŽEçç,ď'āĚĚćléĀž
NodePort	āęČæďIJēIJĀēęAéŽEçç,ď'ād'ŮéČléòŁēŮōījNāRřāžēä;řçTlēŁZäyŁčšžďň
External-Name	æŁŁēŽEçç,ď'ād'ŮéČłčŽĐæIJ■āŁāījŤāĚēāŁřēŽEçç,ď'āĚĚćłījNāŮžä;řāIJlēŽEçç,ď'āĚĚćłā
LoadBalancer	K8S āũēä;IJāIJlāžŚçŎřāčČäy■ījNērČçTlāžŚçŎřāčČāŁZāžžè'şè;ĭāİGèaaāZÍ

8.3 8.3 etDæžŘeõřąıT

SVC_NAME.NS_NAME.DOMAIN.LTD

äĳŇăĉĈiĳŽredis.default.svc.cluster.local.

8.4 8.4 Service æÿĚă■T

- æÿĚă■TçzDæŁŘ

```
apiVersion <string>      # api çL'ÍæIJñăŘűiĳŇv1
kind <string>             # Ľ
→ etDæžŘçszăĹŋiĳŇăăĜèõřăĹŽăžžăzĂăžĹçszăđŇçŽDèTĐæžŘ
metadata <Object>        # POD âĚĈæTřă■ó
spec <Object>            # âĚĈæTřă■ó
```

8.5 8.5 service.spec èĝĐèŇĈ

1. clusterIPiĳŽæŇĜăõŽ Service âd'ĐăžŎ service çĳŚçzIJçŽDăŞläÿł
IPiĳŇézYèõđ'äÿžăĹăĂăĹĚéĚ■
2. typeiĳŽ service çszăđŇiĳŇăŘřĈĹiĳŽExternalName, ClusterIP, NodePort, and LoadBal-
ancer

8.6 8.6 ClusterIP çszăđŇçŽĐ service

```
apiVersion: v1
kind: Service
metadata:
  name: redis
  namespace: default
spec:
  selector:
    app: redis
    role: logstor
  type: ClusterIP
  clusterIP: 10.96.0.100
  ports:
    - port: 6379          # service çŋřăŘĈ
      targetPort: 6379    # pod çŽŚăŘŋçŽĐçŋřăŘĈ
      protocol: TCP
```

8.7 8.7 NodePort ęszăđŇčŽĎ service

NodePort æŸřăĬĬ ClusterIP ęszăđŇăŸĽăċđăĽăăžĚăŸĂăŸĽăŽť éĬšăĬĬăžĚ node
čŽĎč;ŚčžĬăŚ;ăŖ■č'žéŮť äŸĽčŽĎăŸĂăŸĽ nodePortĭĭŇăĽ'ĂăžčťĬăĽăŖřăžăžŎéŽĚč;đ'ăđ' ŮéĬléŏéŮŏăĽ
-> NodeIP:NodePort -> ClusterIP:ServicePort -> PodIP:ContainerPortăĬĬ

ăŖřăžčŖĚğčăŸž NodePort áċđăĭžăžĚ ClusterIP čŽĎăĽšĚč;ĭĭŇăŮĬ'ăŏčăĽŮčŇřăŖřăžăĬĬăŖŖăŸĽéŽĚč;
nodeip äžŎĚĂŇăŮŏéŮŏăĽŖ clusterIPĭĭŇăĽĚ■čťś clusterIP ěŹĚăŇĚť šĚĭĭăĬĚăăĚĜś PODăĬĬ

- æŸĚă■ťčđ'žăĬŇ

```
apiVersion: v1
kind: Service
metadata:
  name: myapp
  namespace: default
spec:
  selector:
    app: myapp
    release: canary
  type: NodePort
  ports:
    - port: 80          # service čŇřăŖč
      targetPort: 80    # pod čŽŚăŖŇčŽĎčŇăŖč
      nodePort: 30080   # service äĭjžăĬĬăŖŖăŸĽ node äŸĽăŮžăĽă
    ↪ iptables/ipvs
    ↪ ěğĎăĽŽĚč■ăŏžăŖŚĚĚžăŸĽčŇăŖččŽĎăŮŏéŮŏĭĭŇăĽ'ĂăžăăĽĚăăžăĽĬĚŖĂăĽ'ĂăĬĬĬ
    ↪ node čŽĎăĽžăŸĽčŇăŖčăśăěčă■ăčťĬ
      protocol: TCP
```

ăĬĬăĽéŽĚč;đ'ăđ' ŮéĬĬăŖśăŖřăžăă;ĲčťĬ: http://172.16.100.102:30080
↪ æĬĚăŮŏéŮŏăĽžăŸĽ service äĬĬăĽăăžĚ

ăĬĬăĽéŽĚč;đ'ăĚĚăŖřăžăă;ĲčťĬ service čŽĎăŚŚăŖ■ăĬĬ coredns
↪ äŸĽĚğčăċăđŖă;ŮăĽŖ service äĬĬăĽăĬ: dig -t A myapp.default.svc.
↪ cluster.local @10.96.0.10

8.8 8.8 loadBalancerIP ęszăđŇ

service äĬĬăŖŖăŖřăžăĬĬčŽĎ iptables/ipvs ěğĎăĽŽăĚĬĭĭŇăŮŏéŮŏăžăăĎŖăŸĂăŖř
node ěč;ăŖřăžăăĽŖĚč; podĭĭŇăĽ'ĂăžăăžťĚŖăĬĬăĽžăžŽ nodeip
ăĽ■ăĽăĚť šĚĭĭăĬĚăăŽĬĭĭŇăěčăċĬăŮĚăĬăĬăĚŇăĬĬăĽăžŚĭĭŇăŖřăžăă;ĲčťĬ k8s
ăĚĚč;ŏčŽĎ loadBalancerIPĭĭŇăŚ■ăĬăĚŇăĬĬăžŚčŽĎĚť šĚĭĭăĬĚăăăŽĬă■ăĬăĽăĭĭŇăŮŏđčŎŖăĬăĂăčŽĎ

ăŖřăžčŖĚğčăŸž loadBalancerIP áċđăĭžăžĚ NodePort ęszăđŇčŽĎ service
ĭĭŇăĬĬăĽéŽĚč;đ'ăđ' ŮéĬĬăŖřăŖŖăŖř nodeip ěŹĚăŇĚť šĚĭĭăĬĚăăăĬĬ

8.9 8.9 æUäéŽẸçŁd'äIjraiÄçŽĐ Service

æUääd't' service èäłçd'ž service æšqæIJL' ClusterIP äžšäy■æYäârĐ Node-PortüijÑèÄÑæYrârE service çŽĐäššâR■çŽt' æÖèëğçæđŘäyž nodeIP äžÖèÄÑçŽt' æÖèëóŁéUő nodeIP äyŁçŽĐ PODäÄĆ

- æyĚä■Tçd'žäŁÑ

```
apiVersion: v1
kind: Service
metadata:
  name: myapp-nohead
  namespace: default
spec:
  selector:
    app: myapp-nohead
    release: canary
  type: ClusterIP
  clusterIP: None
  ports:
    - port: 80          # service çnrârČ
      targetPort: 80    # pod çŽŠâRñçŽĐçnrârČ
```

- æšēçIJN CoreDNS æIJ■äŁäqäŽÍçŽĐäIjraiÄ

```
kubectl get svc -n kube-system
```

- äIjIéŽẸçŁd'äEĚä;ŁçTÍ CoreDNS çŽĐäIjraiÄèğçæđŘæUääd't'çŽĐ serive äššâR■üijÑä;UäŁrçŽĐçŽt' æÖèäyž nodeip äy■çŽĐ pod äIjraiÄüijÑäŁfçTÍ dns çŽĐäd'Žæİq A èörä;Tæİèet'šè;äIĞèaq

```
dig -t A myapp-nohead.default.svc.cluster.local. @10.96.0.10
```

```
;; ANSWER SECTION:
myapp-nohead.default.svc.cluster.local. 5 IN A 10.244.1.75
myapp-nohead.default.svc.cluster.local. 5 IN A 10.244.2.74
```

8.10 8.10 externalName çšžädÑ

ä;Š POD éIJÄèçÄèóŁéUőäyÄäyİéŽẸçŁd'äd'ÚéČÍçŽĐæIJ■äŁäqæUüäÄŽüijNexternalName äRfrazèæYäârĐäyÄäyİéŽẸçŁd'äd'ÚéČÍçŽĐæIJ■äŁäqäŁréŽẸçŁd'äEĚéČIüijÑä;ŽéŽẸçŁd'äEĚ POD èóŁéUőäÄĆ

ärsæYræŁŁäd'ÚéČÍçŽĐäyÄäyİäššâR■äIjraiÄüijÑæYäârĐäyžéŽẸçŁd'äEĚéČI coredns èğçæđŘçŽĐäyÄäyİäEĚéČIäIjraiÄüijÑæRŘä;ŽéŽẸçŁd'äEĚéČIèóŁéUőäÄĆ

CHAPTER 9

āī ingress æŌğāLúāī

æĈædIJ k8s éIJĀæĀæRŘāZāyĀāyĭç;ŚçñZiijNāzūāyTēfZāyĭçñZçCzéIJĀæĀæZē
https èŌfēŪŏiijNèĀN iptables/ipvs āūčā;IJāIJĪ 4 āŚÇiijNāŏcæLūāRŚāGžçŽD ssl
èrūæŚCæāzæIJñāy■ècñèğçædRārsècñèrČāžēāLrāRŌçnr PODāžEāĀĆèğçāEşæŪzæşTæIJL'āyd'āyĭiijŽ

1. āRřāzēāIJĪāĒnæIJL'āžŚçŽDèt'şè;āiĠèāāāZīāyLéĒ■ç;ŏāyL ssl èrAāžēāĀĆ
 2. æŪrāzžāyĀāyĭèt'şè;āiĠèāāāZīçŽD POD iijNāZ.NāçC nignx iijNēfZāyĭ POD
āĒsāžñāyæIJžçŽDç;ŚçzIJāŚ;āR■çl'zéŪt'iijNāzşārşæYřèrt'āRřāzēçŽt'æŌčēĀŽēfĠ
nodeip èŌfēŪŏāLrèt'şè;āiĠèāāāZīiijNssl èrAāžēēĒ■ç;ŏāIJĪēfZāyĭèt'şè;āiĠèāāāZīāyLēiijNārřād'ŪēfĠ
https èĀNārřāEĒçŽDāžççRĒāyž http ā■RèŏŏāLr POD ç;ŚçzIJçŽD POD āyLāĀĆ
- ā■YāIJĪçŽDēŪŏéçY

```
- èt'Sè;āiĠèāāāZīlā;ĲçTĪèLÇçCzçŽDç;ŚçzIJāR■çğřçl'zéŪt',  
→éČčāzĪāŏČāRĪ èČ;āIJĪèfZāyĭ node èLÇçCzāyLèfRēāNāyĀāyĭāžE,  
→āRēāLZārsāGžçŌřçnrāRČāEşçĪā  
- èt'Sè;āiĠèāāāZīāyRāzççRĒ POD ā■yè; ssl èrAāžēççŽDāĒşéTŏèLÇçCz,  
→āŏČāy■èČ;āRĪèfRēāNāyĀāyĭ, āŏČéIJĀæĀāIJĪāL'ĀæIJL'èLÇçCzèfRēāNāyĀāyĭ
```

- èğçāEşæŪzæşT

```
- èt'Sè;āiĠèāāāZīlā;ĲçTĪ DaemonSet āIJĪāRřāyĭ node  
→èLÇçCzèfRēāNāyĀāyĭ,āžççRĒèrūæŚCèğş POD ç;ŚçzIJççŽDāy■çŽD POD āyL  
- æĈædIJéZĒç;èLÇçCzéīdāyŷçŽDād'Ž,āĒūāŏdāy■āfĒāIJĪāRřāyĭ node  
→èLÇçCzéČ;āfĒēāzèfRēāNāyĀāyĭèt'Sè;āiĠèāāāZīlā POD  
- æŌğāLūèt'Sè;āiĠèāāāZīlā POD èfRēāNççŽDæTřéĠRāRřāzēēĀŽēfĠ lables  
→æNĠāŏžèfRēāNéČčāGāāyĭ node èLÇçCzāyL  
- çDūāRŌāRřāzēāIJĪèt'Sè;āiĠèāāāZīlā POD æL'ĀāIJĪçŽD node  
→èLÇçCzāyLāL'SāyL "æśāçCz" ā;āĒūāzŪçŽD POD  
→āy■āiijžāfē■ècñèrČāžēāyLāīē, èĀNārĪāæIJL'èt'Sè;āiĠèāāāZīlā POD  
→āRřāzēāŏžāf■èfZāžŽ "æśāçCz"
```

- et' šè; ; ĩĜèaaāZlāRréĀL'ijNæNL'çĖġaijYāĖLçžġāĖLāRŌæŌšāžR

```
Envoy          # äžšāŌSçTŠénYæĀġèĈ;æIJ■āŁažžččŘĚ, āũšāžŌcncfærTäyž
Traefik        # äÿžāçŏæIJ■āŁaèĀŇçTŠçŽDāŘ■āŘSāžččŘĚ
Nginx          # æTzéĀāāRŌāRřäzëĖĀĆçTlāžŌāçŏæIJ■āŁaçŌrāćĈ
HAproxy       # äÿ■æŌlè■Řä;ĚçTl
```

æŮřāžžäÿĀäÿĽ service āřĖĖIJĀēĖAžžččŘĚçŽDäÿ■āŘNæIJ■āŁaçŽD pod āĽĖçšž

æŮřāžžäÿĀäÿĽ ingress ètDæžRijNāžŌ service äÿ■āRŮā; ŮāĽĖçšžçžSæđIJijNæYāārDèĚŽ
Envoy äÿ■ijNéĜ■è; ; Envoy è;řāžžāĀĆ

9.1 9.1 ingress.spec èġDèNĈ

- API āŠN kind

```
apiVersion: extensions
kind: ingress
```

- ingress.spec

```
backend      # āŘŌçnræIJL'āšlāžž POD
rules        # èřĈāžžèġDāĽž
  host       # èžžæNšäÿžæIJž
  http       # http èŮrāçD
```

9.2 9.2 ingress-nginx äžččŘĚ

- āŘŌçnr service āŠN pods

```
apiVersion: v1
kind: Service
metadata:
  name: service-ingress-myapp
  namespace: default
spec:
  selector:
    app: myapp
    release: canary
  ports:
    - name: http
      port: 80
      targetPort: 80
```

(continues on next page)

(continued from previous page)

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp
  namespace: default
spec:
  replicas: 4
  selector:
    matchLabels:
      app: myapp
      release: canary
  template:
    metadata:
      labels:
        app: myapp
        release: canary
    spec:
      containers:
        - name: myapp
          image: ikubernetes/myapp:v2
          ports:
            - name: http
              containerPort: 80

```

- `az ingress-nginx`

```

kubectl apply -f https://raw.githubusercontent.com/kubernetes/
ingress-nginx/nginx-0.21.0/deploy/mandatory.yaml

```

- `ingress-nginx` `az ingress-nginx`

```

kubectl apply -f https://raw.githubusercontent.com/kubernetes/
ingress-nginx/nginx-0.21.0/deploy/provider/baremetal/service-
nodeport.yaml

```

- `az ingress` `az ingress-nginx` `service` `az ingress-nginx` `service` `az ingress-nginx` `ingress-nginx`

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: ingress-deploy-myapp
  namespace: default
  annotations:
    kubernetes.io/ingress.class: "nginx"
spec:
  rules:
    - host: myapp.kaliarch.com
      http:
        paths:
          - path: /
            pathType: Prefix
            backend:
              serviceName: myapp
              servicePort: 80

```

(continues on next page)

(continued from previous page)

```

http:
  paths:
    - path:
      ↪ çl'žçŽĐæŮŮāĀŽāžčèāíæăžīijNèőéŮőæăžçŽĐæŮŮāĀŽæŸăăřĐăĹř backend
      backend:
      ↪ service çŽĐéĚ■ç;ő
        serviceName: service-ingress-myapp
        ↪ service äžŌèĀNèŌŮāŔŮāĹŕăŔŌçńŕăÿzæIJžçŽĐăŔŸăĹĹ
        servicePort: 80
        ↪ service çŽĐăIJŕăĪĀ

```

- æšççIJN ingress-nginx áŕžăđ'ŮæŽt' éIJšçŽĐçńŕăŔçīijNèĤŽéĜNăÿž30080īijNăŠN 30443 äÿd'äÿĹ

```
kubectl get service -n ingress-nginx
```

- äĵççĹ nodeip + ingress-nginx æŽt' éIJšçńŕăŔçèőéŮőīijNçŤsäžŌăÿĹéĪăĹŽăžççŽĐ ingress äÿžăšžăžŌăÿzæIJžăŔ■çğŕçŽĐīijNæL'ĀăžééIJĀèçAăIJĹèőéŮőæŮŮăIJĹ /etc/hosts āAžăç;æŸăăřĐăĹĹ nodeăĀĆ

```
http://myapp.kaliarch.com:30080/index.html
```

9.3 9.3 ingress-tomcat äžççŔĚ

- āŔŌçńŕ service āŠN pods

```

apiVersion: v1
kind: Service
metadata:
  name: service-ingress-tomcat
  namespace: default
spec:
  selector:
    app: tomcat
    release: canary
  ports:
    - name: http
      port: 8080
      targetPort: http
    - name: ajp
      port: 8009
      targetPort: ajp
---
apiVersion: apps/v1
kind: Deployment

```

(continues on next page)

(continued from previous page)

```

metadata:
  name: deploy-tomcat
  namespace: default
spec:
  replicas: 4
  selector:
    matchLabels:
      app: tomcat
      release: canary
  template:
    metadata:
      labels:
        app: tomcat
        release: canary
    spec:
      containers:
        - name: tomcat
          image: tomcat:8.5.32-jre8-alpine
          ports:
            - name: http
              containerPort: 8080
            - name: ajp
              containerPort: 8009

```

- `cat /dev/urandom | tr -dc 'a-z0-9' | fold -n 64 | xargs printf '%s\n'` ingress-nginx `cat /dev/urandom | tr -dc 'a-z0-9' | fold -n 64 | xargs printf '%s\n'`

```

# openssl genrsa -out tls.key 2048

#
openssl req -new -x509 -key tls.key -out tls.crt -subj /C=CN/ST=Beijing/L=Beijing/O=DevOps/CN=tomcat.kaliarch.com

```

- `kubectl create secret tls tomcat-ingress-secret --cert=tls.crt --key=tls.key`

```

kubectl create secret tls tomcat-ingress-secret --cert=tls.crt --key=tls.key

```

- `kubectl apply -f ingress-tomcat.yaml` ingress service ingress-tomcat

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: ingress-deploy-tomcat-tls
  namespace: default

```

(continues on next page)

```

  annotations:
    kubernetes.io/ingress.class: "nginx"
spec:
  tls:
    - hosts:
        - tomcat.kaliarch.com
      secretName: tomcat-ingress-secret
  rules:
    - host: tomcat.kaliarch.com
      http:
        paths:
          - path:
              backend:
                serviceName: service-ingress-tomcat
                servicePort: 8080

```

- ```
kubectl get service -n ingress-nginx
```

- https://tomcat.kaliarch.com:30443

## CHAPTER 10

### Podman POD and Yacht

ad' geCláLEæIJL' çLúæÄAçŽDāžTçTléČ;æIJL' æNĀāžĒāYāClíijNāIJl Docker  
äyLæLŠāžnārĒāōzāZlæL' ĀéIJĀēæAçŽDāYāClāūæT;āIJlāōfāyžæIJžāyLijNā;ĒæYr  
k8s äyLäyēāNijNāZāyž POD äijŽēcnāIJlāyāRŇçŽD node  
èLČçCžāyLāLZāžzāLāēZd' iijNæL' Āāž k8s éIJĀēæAçYāāēUāRēad' ŪçŽDāYāClāūæIJžāLūijNāōČēČ;èD  
k8s æRŘä; ŽāžĒad' ŽçgāyāRŇçŽDāYāClāūijNk8s  
äyāYāClāūāsdāžŌ POD èĀNāyæYrāōzāZlīijNPOD āRřāžēæNČē; iijNPOD  
äyžāzĀāžLēČ;æIJL' āYāClāūāŚcīijšēfZæYrāZāyžāIJlæL' ĀæIJL' èLČçCžāyLēfRēāNāžĒāyĀy  
Pause çŽDēTijāČRīijNāōČæYr POD çŽDāšžçāĀēđūēđDāōzāZlīijNāōČæNēæIJL' āYāClāūijNāRŇāyĀy  
POD āĒēçŽDæL' ĀæIJL' āōzāZlæYrāyĀyīç;ŠçzIJāRŇçgřçl' zēŪr' çŽDāČ

### 10.1 10.1 Podman and Yacht

æšēçIJN POD æTræNĀçŽDāYāClíçsāđNīijŽkubectl explain pods.spec.volumes

1. HostPathijŽāIJlèLČçCzæIJnāIJræŪrāžžāyĀāyēūfā;DīijNāyŌāōzāZlāžžçnNāĒšēĀTāĒšççzīijNā;ĒēLČ  
node æŪūāĀZāyēČ;ēūlèLČçCzā;ççTÍHostPathāČ
2. LocalijŽçŽt' æŌēā;ççTléLČçCzçŽDēō;ad' GāĀāžžæTræNĀāyĀāyīçZōā;TçsāijijāžŌ  
HostPathāČ
3. EmptyDirijŽāRlāIJlèLČçCzæIJnāIJrā;ççTīijNāyĀæŪç POD  
āLāēZd' iijNāYāClāūāžšāijŽāLāēZd' iijNāōČāyāĒūæIJL' æNĀāžĒæĀgīijNā;Šāyt' æŪūçZōā;TæLŪē  
4. ç;ŠçzIJāYāClíijŽiSCSIāĀNFSāĀCifsāĀAglusterfsāĀAcephfsāĀĒBSīijLAWs)āĀDiskīijLazo

## 10.2 10.2 áóžáZíæŇCè;éĀL'éąž

āIJÍ K8S äy■ā■ūæYřásdāžŮ POD çŽDiiŇNèĀŇäy■æYřáóžáZíiŇNæL'Āäžēā■ūçŽDāóŽázL'āIJÍ  
POD äy■iŇNäyĀäyĪ POD äy■āRřázēāóŽázL'ād'ŽäyĪā■ūāĀĆ

- āIJÍ POD äy■æŇCè;ä;ŁçTīiŇNkubectl explain pods.spec.containers.volumeMounts

```
apiVersion: v1
kind: Pod
metadata:
 name: myapp
 namespace: default
 labels:
 app: myapp
spec:
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 volumeMounts
 <[]Object> # ā■ūæŇCè;;āřžèśā
 mountPath
 <string> # æŇCè;;èûřā;Ď
 mountPropagation <string> # _
 ↪çāóāóŽæŇCè;;āēĆä;TāžŮäyžæIJžāijāāš■āLřāóžáZí
 name
 <string> # æŇCè;;āšlāyĪā■ū
 readOnly
 <boolean> # æYřāRřæāRlèérzæŇCè;;
 subPath
 <string> # æŇCè;;āIJlā■Rèûřā;ĎäyŇ
 subPathExpr
 <string> # äyŮ subPath_
 ↪çšžāijijijijNæŇCè;;āIJlā■Rèûřā;ĎäyŇiŇNäy■āRŇçŽDæYřāRřázēā;ŁçTī
 ↪$(VAR_NAME) ēāĪçd'žāóžáZíæL'1'āsTēŁžäyĪāRřéĜR
```

## 10.3 10.3 èŁĆçCzā■YāĆÍ

### 10.3.1 10.3.1 hostpathā■YāĆíā■ū

āIJlāóŁäyžæIJçŽDèûřā;ĎæŇCè;āLř POD äyLiiŇNPOD  
āLāéZd'āRŮiŇNā■ūæTřæ■óæYřäy■āijŽéŽRázNāLāéZd'çŽDiiŇNä;EāēĆæđIJ  
node èŁĆçCzāæŇCæŮL'iiŇNéCčázLæTřæ■óæIJL'āRřèČ;äyčād'siŇNāēĆæđIJ POD  
ècnèrČāžēāLřāĒūāzŮçŽDèŁĆçCzāiŇNéCčázLāŮšæĪēā■ūçŽDæTřæ■ōāřsèŮŁēŮöäy■āLřāzEāĀĆ

<https://kubernetes.io/docs/concepts/storage/volumes/#hostpath>

- áóŽázL'āRĆæTřiiŇNkubectl explain pods.spec.volumes.hostPath

```
path <string> # äyžæIJžäyŁçŽóā;TçŽDèûřā;ĎāĀĆ_
 ↪āēĆæđIJèûřā;ĎæYřçŇæāRūéš;æŮēiŇNāĪžāijžèûséŽRçIJšāóđèûřā;ĎçŽDés;æŮēāĀĆ
type <string> # ēĝAäyŇēāĪ
```

- çd'žā;Ň



```

apiVersion: v1
kind: Pod
metadata:
 name: myapp
 namespace: default
 labels:
 app: myapp
spec:
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 volumeMounts:
 - name: webstore
 mountPath: /usr/share/nginx/html
 readOnly: false
 volumes:
 - name: webstore
 hostPath: /data/myapp
 type: DirectoryOrCreate

```

### 10.3.2 10.3.2 gitRepo

gitRepo is a Pod that runs a container named myapp. The container has a volume named webstore that is a hostPath volume. The volume is mounted to /data/myapp. The volume type is DirectoryOrCreate.

### 10.3.3 10.3.3 emptyDir

emptyDir is a Pod that runs a container named myapp. The container has a volume named webstore that is an emptyDir volume. The volume is mounted to /data/myapp. The volume type is emptyDir.

- `emptyDir` is a Pod that runs a container named myapp. The container has a volume named webstore that is an emptyDir volume. The volume is mounted to /data/myapp. The volume type is emptyDir.

```

medium <string> # "Memory" Disk
sizeLimit <string> # "1Gi" 1Gi

```

- `memory` is a Pod that runs a container named myapp. The container has a volume named webstore that is a memory volume. The volume is mounted to /data/myapp. The volume type is memory.

```

apiVersion: v1
kind: Pod
metadata:
 name: pod-volume-demo
 namespace: default
 labels:
 app: myapp
 tier: frontend
spec:
 volumes:
 - name: html
 emptyDir: {} # ä;ŁçŤĺčĈAçŽŸiijŇăŸŦæšæIJL'ăőžéĜŔéŽŘăĹŮ
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - name: html
 mountPath: /usr/share/nginx/html/
 ports:
 - name: http
 containerPort: 80
 - name: https
 containerPort: 443
 - name: busybox
 image: busybox:latest
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - name: html
 mountPath: /data/
 command:
 - "/bin/sh"
 - "-c"
 - "while true; do date >> /data/index.html; sleep 10; done"

```

- ä;ŁçŤĺčĈäŮŇ

```

apiVersion: v1
kind: Pod
metadata:
 name: pod-volume-demo
 namespace: default
 labels:
 app: myapp
 tier: frontend
spec:
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 imagePullPolicy: IfNotPresent

```

(continues on next page)

(continued from previous page)

```

volumeMounts:
 - name: html
 mountPath: /usr/share/nginx/html/
ports:
 - name: http
 containerPort: 80
 - name: https
 containerPort: 443
- name: busybox
 image: busybox:latest
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - name: html
 mountPath: /data/
 command:
 - "/bin/sh"
 - "-c"
 - "while true; do date >> /data/index.html; sleep 10; done"
volumes:
 - name: html
 emptyDir:
 medium: ""
 sizeLimit: 1536Mi

```

## 10.4 10.4 çİŞçzIJâ■YâĆÍ

çİŞçzIJâ■YâĆÍijNârsæYrèDşççzâžEèLCçCzçTşâS;âSİlæIJşçZDâ■YâĆİèöçad'GrijNâ■şaj£  
pod èçñèrÇâžçÅLřåLńçZD node èLCçCzçäyLrijNâž■çDüâRřäžæNÇè;ä;£çTİåEüäy■çZDæTřæ■õãĂĆ

### 10.4.1 10.4.1 nfs

nfs æIJ■åLqâZİæYřâ■YâIJläžÖéZEççd'ázNâd'ŮçZDæIJ■åLqâZİrijNâóCäy■åRŮ node  
èLCçCzçZDâ;şâ\$■rijNâZæĀNâIJ node èLCçCzçäoTæIJzâRÖäz■çDüèÇç;äd'şæRŘäçZæNqâžEâ■YâĆİçzZâE  
PODãĂĆ

- âIJ k8s çZD node æLçäyÄäyläyæIJžrijNâóL'èçEéE■ç;õ nfs æIJ■åLqâZİläžüâRřåLÍ

```

$ yum install nfs-utils
→ # åóL'èçE nfs æIJ■åLq
$ mkdir -p /data/volumes
→ # åÍžâžž volume å■ũçZõâ;T
echo '/data/volumes 172.16.100.0/16(rw,no_root_squash)' >> /etc/
→ exports # éE■ç;õ nfs æIJ■åLqâZİ
$ systemctl start nfs
→ # åŘřåLÍ nfs æIJ■åLqâZİ

```

(continues on next page)

(continued from previous page)

```
$ ss -tnl
çàøèød' çŽŠăŘňçnrăRČiijŇnfs çŽŠăŘň TCP 2049 çnrăRČ
```

- âĬĬ k8s éŽEçĭd' çŽĐ node èĽČćČžăŎĽ' èčĚ nfs éĽ sâĽĬiijŇætĬNërŢæŇCèĭĭæŸrăŘæ■čăÿÿ

```
$ yum install nfs-utils
$ mount -t nfs 172.16.100.104:/data/volumes /mnt
```

- âŎŽăžĽ' nfs âRČæŢriijŇkubectĭ explain pods.spec.volumes.nfs

```
path <string> # nfs æĬJ■âĽăăŽĬçŽĐèŭrăĭĐ
readOnly <boolean> # æŸrăŘæâĽèérž
server <string> # nfs æĬJ■âĽăăŽĬâĬJřăĭĂ
```

- äĭ£çŢĭçd' žăĭŇ

```
apiVersion: v1
kind: Pod
metadata:
 name: pod-vol-nfs-demo
 namespace: default
spec:
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 volumeMounts:
 - name: html
 mountPath: /usr/share/nginx/html/
 volumes:
 - name: html
 nfs:
 path: /data/volumes
 server: 172.16.100.104
```

## 10.5 10.5 âĽEăÿČăijRă■ŸăĆĬ

âĽEăÿČăijRă■ŸăĆĭčĭæRŘăĭŽèĐšçèzèĽČćČžćŢšăŢ;ăŢĭæĬJšçŽĐă■ŸăĆĭriijŇăRĽæŢŢçĭŢçzĬJă■ŸăĆĭæZ  
NFS æRŘăĭŽçŽĐçĭŢçzĬJăĆĭă■Ÿăÿ■riijŇçŢĭæĽüéĬJăèçAçšééAšăĽEéĖ■çžŽ POD çŽĐ NFS  
ă■ŸăĆĭçŽĐăĬJăĭAæĽ'■èčĭ;äĭ£çŢĭriijŇèĂŇăĬJĭçŢšăĽEăÿČăijRăæRŘăĭŽçŽĐă■ŸăĆĭčĭ;ăĽZçŽĐă■ŸăĆĭăÿĽriij

çŢšæ■d' K8S æRŘăĭŽăžĖ PVăĀAPVC äÿd' çğ■æĬJžăĽüriijŇèŏĭ' æŽŏéĂŽçŢĭæĽüæŰăéĬJăăĖšăĽČăžŢšĆ

äÿĂèĽŇ PV äÿŎ PVC æŸrăÿĂăržçzŢšăŎŽçŽĐriijŇNPVăšđăžŎăĖĭăšĂriijŇNPVC  
ăšđăžŎăšŢrăÿĭăR■çğŕçĭ'žéŰŢ'riijŇăĭŢăÿĂăÿĭ PV èčŇăÿĂăÿĭ PVC  
çžŢšăŎŽriijŇăĽŇçŽĐăR■çğŕçĭ'žéŰŢ' PVC âŢsăÿ■âŢŢrăžèăĖ■çžŢšăŎŽăžĖăĂČèŕüăšĆçžŢšăŎŽăšŢrăÿĭ  
PV âŢsăŸŕçŢŢ PVC æĭèăŏŇăĽŢŢçŽĐriijŇèčŇ PVC çžŢšăŎŽçŽĐ PV çğŕăĭĬJ PV  
çŽĐçžŢšăŎŽçĽüăĂăăĂČ

PVC çZŠăŮŽăEăYĂăYł PVİİjŇéČčăZŁ PVC æL'Ăăd'ĐăŘ■çğřçł'žéŮt'ăŮŽăZŁçŽĐ  
 POD ārŝăRřăžăē;łçł persistentVolumeClaim çŝžăđŇçŽĐ volumes  
 äžEİİjŇçĐŮăRŎăŮăZăłărŝăRřăžăēĂžēłĞ volumeMounts æŇČè; PVC  
 çŝžăđŇçŽĐă■ăžEăĂČ

persistentVolumeClaim â■ăYřăŘăĂăŮăđ'ŽëŮrëržăEžİİjŇëłZăRŮăEŝăžŮ PV  
 âŮŽăZŁ'æŮŮăĂžçŽĐëržăEžçŁ'zæĀgİİjŽă■TëŮrëržăEžăĂĂăđ'ŽëŮrëržăEžăĂĂăđ'ŽëŮrăRłëržăĂČ

ăēČăđİJăŝŘăYł POD äy■ăİJłİJăĕĂăžEİİjŇăŁŝăžŇăŁăŮČăŁăéŽđ'ăžEăĂĂăRŇăŮŮăžŝăŁăéŽđ'ăžE  
 PVCăĂĂăČčăZŁă■đ'æŮŮ PV èłYăRřăžăēİJŁ'èĞłăŮŝçŽĐăŽđăTŭç■ŮçTëİİjŽ  
 deleteăŁăéŽđ'PVăĂĂăRetainăžĂăžŁéČ;äy■ăĂžăĂČ

## 10.5.1 10.5.1 PersistentVolume

çłŤŝŮăçŘEăŝYăŮăZăŁăçŽĐçŽĐăYăĂăYł■YăČłçŽĐăRŘëçİİjŇăYřăYĂăYłéZEçŁđ'çžğăŁŇçŽĐăĬăŝĂĕł  
 Pod éłĂăřĂăŮŮăřZ PV æŝăæİJŁ'ă;ŝăŝ■ăĂČ

èğĂİİjŽkubectł explain PersistentVolume.spec

- âİJł nfs äYŁăŮŽăZŁ'ă■YăČłİİjŇ/etc/exportsİİjŇăžŮăYłŤăřİjăĞž nfs âŮŽăZŁ'

```
/data/volumes/v1 172.16.100.0/16(rw,no_root_squash)
/data/volumes/v2 172.16.100.0/16(rw,no_root_squash)
/data/volumes/v3 172.16.100.0/16(rw,no_root_squash)
/data/volumes/v4 172.16.100.0/16(rw,no_root_squash)
/data/volumes/v5 172.16.100.0/16(rw,no_root_squash)
```

```
exportfs -arv
```

- ārE nfs âİJł k8s äy■ăŮŽăZŁ'ăYž PersistentVolumeİİjŇëřëğĂİİjŽkubectł explain PersistentVolume.spec.nfs

```
apiVersion: v1
kind: PersistentVolume
metadata:
 name: pv-001
 labels:
 name: pv001
spec:
 accessModes:
 - ReadWriteMany
 - ReadWriteOnce
 capacity:
 storage: 1Gi
 nfs:
 path: /data/volumes/v1
 server: 172.16.100.104

```

```
apiVersion: v1
```

(continues on next page)

```
kind: PersistentVolume
metadata:
 name: pv-002
 labels:
 name: pv003
spec:
 accessModes:
 - ReadWriteMany
 - ReadWriteOnce
 capacity:
 storage: 2Gi
 nfs:
 path: /data/volumes/v2
 server: 172.16.100.104

apiVersion: v1
kind: PersistentVolume
metadata:
 name: pv-003
 labels:
 name: pv003
spec:
 accessModes:
 - ReadWriteMany
 - ReadWriteOnce
 capacity:
 storage: 3Gi
 nfs:
 path: /data/volumes/v3
 server: 172.16.100.104
```

|                              |              |              |                |           |  |   |
|------------------------------|--------------|--------------|----------------|-----------|--|---|
| kubectl get persistentvolume |              |              |                |           |  |   |
| NAME                         | CAPACITY     | ACCESS MODES | RECLAIM POLICY | STATUS    |  |   |
| ↪ CLAIM                      | STORAGECLASS | REASON       | AGE            |           |  |   |
| pv-001                       | 1Gi          | RWO,RWX      | Retain         | Available |  | ↪ |
| ↪                            |              |              | 3m38s          |           |  |   |
| pv-002                       | 2Gi          | RWO,RWX      | Retain         | Available |  | ↪ |
| ↪                            |              |              | 3m38s          |           |  |   |
| pv-003                       | 3Gi          | RWO,RWX      | Retain         | Available |  | ↪ |
| ↪                            |              |              | 3m38s          |           |  |   |

æYř Namespace çğǵÁŁńçŽĐēĭDæžŘĭijNæŘŘèřřǎř PV  
çŽĐäyÄäyĭèrûæśCăĂĈèrûæśCăřæAřăNĚăRňă■YăĆĭăđ gǎřŘĭijNěôĚUôăĭăajřRç■LăăĂĈ

```

accessModes <[]string> # ReadWriteOnce, ReadOnlyMany, ReadWriteMany
dataSource <Object> # VolumeSnapshot
resources <Object> # PersistentVolumeClaim
selector <Object> # PersistentVolumeClaim
storageClassName <string> # PersistentVolumeClaim
volumeMode <string> # PersistentVolumeClaim
volumeName <string> # PersistentVolumeClaim

```

- `kubectl explain pods.spec.volumes.persistentVolumeClaim`

```

persistentVolumeClaim
 claimName <string> # PersistentVolumeClaim
 readOnly <boolean> # PersistentVolumeClaim

```

- `kubectl explain PersistentVolumeClaim.spec`

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
 name: my-pvc
 namespace: default
spec:
 accessModes:
 - ReadWriteMany
 resources:
 requests:
 storage: 2Gi

```

- `kubectl explain pod.spec.volumes.persistentVolumeClaim`

```

apiVersion: v1
kind: Pod
metadata:
 name: pod-vol-nfs-demo
 namespace: default
spec:

```

(continues on next page)

(continued from previous page)

```

containers:
- name: myapp
 image: ikubernetes/myapp:v1
 volumeMounts:
 - name: html
 mountPath: /usr/share/nginx/html/
volumes:
- name: html
 persistentVolumeClaim:
 claimName: my-pvc # ä;£çŧŀçŽŽ PVC çŽĐăŘ■çğř

```

### 10.5.3 StorageClass

PVC çŧŧèŕŭ PV çŽĐăŮŭăŽŭijŊæIJL'çņęăŔŀŁæĹăžŭçŽĐ PVijŊk8s  
 äÿæŁŖăžŭăĖăđ'ĠăžĖ StorageClass âŔŕăžęăĹĹ PVC çŧŧèŕŭ PV çŽĐăŮŭăŽŭéĂŽèĖĠ  
 StorageClass âŁŀăĂŖçŧŧæŁŔ PVăĂĈ

StorageClass âŔŕăžęăŁŀăĂŖçŽĐăŁŕ CephFS āĂĀNFS  
 ç■Łă■ŸăĈŭijŁăŁŮëĂĖăžŖŕă■ŸăĈŭijŁăžğçŧŧŕŕăŸŕ PVijŊĖęĂæŖŖă■ŸăĈŭŕăđ'ĠăĖĖăžæŧŕăŊĀ  
 RESTfull ěĈăŕijçŽĐăŖăŔčăĂĈ

## 10.6 StorageClass Ceph RBD

### 10.6.1 éĖ■çŕŕ Ceph âĈŀă■ŸæŖă

- âŁŽăžž ceph â■ŸăĈŀăŖă

```

yum install -y ceph-common
âIJĹăŁ'ĂæIJL'èĹĈçĈăăŖŀ'ěĈĖ ceph-common

```

```

ceph osd pool create kube 4096
âŁŽăžž pool
ceph osd pool ls
æŖĖçIJŊ pool

ceph auth get-or-create client.kube mon 'allow r' osd 'allow rwx,
pool=kube' -o /etc/ceph/ceph.client.kube.keyring
ceph auth list
æŖŀăŕĈ client.kube çŧŀăŮŭèŖŕŕăŕ kube,
èĖŖăžŕŕ pool

scp /etc/ceph/ceph.client.kube.keyring node1:/etc/ceph/
âŕĖçŧŀăŮŭ keyring
æŮĠăžŭăŊŭèt'ĭăŁŕăŔĐăŸŕ ceph èĹĈçĈă
scp /etc/ceph/ceph.client.kube.keyring node1:/etc/ceph/

```



```
$ git clone https://github.com/kubernetes-incubator/external-
→ storage.git # äÿÑè;j rbd-provisioner
$ cat >>external-storage/ceph/rbd/deploy/rbac/clusterrole.yaml<<EOF
→ # åĚÅèőÿ rbd-provisioner èőféűő ceph
→ çŽĎářĚěšě
 - apiGroups: [""]
 resources: ["secrets"]
 verbs: ["create", "get", "list", "watch"]
EOF
$ kubectl apply -f external-storage/ceph/rbd/deploy/rbac/
→ # åŁ'èčĚ rbd-provisioner
```

```
https://github.com/kubernetes-incubator/external-storage/tree/
↪ master/ceph/rbd/examples # rbd-provisioner ä;ŧčŧ ceph rbd_
↪ čŧčd' žä;Ŧ
```

```

apiVersion: v1
kind: Secret
metadata:
 name: ceph-admin-secret
 namespace: kube-system
type: "kubernetes.io/rbd"
data:
 # ceph auth get-key client.admin | base64
 ↪ # äzŒèŁŽäÿłåŚ;äzd'äÿ■åŘŮå¿Ů keyring_
 ↪ èöd'èrAçŽĎ base64 årEëŠěäÿšåd'■åŁŮåŁřäÿNéİć
 key: QVFER3U5TmM1NXQ4SlJBQXhHMGltdXZlNFZkUXRvN2tTZ1BENGc9PQ==

apiVersion: v1
kind: Secret
metadata:
 name: ceph-secret
 namespace: kube-system

```

(continued from previous page)

```

type: "kubernetes.io/rbd"
data:
 # ceph auth get-key client.kube | base64
 # äžŒēŁžäyłāš;äzd'äy■āŘŮāčŮ keyring
 → èŒd'èřAçŽĎ base64 āřĚéšěäyšād'■āĹŮāĹřäyŇéĭč
 key: QVFCcUM5VmNWVDdQRlJBQWR1NUxFNzVKeThiazdUWVhOa3N2UWc9PQ==

```

- āĹŽāžž StorageClass æŇĜāŘŠ rbd-provisionerĭĭjŇ

```

kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
 name: ceph-rbd
provisioner: ceph.com/rbd
reclaimPolicy: Retain
parameters:
 monitors: 172.16.100.9:6789
 pool: kube
 adminId: admin
 adminSecretName: ceph-admin-secret
 adminSecretNamespace: kube-system
 userId: kube
 userSecretName: ceph-secret
 userSecretNamespace: kube-system
 fsType: ext4
 imageFormat: "2"
 imageFeatures: "layering"

```

- āĹŽāžž PersistentVolumeClaim

```

kind: PersistentVolumeClaim
apiVersion: v1
metadata:
 name: ceph-rbd-pvc data-kong-postgresql-0
spec:
 storageClassName: ceph-rbd
 accessModes:
 - ReadWriteOnce
 resources:
 requests:
 storage: 1Gi

```

- āĹĲ POD äy■ā;ŁçŤĲ PVCĭĭjŇæĲĲāŘŮāĲĲāŒžāŽĹäy■æŇĈèĭĭ PVCāĲĲ

```

apiVersion: v1
kind: Pod

```

(continues on next page)

(continued from previous page)

```
metadata:
 name: ceph-sc-pvc-demo
 namespace: default
spec:
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 volumeMounts:
 - name: pvc-volume
 mountPath: /usr/share/nginx/html/
 volumes:
 - name: pvc-volume
 persistentVolumeClaim:
 claimName: ceph-rbd-pvc
```

# CHAPTER 11

## ■AäyÄ éĚ■ç;őäŁæAřăőăZÍlăŇŮ

k8s æRŘä;ZăĚ configMapăĂsecret èŁŽăyď'çġ■çL'zæŁçşzăđŇçŽĐă■ŸăĆlă■üiijŇăď'ŽæŤræĈĚăĚăy  
POD æRŘä;Ză■ŸăĆlč'žéŮŤ'ijŇëĀŇæŸřăyžçŤlæŁăæRŘä;ZăĚăžŌéŽĚç;ď'ăď'ŮéĆlăĹŤ  
POD âĚĚĈlăşlăĚĚéĚ■ç;őäŁæAřçŽĐăŮăijRăĂĆ

- éĚ■ç;őäŁæAřăőăZÍlăŇŮæIJL'ăŞlăžZæŮăijR
- 1. èĠlăŌŽăZL'ăŚ;ăzd'ëăŇăŖĆæŤriijŇă;ŇăçĈiijŽcommandăĂAargsiijŇăăžă■ő args  
ăijăéĂŞăy■ăŖŇçŽĐăŖĆæŤræĹăřĚăőăZÍlăĚŤrëăŇăyžăy■ăŖŇçŽĐçL'zæĂġ
- 2. çŽŤ æŌĚăĹĹéĚ■ç;őäŁæAřăĹă;IJăyž image äy■iijŇă;ĚăŸŤēŁçġæŮăijRēĹăyăy■çAŤăŤ'zriijŇēŁŽ
- 3. çŌŖăçĈăŖŸéĠŖiijŇCloud Native æŤræŇăĚăŽēŁĠçŌŖăçĈăŖŸéĠŖăĹăĹă;ēĚ■ç;őiijŇăĹŮēĂĚă;Łç  
ENTRYPOINT èĐŽăIJŇăĹēēĈĐăď'ĐçŖĚçŌŖăçĈăŖŸéĠŖăyžéĚ■ç;őäŁæAř
- 4. â■ŸăĆlă■üiijŇăIJlăőăZÍlăŖăĹăŮăăĂZăŇĈè;ăyĂăyĹă■ŸăĆlă■üiijŇăĹŮēĂĚăyŞçŤlçŽĐéĚ■ç;őă■Ÿ
- SecretăyŌConfigMapăřăŤ

```
çŽŸăŖŇçĈZiijŽ
- key / value çŽĐă;ćăijR
- âşďăžŌăŞŖăyłçL'zăőžçŽĐ namespace
- âŖŖăžēăŖijăĠăŹăĹŖçŌŖăçĈăŖŸéĠŖ
- âŖŖăžēăĂžēŁĠçŽŌă;Ť/
→æŮĠăžăŮă;ćăijRăŇĈè; (æŤræŇăŇĈè;ăĹ'ĂăIJL'keyăŞŇĈĹăĹĚkey)

ăy■ăŖŇçĈZiijŽ
- Secret âŖŖăžēēĈŇ ServerAccount âĚşëĂŤ(ă;ŁçŤl)
- Secret âŖŖăžēă■ŸăĆl register çŽĐĚL't'ăĹĈăŁæAřiijŇçŤlăIJĹ
→ImagePullSecret âŖĆæŤŖăy■iijŇçŤlăžŌăŇL'ăŖŮçġĂæIJL'ăžşăžşçŽĐĚŤIJăĈŖ
- Secret æŤræŇă Base64 âĹăăŖĚ
- Secret âĹĚăyž kubernetes.io/Service AccountiijŇkubernetes.io/
→dockerconfigjsoniijŇOpaqueăyĹ'çġ■çşzăđŇ, Configmap äy■ăŇžăĹĚçşzăđŇ
```

(continues on next page)

(continued from previous page)

```
- Secret æŮĠăzŭă■ŸăĈlăIŮlŮmpfsæŮĠăzŭçşzçzŞăŸ■iijŃPod âĹăéŹd'ăŘŮ
→SecretæŮĠăzŭăzŞăijŹăřzăžŤčŹĎăĹăéŹd'ăĂĈ
```

## 11.1 11.1 POD èŮăRŮçŮřăĈĈăRŸéĠŖ

- enviiŹŇěřęġAiiŹkubectl explain pods.spec.containers.env

```
name <string> # âŖŸéĠŖăĤ■çġŕ
value <string> # âŖŸéĠŖçŹĎăĀij
valueFrom <Object> # âijŤčŤlăĀijiiŹŇăęĈiiŹŹconfigMap
→çŹĎăŞŖăŸlêŤŮăĀPOD âŮŹăžĹ'ăŸ■çŹĎăŮăŮŮăŖ■iijŇăęĈiiŹŹmetadata.labels
resourceFieldRef <Object> # âijŤčŤlêŤĎăžŖéŹŖăĹŮăŸ■çŹĎăĀij
secretKeyRef <Object> # âijŤčŤl secretKey
```

## 11.2 11.2 configMap

ăĀĠăęĈăĹŤăžŇçŮŮăĬlêęĀăŖŖăĹăŸăĀŸĹ POD iijŇěŤŹăŸĹ POD  
ăŖŖăĹăĹăŮăĀŹiijŇěĬĀêęĀęŖzăŖŮăŸ■ăŖŇçŹĎăĤ■ç;ŮăŤăęĀŖiijŇěĈăžĹăĹŤăžŇăĬĹăŸd'çġ■ăŮăiŹŖiijŹ

1. âŖŖăžēăŖę configMap èŤĎăžŖăĤŖăĤăĹŖă;ŞăĹ■ POD äŸĹiijŇ-  
POD äžŮ configMap ěŖzăŖŮăŸăĀŸăĤŖă■ŮiijŇăijăéĀŤçŹŹ POD  
ăĤĤéĈĹăŮăŹĹčŹĎăŸăĀŸăŖŸéĠŖiijŇăŖŸéĠŖěĈŇăęŤăĤăŖŮiijŇăŖŖăžēéĠ■ăŖŖăŮăŹăŹĹăĈ
2. âŖŖăžēăŖę configMap èŤĎăžŖăĤŖăĤăĹŖă;ŞăĹ■ POD  
äŸĹiijŇăĬĀŸăžăŸăĀŸăĤŮĠăžŭçşzçzşçŹĎăŮŖăĬŮiijŇěŤŹăŸĹčŹŮăĬŤă■Ĉăě;ăŸŖăžŤčŤlĈĹăŹŖěŖzăŖŮăĤé  
configMap âŤŮăŤăžĤiijŇěĈăžĹăŖŖăiŹéĀŹçşġ POD iijŇPOD  
ăŖŖăžēěŤŹăŇěĠ■ěĤ■ç;ŮăĈ

ăĬĹăŖŖăŸĹ configMap äŸ■ăĹăĀĬĹčŹĎăĤ■ç;ŮăŤăęĀŖěĈ;ăŤĹă■ŸăŸŖéŤŮăĀijçŹĎăĤ■ç;ŮăĬăiŹŖăĈ

- æŸĤă■ŤăăiŹăiŹŖiijŇěřęġAiiŹkubectl explain configMap

```
apiVersion <string> # çĹ'ĹăĬŮăŖăŮ
binaryData <map[string]string> # äžŇěŤŹăĹŮçŹĎăŤŖă■Ů
data <map[string]string> # éŤŮăĀijăŖŖçŹĎăŤŖă■Ů
kind <string> # âŖŖăŤăçşşăĎŇ
metadata <Object> # âŖŖăŤăçşşăĤŖă■Ů
```

- âŤăžd'ěăŇăŮăiŹŖăĹŤăžž

```
âĹăžăžăŖă■ăŸž my-config çŹĎ
→configMapiijŇăŮĈçŹĎăŤŖă■ŮăĤěěĠçŹŮăĬŤăŸ■çŹĎăŮĠăžŭiijŇěŤŮăŸŖăŮĠăžŭăŖ■iijŇăĀijăŸž
kubectl create configmap my-config --from-file=path/to/dir

âĹăžăžăŖă■ăŸž my-config çŹĎ
→configMapiijŇăŮĈçŹĎăŤŖă■ŮăĤěěĠăŮĠăžŭăŸ■çŹĎăŤŮăĀijăŖŖ
```

(continues on next page)

(continued from previous page)

```
kubectl create configmap my-config --from-file=path/to/file

ĀĹŽāžžāŘ■äÿž my-config çŽĎ
→configMapiijNăzŞăŘřăžěæL'NăĹĹăŇĜăŌŽéTŌçŽĎăŘ■çĝř
kubectl create configmap my-config --from-file=key1=/path/to/bar/
→file1.txt --from-file=key2=/path/to/bar/file2.txt

äžŌă■ŮéÍcéĜŘăÿ■āĹŽāžž
kubectl create configmap my-config --from-literal=key1=config1 --
→from-literal=key2=config2

äžŌenvăŮĜăžüäÿ■āŚ;āŘ■ my-config
kubectl create configmap my-config --from-env-file=path/to/bar.env
```

## 11.2.1 11.2.1 æşĹăĚĚ POD ENV

- āĹŽāžž ConfigMap āžŭāĪĪ POD ENV äÿ■äÿçŤĪ

```
apiVersion: v1
kind: ConfigMap
→ConfigMap āřžèşā # āĹŽāžž
metadata:
 name: nginx-config
 namespace: default
data:
 server_name: myapp.kaliarch.com #
→éŤŌăĀijăřžæŤřă■Ō #
 nginx_port: | #
→éŤŌăĀijăřžæŤřă■ŌiijNăd'ăd'Ďăÿž nginx
→éĚ■ç;ŌăŮĜăžüiijNéIJĂëçAæşĹăĎŘă■céăŇçŽĎăĚžæşŤ
 server {
 server_name myapp.kaliarch.com;
 listen 80;
 root /data/web/html;
 }

apiVersion: v1
kind: Pod
metadata:
 name: pod-configmap-demo
 namespace: default
 labels:
 app: myapp
 tier: frontend
 annotations:
 kaliarch.com/created-by: "cluster amdin"
```

(continues on next page)



(continued from previous page)

```

<h1>this is a test page</h1>
vhost: |
→ æŦræ■ō2iijŇăóĈăŦrăžěăIJÍ container äÿ■ä;ŧçŦÍ ENV
→ æşláĖĖçŎŕăĉĈăŦŦēĠŦiijŇăžŞăŦŦăžěăIJÍ container äÿ■ä;ŧçŦÍ
→ volumeMounts æŇĈè; ;æĹŦăÿžæŨĠăžű
server {
 listen 80;
 server_name localhost;

 location / {
 root /usr/share/nginx/html;
 index index.html index.htm;
 }

 error_page 500 502 503 504 /50x.html;
 location = /50x.html {
 root /usr/share/nginx/html;
 }

 location = /hostname.html {
 alias /etc/hostname;
 }
}
server {
 server_name myapp.kaliarch.com;
 listen 80;
 root /data/web/html;
}

apiVersion: v1
kind: Pod
metadata:
 name: pod-configmap-volumes-demo
 namespace: default
 labels:
 app: myapp
 tier: frontend
 annotations:
 kaliarch.com/created-by: "cluster amdin"
spec:
 containers:
 - name: myapp
 image: ikubernetes/myapp:v1
 ports:
 - name: http
 containerPort: 80
 volumeMounts:
 - name: nginx-conf

```

(continues on next page)



(continued from previous page)

```

 mountPath: /etc/nginx/conf.d
 readOnly: true
 - name: nginx-page
 mountPath: /data/web/html/
 readOnly: true
volumes:
 - name: nginx-conf
 configMap:
 name: nginx-config-volumes
 items:
 - key: vhost
 path: www.conf
 mode: 644
 - name: nginx-page
 configMap:
 name: nginx-config-volumes
 items:
 - key: index
 path: index.html
 mode: 644

```

- `kubectl exec -it pod-configmap-volumes-demo -c myapp -- /bin/sh`

```
kubectl exec -it pod-configmap-volumes-demo -c myapp -- /bin/sh
```

- `curl 10.244.2.104`

```

$ curl 10.244.2.104
Hello MyApp | Version: v1 | Pod Name

$ curl -H "Host:myapp.kaliarch.com" 10.244.2.104
<h1>this is a test page</h1>

```

## 11.3 secret

```

configMap æYræYÖæÜGãYäCíæTæöçZDijNæCædIJéIJÄæAãYäCíæTæDæTæöiijNäLZéIJÄ
secret iijNsecret äyÖ configMap çZDäIJçTlâ$æIJnäyÄèGt' iijNäyT secret
äyçZDæTæöäyæYræYÖæÜGãYäTçZDijNæNæYr base64 çijÜçäAäfläYçZDäÄC

```

- `secret çszädn`

```
docker-registry # docker registry secret
generic #
→ äzÖæIJñâIJræÜGäzÜiijNçZôâ;TæLÜâ■ÜéíçâÄijâLZâzZäyÄäyI secret
tls # docker TLS secret
```

- æyËâ■TæäijäijRiijNëfëgÄijZkubectl explain secret

```
apiVersion <string> # API çL'ÍæIJñ
data <map[string]string> #
→ äzëéTôâÄijârZâLÜâGžæTṛæ■öiijNâÄijéIJÄëeAçZRè£G base64 âLäârE
kind <string> # árZèsaçsZâdN
metadata <Object> # äËÇæTṛæ■Ö
stringData <map[string]string> # æYÖæÜGçZDæTṛæ■Ö
type <string> # æTṛæ■öçsZâdN
```

### 11.3.1 11.3.1 çgAæIJL'âzŞâzŞèód'èrA1

- éçÜâËLéÄZë£GâS;äzd'èqNâLZâzZâGžæIë secret

```
kubect create secret docker-registry regsecret --docker-
→ server=registry-vpc.cn-hangzhou.aliyuncs.com --docker-
→ username=admin --docker-password=123456 --docker-
→ email=420123641@qq.com
```

- âçCædIJæÇsä£Iâ■YäyZæÜGäzÜâRfäZë

```
kubectl get secret regsecret -o yaml
```

- POD âLZâzZæÜüâÄZiijNäZÖ docker hub æNL'ârÜéTijâÇRä;£çTíçZDçTíæLüâR■ârEçäÄijNkubectl explain pods.spec çZD imagePullSecrets â■Üæöt

```
apiVersion: v1
kind: Pod
metadata:
 name: secret-file-pod
spec:
 containers:
 - name: mypod
 image: redis
 imagePullSecrets:
→ èÖüâRÜéTijâÇRÉIJÄëeAçZDçTíæLüâR■ârEçäÄ #
 - name: regsecret # secret árZèsa
```

### 11.3.2 11.3.2 çgAæIJL'âzŞâzŞèód'èrA2

- éçÜâËLéÄZë£GâS;äzd'èqNâLZâzZâGžæIë secret

- ```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: admin
  namespace: default
imagePullSecrets:
- name: regsecret # æŃĜåŮŽ secret
```

- ```
apiVersion: v1
kind: Pod
metadata:
 name: pod-serviceaccount-demo
 namespace: default
 labels:
 app: myapp
 tier: frontend
spec:
 containers:
 - name: nginx
 image: ikubernetes/myapp:v1
 ports:
 - name: http
 containerPort: 80
 serviceAccountName: admin
```

```
kubectl create secret tls nginx-secret --cert=tls.crt --key=tls.key
```

- ```
apiVersion: v1
kind: Pod
```

88

(continued from previous page)

```

metadata:
  name: pod-configmap-volumes-demo
  namespace: default
  labels:
    app: myapp
    tier: frontend
  annotations:
    kaliarch.com/created-by: "cluster amdin"
spec:
  containers:
    - name: myapp
      image: ikubernetes/myapp:v1
      ports:
        - name: http
          containerPort: 80
      volumeMounts:
        - name: nginx-conf
          mountPath: /etc/nginx/secret
          readOnly: true
  volumes:
    - name: nginx-conf
      configMap:
        name: nginx-secret
      items:
        - key: tls.key
        - key: tls.crt
        - path: www.conf
        - mode: 644

```

CHAPTER 12

■AäžÑ StatefulSet æŌğáLúāŽí

StatefulSet éĀĆçŦíāžŌæIJL'çŁúæĀAçŽDāžŦçŦíiijNāyĀēLñāōČçōaçŘEçŽDāĒūæIJL'äyĀäyNçL'žçCžçç
POD èŦDæžŘ

1. çĺşăőŽăyŦăŦŕăyĀçŽDç;ŚçzIJæăĠerEçņę
2. çĺşăőŽăyŦæŊAăžĒçŽDă■ŸăĆí
3. æIJL'ăžŔăĀAăžşæzŚçŽDēĆíç;şăŞŊæL'ŦăŦ
4. æIJL'ăžŔăĀAăžşæzŚçŽDçzŁæ■căŞŊăŁăéŽd'
5. æIJL'ăžŔçŽDæzŽăŁăēŽt'æŨŦ

äyĀäyġăĒyădNçŽD StatefulSet äžŦçŦíāyĀēLñăŊĒăŔŋăyL'äyŦçzDăžūiijŽ

1. headless service iijŁæŨăăd't' serviceiijL'
2. StatefulSet iijŁæŌğáLúāŽíiijL'
3. volumeClaimTemplateiijŁă■ŸăĆíă■ũçŦşèrûæġæĬiijL'

ăŔDăyŦ POD çŦíăŁŦçŽDă■ŸăĆíă■ũăĤĒéąžă;ŦçŦíçŦŦ StorageClass
ăŁăæĀAăç;ŽçzŽæŁŨēĀĒçŦşçōaçŘĒăŦŸăžŊăĒŁăŁŽăžzăē;çŽD PVăĀĆ

ăŁăéŽd' StatefulSet æŁŨēĀĒçijl'ăĠŔăĒŨēğDăġăŦijèĠt' POD
èçŋăŁăéŽd'æŨŸäy■äijŽēĠăŁăŁăéŽd'ăĒŨă■ŸăĆíă■ũăžççăōăĬăŦŕæ■ōăōL'ăĒĬăĀĆ

StatefulSet æŌğáLúāŽíăç;ĲetŨăžŌăyĀäyġăžŊăĒŁă■ŸăĬĲçŽD headless Service
ăržèşăăōđçŌŦ POD äržèşăçŽDăŊAăžĒăĀAăŦŕăyĀçŽDæăĠerEçņęēĒ■ç;ŋiijŽă■d'
headless Service éIJăēçAçŦŦçŦíăŁúăL'ŊăŁĲē■ç;ŋiijŊăōČēČ;ăōđçŌŦăĬĲ POD
ăĠžçŌŦăŦĒēŽIJèçŋēĠ■ădDăŨăăŽiijŊăç;ĲçDŨēČ;ăd'şăç;ŦçŦíāzŊăL'■çŽDăyžæIJžăŔ■ăĀĆ

12.1 12.1 æÿĖā■ṪæāijāijṚ

```

podManagementPolicy    <string>          #
→æŌğāĹúæL'1'āsṪæŪūāĴçŽĐéāžāžṚç■ŪçṪē
replicas                <integer>         # æĹāæĪēēĪRēāŅçŽĐāL'ṛæIJñæṪṛ
revisionHistoryLimit   <integer>         #
→æŽt'æŪṛāŌĒāRšæIJĀāđ'ğāĪā■ŸæṪṛéĠṚ
selector               <Object>          # æāĠç■ēĀL'æŅl'āŽĪ
serviceName            <string>          # headless service
→çŽĐāR■çğṛīijŅāšžāžŌēēŽāyĪ service äyž POD āĹēēĒ■æāĠēṛEçñē
template               <Object>          # POD
→āržēsæāĹāæĪēīijŅéIJĀēēAēĒ■ç;ōæŅCē;ā■ŸāĊĪā■ūīijŅāžṪērēä;ēçṪĪ PCV
→çšžāđŅ
updateStrategy         <Object>          # StatefulSet æŽt'æŪṛç■ŪçṪē
volumeClaimTemplates   <[]Object>        # pvs çŽĐāĹŪēāĪ

```

- POD āĖšèĀṪā;ēçṪĪ PVC éĀžè;Ś

æṛṚāyĪ POD äy■āžṪērēāŌŽāZL'äyĀäyĪ PVC çšžāđŅçŽĐ volume īijŅēēŽāyĪ PVC çšžāđŅçŽĐ volume āžṪērēāĖšèĀṪāĹṛāyĀäyĪā;šāL■āṚŅāyĀäyĪāR■çğṛĪ'žēŪṛ çŽĐ PVCīijŅēēŽāyĪ PVC āžṪērēāĖšèĀṪāĹṛēŽEç;đ'çžgāĹŅçŽĐ PV äyĹāĀĆ

statefullset āijŽāyž POD ēĠāĹāĹāĹZāžž PVC çšžāđŅçŽĐ Volume īijŅāžūāyṪāĪĪ POD æĹ'ĀāĪĪçŽĐāR■çğṛĪ'žēŪṛ äy■ēĠāĹāĹāĹZāžž PVCāĀĆ

āĪĪ StatefulSet äy■īijŅæṛṚāyĀäyĪ POD çŽĐāR■ā■ŪæŸṛāŽžāŌŽāyṪāṪṛāyĀçŽĐīijŅā■şæĪĪ'āžṚçŽĐæṪ 0 æŅCāžEīijŅēĠ■āžžçŽĐ POD ēēŸāṚŅāAŽ web-0āĀĆ

ēŌēēŪŌ Service æŪūāĴçŽĐæāijāijṚīijŽ(serviceName).(namespace).svc.cluster.localīijŅēēŽāyĪæŪāā Service āṚ■ā■ŪāĪĪēğçæđṚæŪīijŅēğçæđṚāyž POD āṚ■çğṛçŽĐāĹŅāR■āĀĆ

headless ēÇ;āĪĪēṛĀīijŅārž service çŽĐēŌēēŪŌēÇ;āđ'şēğçæđṚāyž POD IPīijŅā;EæŸṛçŌṛāĪĪēĪĪēēAæāĠēṛEçŽĐæŸṛæṛāyĪ POD çŽĐāR■ā■ŪīijŅæL'ĀāžēīijŅāṚĪēĪĪēēAāĪĪ Service āL■āĹāyĹ POD çŽĐāR■çğṛā■şāṚṛāĀĆ

äĹŅāēÇīijŽpod āṚ■çğṛāyž web-0īijŅæĪ■āĹāṚ■āyžīijŽmyappīijŅēCçāžĹēŌēēŪŌēēŽāyĪ POD āṛsä;ēçṪĪ

```
web-0.myapp.default.svc.cluster.local
```

12.2 12.2 āĹZāžž NFS PV

```

apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-001
  labels:
    name: pv001
spec:

```

(continues on next page)

(continued from previous page)

```

accessModes:
  - ReadWriteMany
  - ReadWriteOnce
capacity:
  storage: 5Gi
nfs:
  path: /data/volumes/v1
  server: 172.16.100.104
---

apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-002
  labels:
    name: pv003
spec:
  accessModes:
    - ReadWriteMany
    - ReadWriteOnce
  capacity:
    storage: 5Gi
  nfs:
    path: /data/volumes/v2
    server: 172.16.100.104
---

apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-003
  labels:
    name: pv003
spec:
  accessModes:
    - ReadWriteMany
    - ReadWriteOnce
  capacity:
    storage: 5Gi
  nfs:
    path: /data/volumes/v3
    server: 172.16.100.104
---

apiVersion: v1
kind: PersistentVolume
metadata:

```

(continues on next page)

(continued from previous page)

```

name: pv-004
labels:
  name: pv004
spec:
  accessModes:
    - ReadWriteMany
    - ReadWriteOnce
  capacity:
    storage: 10Gi
  nfs:
    path: /data/volumes/v4
    server: 172.16.100.104
---
apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-005
  labels:
    name: pv005
spec:
  accessModes:
    - ReadWriteMany
    - ReadWriteOnce
  capacity:
    storage: 10Gi
  nfs:
    path: /data/volumes/v5
    server: 172.16.100.104

```

12.3 12.3 12.3 12.3 statefulSet

```

apiVersion: v1
kind: Service
metadata:
  name: myapp
  labels:
    app: myapp
spec:
  ports:
    - port: 80
      name: web
  clusterIP: None
  selector:
    app: myapp-pod

```

(continues on next page)

(continued from previous page)

```

---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: myapp
spec:
  serviceName: myapp
  replicas: 3
  selector:
    matchLabels:
      app: myapp-pod
  template:
    metadata:
      labels:
        app: myapp-pod
    spec:
      containers:
        - name: myapp
          image: ikubernetes/myapp:v1
          ports:
            - containerPort: 80
              name: web
          volumeMounts:
            - name: myappdata
              mountPath: /usr/share/nginx/html
  volumeClaimTemplates:
    - metadata:
        name: myappdata
      spec:
        accessModes:
          - ReadWriteOnce
        resources:
          requests:
            storage: 5Gi

```

- `echo pod`

```
pod_name.service_name.ns_name.svc.cluster.local
```

12.4 12.4 æL'f'ăőzăŠŇă■Ğçžğ

- æL'f'ăőzăŠŇçijl'ăőz

```
kubectl scale sts myapp --replicas=5
```

- `■Ğçžğç■ŮçŤëijŇkubectl explain sts.spec.updateStrategy.rollingUpdate.partition`

```
āŖŕäzēāōđçŎřéĜŚäÿİéŽĀāŖŚäÿČiijŇéęŪāĚŁäzĚäzĚæŽt'æŮřād'ğăžŎç■L' äžŎād'ŽāŕŚçŽĎéČlááĚ  
↪0 çŽĎiijŇāŕśāŖŕäzēāĚİéČlæŽt'æŮřäžĚ  
kubectl patch sta myapp -p '{"spec":{"updateStrategy":{"  
↪"rollingUpdate":{"partition":4}}}}'
```

```
kubectl set image statefulset/myapp myapp=ikubernetes/  
↪myapp:v2kubectl
```

```
kubectl patch sta myapp -p '{"spec":{"updateStrategy":{"  
↪"rollingUpdate":{"partition":0}}}}'
```

CHAPTER 13

API Server ʈʈlæLũèõd'érAçşzçzş

apiserver æYræL'ÄæIJL'èrũæsCèõæUõçZDç;SâEşæÕæRçijNèrũæsCèæGçlNäy■ijNèõd'érAçTlæZÕæõ
apply -f ment.yamlijNãõðéZËäyLæYrè;ñæ■cäyZ HTTP åRèõõåRŞ apis-
erver åRŞèrũæsCçZDijNèÄNèõd'érAæYrææAçTş ~/.kube/config
èZäyLæÜGäzũæRŘä;ZçZDijNèfZäyLæÜGäzũèõrâ;TazEçõaçRÊåSÿæiCèZRçZDçTlæLũææAçrãÄC

- k8s çZD API æYf RESTfull éçÕæaijçZDijNæL'ÄäzèèDæzRæYrçTşèurâ;DæãGæYÕçZDijNãIJl
k8s äy■ijNèDæzRæRtèC;åsdazÕäy'd'äylâIJræÜzijZåsdazÕéZEç;d' æLÜ
åsdazÕæR■çgrçl'zéÜt'ãÄC

```
éZEç;d' çžgålñijžnamespaceãÄApv  
åR■çgrçl'zéÜt'ijžPODãÄAdeploymentãÄAdaemonSetãÄA serviceãÄAPCV
```

ä;NæCijZèrũæsC default åR■çgrçl'zéÜt'äyNçZD myapp-deploy
æÕgålũãZlrijNãrsæYræYnéicçZDæEZæşT

```
http://172.16.100.100/apis/apps/v1/namespaces/default/deployments/  
→myapp-deploy
```

äyLéicèalçd'žijZhttp://172.16.100.100:6443 éZEç;d'çZD apis äyNçZD apps çZDçZD v1
çL'LæIJçZD namespaces äyNãrzæL; default äyNçZD myapp-deploy æÕgålũãZl

13.1 13.1 ʈʈlæLũçZDçşzãdN

æLŠäznä;ççTl kubectl èfðæÕè k8s éZEç;d'èfZæaNæÕgålũijNãõðéZËäyLæYræ;ççTlçTlæLũæõçZõã;T
.kube/config èfZäyLæÜGäzũäy■ZDçTlæLũèfðæÕæãLr apis-
erver åõðçÕrèõd'érAçZDijNèÄNæIJL'äzZ POD ijLæCijZ-
CoreDNSijL'äzşéIJÀèæAèÕåRÜéZEç;d'çZDæfæAçrijNãõCäznãzşéIJÀèæAèfðæÕæãLr
k8s éZEç;d'äy■ijNæL'Ääzè k8s äy■çTlæLũçZDçşzãdNæIJL'äy'd'çg■ijZ

1. `azhcszaj;fcTlczDcTlælLuijZuseraccountijNad'DazOcTlælLuaouczZoa;T .kube/config æUgazuay■ijNâRrä;fcTl kubectl config -help èOuaRÜayôâL'âlZâzz`
2. `POD aj;fcTlczDcTlælLuijZserviceaccountijNæYräyÄçg■ k8s ârzèsajijNâoČâRfäzëaj;fcTl kubectl create serviceaccount -help èOuaRÜayôâL'âlZâzz`

13.2 13.2 PODâeČä;Tè£dæÖëéZÈç;d'

POD éIJÄèeAaj;fcTl serviceaccount è£dæÖëázüèöd'érAâlRéZÈç;d'ijNPOD äzNæL'ÄäzëèČ;äd'šè£dæÖëâlRéZÈç;d'æYräZäyžæIJL'äyÄäyIaEËç;ocZD service ârE POD çZDèrûæšCäzççRÈèGš apiserver çZDâIJräIÄäzEäÄČ

- âR■â■Üäyž kubernetes çZD servie äyž POD è£dæÖëâlR apiserver æRRä;ZäzEéÄZäfa

```
$ kubectl describe service kubernetes #
→éZÈç;d'âEËèČlçZD POD äyO apiserver éÄZäfaaj;fcTlçZD service
→ijNä;EæYräşlæDR apiserver éIJÄèeAèöd'érAçZD

Name:                kubernetes
Namespace:           default
Labels:              component=apiserver
                    provider=kubernetes
Annotations:         <none>
Selector:            <none>
Type:                ClusterIP
IP:                  10.96.0.1 #
→éZÈç;d'âEËèČlèöfèÜO apiserver çZDç;ŠâËš
Port:                https 443/TCP
TargetPort:          6443/TCP
Endpoints:           172.16.100.101:6443 #
→apiserver âüëä;IJçZDâIJräIÄ
Session Affinity:    None
Events:              <none>
```

13.3 13.3 serviceaccount ârzèsä

k8s çZDèöd'érAæIJL'äyd'çg■äyÄçg■æYriijZhuman useräÄäyÄçg■æYr serviceaccountijNäyNéIcâršæYräLZâzz serviceaccount äoČæYr POD èöfèÜO apiserver æL'ÄçTlçZDäyÄçg■ârzèsajijNèÄN human userijNâ■şä;f kubectl âS;äzd'èaNeÄZèfGèrzâRÜ config äy■çZDcTlælLüèÄNèöd'érAâlR apiserver çZDäÄČ

- âLZâzzäyÄäyI serviceaccount ârzèsajijNâoČäijZeGlâLlâLZâzzâzüâËšèAŤäyÄäyI secretijNè£ZäyI serviceaccount âRfäzëâlR apiserver äyLè£ZëaNeöd'érAijNä;EæYrèöd'érAäy■äzçèalæIJL'æIČéZRijNæL'ÄäzëéIJÄèeAæÖLæIČ

```
$ kubectl create serviceaccount admin
$ kubectl get secret
```

- `apiVersion: v1` `kind: Pod` `metadata: name: pod-serviceaccount-demo namespace: default labels: app: myapp tier: frontend spec: containers: - name: nginx image: ikubernetes/myapp:v1 ports: - name: http containerPort: 80 serviceAccountName: admin`

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-serviceaccount-demo
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
  - name: nginx
    image: ikubernetes/myapp:v1
    ports:
    - name: http
      containerPort: 80
  serviceAccountName: admin
```

13.3.1 13.3.1 `apiVersion: v1` `kind: Pod` `metadata: name: pod-serviceaccount-demo namespace: default labels: app: myapp tier: frontend spec: containers: - name: nginx image: ikubernetes/myapp:v1 ports: - name: http containerPort: 80 serviceAccountName: admin`

- `Pod` `apiVersion: v1` `kind: Pod` `metadata: name: pod-serviceaccount-demo namespace: default labels: app: myapp tier: frontend spec: containers: - name: nginx image: ikubernetes/myapp:v1 ports: - name: http containerPort: 80 serviceAccountName: admin`

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-serviceaccount-demo
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
  - name: nginx
    image: ikubernetes/myapp:v1
    ports:
    - name: http
      containerPort: 80
  serviceAccountName: admin
```

`serviceaccount` `apiVersion: v1` `kind: ServiceAccount` `metadata: name: admin namespace: default` `secrets: - name: admin-token` `apiVersion: v1` `kind: Secret` `metadata: name: admin-token namespace: default` `data: token: eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXLTQ9In0=`

13.4 13.4 `kubectl` `explain` `Pod` `spec`

- `kubectl explain Pod spec`

```
apiVersion: v1
kind: Config
clusters:
- cluster:
  certificate-authority-data: DATA+OMITTED
  server: https://172.16.100.101:6443
  apiserver: https://172.16.100.101:6443
```

(continues on next page)

(continued from previous page)

```

name: kubernetes
users:
- name: kubernetes-admin
  → ĄŁŰĕąłäÿ■çŽĎäÿÄäÿłçŦłæŁŰåržèśą
  user:
    client-certificate-data: REDACTED
    → ąŃćæŁŰçńŕĕŕAăžę
    client-key-data: REDACTED
    → ąŃćæŁŰçńŕçġAéŠě
contexts:
  → äÿŁäÿNæŮĜăĹŰĕął
- context:
  → ąĹŰĕąłäÿ■çŽĎäÿÄäÿłäÿŁäÿNæŮĜåržèśą
    cluster: kubernetes
    user: kubernetes-admin
    name: kubernetes-admin@kubernetes
    → äÿŁäÿNæŮĜăŘ■çġř
current-context: kubernetes-admin@kubernetes
→ ą;ŚăĹ'■äÿŁäÿNæŮĜ
preferences: {}

```

- éŽĚç;d' āŘ■çġř
āŦŕāžĕă;çŦłäÿ■āŦŦçŽĎçŦłæŁŰĕŏféŮŏäÿ■āŦŦçŽĎéŽĚç;d' āĈ

```

éŽĚç;d' ąĹŰĕąłiijžéŽĚç;d' āŦžèśąąĹŰĕął
çŦłæŁŰăĹŰĕąłiijžçŦłæŁŰåržèśąąĹŰĕął
äÿŁäÿNæŮĜiijžæÿŕăŦŦŕĕŕéŕéŽĚç;d' äÿŮçŦłæŁŰçŽĎăĚşçşăĹŰĕąłăĈ
ą;ŚăĹ'■äÿŁäÿNæŮĜiijžĕąłçd' žă;ŚăĹ'■ă;çŦłăşłäÿłçŦłæŁŰĕŏféŮŏăşłäÿłéŽĚç;d'

```

```

ĕĜłăŏžăžĹ' éĚ■ç;ŏăŕăăŦŦiijžĕŕĕĕġĀiijžkubectl config --help
ca ăŦŦĕŕAăžęăŦăĹ■ÿĕŕă;Ďiijž/etc/kubernetes ăŦăĹ■ÿăžĚæĹ'ĂæIJĹ'çŽĎ ca_
→ ăŦŦç■;ăŦŦçŽĎĕŕAăžęăŦăăŦăăĈ

```

13.5 13.5 æŰžăĹăĕŕAăžęçŦłæŁŰăĹŦ config

k8s apiserver ĕŏd'ĕŕAæŮžăĹŦŦăĹĹ'äÿd'çġ■iijŽsslĕŕAăžę ăŦŦ token
ĕŏd'ĕŕAĭijŦăĹŦăĹă;çŦł ssl ĕŕAăžęăĹŽăžçŦłæŁŰ

13.5.1 13.5.1 ăĹŽăžŽSSLĕŕAăžęçŦłæŁŰ

- ăĹŽăžžĕŦăŮĕ apiserver çŽĎçŦłæŁŰĕŕAăžę

```

# ăĹŽăžžçġAéŠě
(umask 077; openssl genrsa -out kaliarch.key 2048)

```

(continues on next page)

13.6 13.6 áĹŽăžžæŮř config æŮĜăžŮ

ăĵĕĵŤĭ kubectĭ config set-cluster áĹŽăžžăŷĂăŷĭæŮřĕŽĎ config
æŮĜăžŮĭĭĭŇæĈşèèAèőĵăőŽèĕŽăŷĭæŮřăĹŽăžžĕŽĎ config æŮĜăžŮăŔăžžăĵĕĵŤĭ --kube-
config=/tmp/test.conf æŇĜæŸŎăĂĈ

- èőĵĕĵăőéŽĖĕĵĵĎĕĎĕĎăŎĕĵĎ ca æĬĴăĎĎĎĕŔĂăžĕĭĭŇ-kubeconfig âŔăžžæŇĜăőŽ
kubectĭ äĵĕĵŤĭĵĎĎĕĎĕĵĵăŮăĵĵĭĭŇĕžŸĕĎĎăŷĵĕĵŤĭăĹăőŮĵĎăĵĵ .kub
ĕĴăăĵĵăĵĎ config

```
kubectĭ config set-cluster k8s-cluster --server=https://172.16.100.  
101:6443 --certificate-authority=/etc/kubernetes/pki/ca.crt --  
embed-certs=true --kubeconfig=/tmp/test.conf
```

- âŔĖ kaliarch ĕĴĭăĹăŮăŮăĹăăĹŕ k8s ĕĴĎ config äŷĭĭĭŇĕőĵĕĵăőĕăĹŮĕŇŕĕŕĂăžĕăŷĵ
kaliarch.crtĭĭĭŇĕőĵĕĵăőĕăĹŮĕŇŕĕĝĂĕŞĕăŷĵĭĭŇĴkaliarch.keyĭĭŇăĵĕĵŤĭ --embed-
certs=true æĬĕĕŽŔĕŮŔĕĴăžŽăĬĴăŕĖăĴăăĴŕ

```
kubectĭ config set-credentials kaliarch --client-certificate=./  
kaliarch.crt --client-key=./kaliarch.key --embed-certs=true
```

- áĹŽăžžăŷĹăŷŇæŮĜăŕžžĕăĭĭŇæŎĹăĬĈ kaliarch ĕĴĭăĹăŮăőĕĕŮăăŔăĕĝŕăŷĵ kubernetes
ĕĴĎĕŽĖĕĵĵĎ

```
kubectĭ config set-context def-ns-admin@k8s-cluster --cluster=k8s-  
cluster --user=def-ns-admin --kubeconfig=/tmp/test.conf
```

- áĹĜăăĕăĵăĹăĴăĵăĵĕĵŤĭĵĎăŷĹăŷŇæŮĜĭĭŇăĹŕăŎĹăĬĈ kaliarch áĹŕ kubernetes
ĕĴĎăŷĹăŷŇæŮĜăŷĹ

```
kubectĭ config use-context def-ns-admin@k8s-cluster --kubeconfig=/  
tmp/test.con
```

13.7 13.7 áŞžăžŎ token èőĎ'ĕŕĂ

13.7.1 13.7.1 áĹŽăžž serviceaccount

- äŷĵ POD áĹŽăžžăŷĂăŷĭ serviceaccount âŕžžĕăĭĭŇăăĈăŸŕ POD èőĕĕŮă apiserver
ĕĴĎăĜăĕŕĂ

```
kubectĭ create serviceaccount dashborad-admin -n kube-system
```


13.7.2 13.7.2 çZŠăŏŽéŽEçŁd'çŏaçŘEăŠŸèğŠèL's

- áĹŽăžž clusterrolebinding âřEçŤĹăĹŭçZŠăŏŽèĞş cluster-admin
éŽEçŁd'çŏaçŘEăŠŸĭĹăĹĴĂénŸăĬĆéŽŘĭĹL'

```
kubectl create clusterrolebinding dashborad-cluster-admin --
→clusterrole=cluster-admin --serviceaccount=kube-system:dashborad-
→admin
```

13.7.3 13.7.3 éĂŽèĚĞ serviceaccount âŁŮăĹř Token

- æL'ŁăĹřăĹŽăĹ■ăĹŽăžžçŽĎ serviceaccount âřžèśă

```
kubectl get secret -n kube-system
```

- âŁŮăĹř serviceaccount âřžèśăăŸ■çŽĎ Token

```
kubectl describe secret -n kube-system dashborad-admin-token-skz95
```

CHAPTER 14

ÅAŽŽ ċŦíæLúæiČéŽŘčšžčžš

åIJÍ k8s äy■čŽĎčŦíæLúæiČéŽŘčšžčžšæŸřä;čŦí RBAC æláaijŘčŽĎiijŇRBAC æŸř
Role-Based AC ċŽĎčijl' åEŽiijŇåĒlčgřiiijŽåšžžžŌèğŠèL'ščŽĎèóéUóæŌğåLúāĀĆ

æĹSžžñåŖřžžèŏl'äyÄäylčŦíæLúæL'ŏæijŦäyÄäylèğŠèL'siijŇèĀŇèčŽäylèğŠèL'sæŇæIJL'æiČéŽŘiijŇæ
RBAC äy■iijŇčŦíæLúæŌĹæiČåřšæŸřæŌĹæiČæšŘäylèğŠèL'sāĀĆ

```
čŦíæLúiiijLuseriijL' iijžčŦíæLúåŖřžžèæŇæIJL' æšŘäylèğŠèL' šāĀĆ
èğŠèL' šiiijLroleiijL' iijžèğŠèL' šåŖřžžèæŇæIJL' æšŘäžžèŏyåŖřāĀĆ
1. æš■ä;IJ
2. åŖžèšä
èŏyåŖřiiijLpermissioniijL' iijž_
→åIJläyÄäylåŖžèšääyLèČ;æŮ;åĹăčŽĎæš■ä;IJčžĎåŖĹèŧúæiĒiijŇčğŖžžžNäyžäyÄäylèŏyåŖřæiČ
```

• ċŦíæLúčšžāđŇ

Human Useriijž	# ċŦíæLúèŧ' ēāŖŮ
Pod Service Accountiijž	# æIJ■åŁæŧ' ēāŖŮ

• èğŠèL'ščšžāđŇ

```
- ruleiijLèğŠèL' šiiijL' āĀArolebindingiijLèğŠèL' ščžšāŏžiiijL'
- _
→clasterroleiijLéŽEç;d' èğŠèL' šiiijL' āĀAclusterrolebindingiijLéŽEç;d' èğŠèL' ščžšā
```

• æŌĹæiČčšžāđŇ

```
- ċŦíæLúéĀžèĚĚ rolebinding āŏž bind ruleiijŇrolebinding_
→åŖĹèČ;æŸřå;šåL' ■åš;åŖ■čl' žéŮŧ' äy■
```

(continues on next page)

(continued from previous page)

```
- éĀŽèĚĜ clusterrolebinding áŌž bind_
→ clauſterroleiijŇclusterrolebindingäijžāIJlæL' ĀæIJL' āŘ■çğřçl' žéŮt' çŤſæŤl
- éĀŽèĚĜ rolebinding áŌž bind clauſterroleiijŇçŤſäžŌ rolebinding_
→ āŘlāIJlā; ſāL' ■āŘ■çğřçl' žéŮt' iijŇæL' Āäžē clauſterrole_
→ æiČéŽŘècñéŽŘāLüäÿžā; ſāL' ■āŘ■çğřçl' žéŮt'
```

- éĀŽèĚĜ rolebinding áŌž bind clauſterrole çŽĎäë;äd'D

```
äëČæđIJæIJL' ā;Ĺād' žāŘ■çğřçl' žéŮt' āĀāäëČæđIJçŤl rolebinding çžſāōž_
→ ruleiijŇéČčäžĹāĹžéIJĀëæAāIJlāŕŕäÿlāŘ■çğřçl' žéŮt' éČ; āōžāžL' role
äëČæđIJä; ěçŤl rolebinding çžſāōžäÿĀäÿl clauſterrole iijŇçŤſäžŌ_
→ clauſterrole æŇææIJL' æL' ĀæIJL' āŘ■çğřçl' žéŮt' çŽĎäiČéŽŘiijŇæĀŇ_
→ rolebinding _
→ āŘlāēČ; çžſāōžā; ſāL' ■āŘ■çğřçl' žéŮt' iijŇéČčäžĹāŕſçIJAāŌžäÿžæŕŔäÿlāŘ■çğřçl' žéŮt' éČ
→ role çŽĎèĚĜçlŇäžĚāĀČ
```

14.1 14.1 æiČéŽŘāLŮèāĹ

```
kubectl get clusterrole admin -o yaml
```

14.2 14.2 āĹŽāžž Role

- āſ;äzd'eāŇāŌŽāžL'

```
kubectl create role pods-reader --verb=get,list,watch --
→ resource=pods
```

- ä;ěçŤlāÿĚā■ŤæŮžāijŔāŌŽāžL'

```
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: pods-reder
  namespace: default
rules:
- apiGroups: # āŕžāſlāžž api_
→ ç;çd' çžĎäĚĚçŽĎèĚĎæžŘèĚžæŇæſ■ä; IJ
  - ""
  resources: # āŕžāſlāžžèĚĎæžŘæŌĹæiČ
  - pods
  verbs: # æŌĹæiČāĀžāſlāžžæſ■ä; IJ
  - get
  - list
  - watch
```

14.3 14.3 aŁZǎžž rolebinding

- ä;fçŦí rolebinding ářžèšǎaŁZǎžžijŇçŦíæŁüäyŎ role çŽĐçŠǎŏŽ

```
kubectl create rolebinding kaliarch-read-pods --role=pods-reader --
↪user=kaliarch
```

- ä;fçŦíæyĚǎ■ŦæŰžǎijRǎŏŽǎžL'

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: kaliarch-read-pods
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: pods-reader
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: kaliarch
```

- aŁGǎ■ççŦíæŁüǎŠŇçŎřǎčČäyŁäyŇæŰĞ

```
$ kubectl config use-context kaliarch@kubernetes
```

- æŦŇërŦçŦíæŁüæŶřǎŘæŇæIJL' get æIČéŽŘ

```
kubectl get pods
```

14.4 14.4 aŁZǎžž clusterrole

- āŠjǎzd'èǎŇǎŏŽǎžL'

```
kubectl create clusterrole cluster-reader --verb=get,list,watch --
↪resource=pods
```

- ä;fçŦíæyĚǎ■ŦæŰžǎijRǎŏŽǎžL'

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: cluster-reader
rules:
- apiGroups:
  - ""
  resources:
```

(continues on next page)

(continued from previous page)

```

- pods
verbs:
- get
- list
- watch

```

- `clusterrole get clusterrole`

```

NAME
→ AGE
admin
→ 5d16h
cluster-admin
→ 5d16h
cluster-reader
→ 4m32s
edit
→ 5d16h
flannel
→ 5d6h
system:aggregate-to-admin
→ 5d16h
system:aggregate-to-edit
→ 5d16h
system:aggregate-to-view
→ 5d16h
system:auth-delegator
→ 5d16h
system:aws-cloud-provider
→ 5d16h
system:basic-user
→ 5d16h
system:certificates.k8s.io:certificatesigningrequests:nodeclient
→ 5d16h
system:certificates.k8s.
→io:certificatesigningrequests:selfnodeclient 5d16h
system:controller:attachdetach-controller
→ 5d16h
system:controller:certificate-controller
→ 5d16h
system:controller:clusterrole-aggregation-controller
→ 5d16h
system:controller:cronjob-controller
→ 5d16h
system:controller:daemon-set-controller
→ 5d16h

```

14.5 14.5 aŁŻaż clusterrolebinding

- aŚjāzd'ēāŃāōŽāzL'

```
kubectl create clusterrolebinding kaliarch-read-all-pods --
→clusterrole=cluster-reader --user=kaliarch
```

- æyĒā■TāōŽāzL'

```
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRoleBinding
metadata:
  name: kaliarch-read-all-pods
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-reader
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: kaliarch
```

- āŁĠæ■çTīāLūāŚŃçŌrācČäyLäyŃæŮĜ

```
$ kubectl config use-context kaliarch@kubernetes
```

- ætŃērTçTīāLūāŸrāŘæŃæIJL' get ælČéŽŘ

```
$ kubectl get pods -n kube-system
$ kubectl config use-context kubernetes-admin@kubernetes
```

14.6 14.6 rolebinding äyŌ clusterrole

āēČædIJā;ŁçTī rolebinding çzŚāōŽäyÄäyŁ clauſterrole iijŃçTśāžŌ
 clauſterrole æŃææIJL' æL' ÄæIJL' āŘ■çğřçł' žéŮt' çŽDæIČéŽŘiijŃæĀŃ rolebinding
 āRlēČjçzŚāōŽā;ŚāL'■āŘ■çğřçł' žéŮt' iijŃéČčāzLāřſçIJAāŌžäyžæfRäyŁāŘ■çğřçł' žéŮt' éČjæŮrāžžäyÄäyŁ
 role çŽDēŁĠGçlŃāžEāĀČ

- aŚjāzd'āōŽāzL'

```
$ kubectl create rolebinding kaliarch-cluster-reader --
→clusterrole=cluster-reader --user=kaliarch
```

- æyĒā■TāōŽāzL'

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
```

(continues on next page)

(continued from previous page)

```

metadata:
  name: kaliarch-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: kaliarch

```

- `kubectl config use-context kaliarch@kubernetes`

```
$ kubectl config use-context kaliarch@kubernetes
```

- `kubectl get rolebinding --namespace=kube-system cluster-reader`

```

$ kubectl get pods -n kube-system
$ kubectl config use-context kubernetes-admin@kubernetes

```

14.7 RBAC

Bind `ClusterRole` to `ClusterRoleBinding` to `ServiceAccount` or `User` or `Group`

- `kubectl create rolebinding clusterrolebinding --namespace=kube-system`

```

$ kubectl create rolebinding clusterrolebinding --namespace=kube-system \
  --clusterrole=cluster-reader --user=kaliarch --serviceaccount=kaliarch
rolebinding.rbac.authorization.k8s.io/clusterrolebinding created

```

- `kubectl create rolebinding clusterrolebinding --namespace=kube-system`

```

# Generate private key
openssl genrsa -out kaliarch.key 2048

# Create Certificate Signing Request (CSR)
openssl req -new -key kaliarch.key -out kaliarch.csr -subj \
  "/O=system:masters/CN=kaliarch/"

# Sign the CSR with the CA
openssl x509 -req -in kaliarch.csr -CA /etc/kubernetes/pki/ca.crt -CAkey \
  /etc/kubernetes/pki/ca.key -CAcreateserial -out kaliarch.crt -days 1800

```

(continues on next page)

(continued from previous page)

```
# æ$ěçIJÑèíAăžę  
openssl x509 -in kaliarch.crt -text -noout
```


CHAPTER 15

■ Až dashboard

áoČä;IJäyž k8s éŽEç; d' çŽĐéŽĐäzúā■YāIJliijNæYř kubernetes
áoYæŮžçŽĐéazçŽĐäzNäyÄrijNëřçègArijŽhttps://github.com/kubernetes/dashboard

15.1 15.1 éČíç;šæťAçlín

- äyž dashboard æRRä;Ž ssl èřAžę

```
# çŤSæĹŔçġAéŠě
(umask 077; openssl genrsa -out dashboard.key 2048)

# çŤSæĹŔäyÄäyĹèĠç;ç;èřAžęiijNæšĹæĐŔ CN
→çŽĐäÄi jâĹÉéazèeAäyŎèĠâûsçŽĐäSSâŔ■áoNăĚĹäyĂèĠt'
openssl req -new -x509 -key dashboard.key -out dashboard.crt -subj
→"/O=dashboard/CN=k8s.dashboard.com"

# æSěçIJNëřAžę
openssl x509 -in dashboard.crt -text -noout
```

- äyNë; dashboard çŽDæyĚā■TæŮĠzű

```
wget https://raw.githubusercontent.com/kubernetes/dashboard/v1.10.1/
→src/deploy/recommended/kubernetes-dashboard.yaml
```

- äyž dashboard āĹŽāžž secret ářžèšą

```
kubectl -n kube-system create secret generic kubernetes-dashboard-
→certs --from-file=dashboard.crt=./dashboard.crt --from-
→file=dashboard.key=./dashboard.key
```

- edit dashboard service to use nodeport

```
sed -i '/targetPort: 8443/a \ type: NodePort' kubernetes-dashboard.yaml
```

- create kubernetes-dashboard.yaml service to use Dashboard Secret

```
# ----- Dashboard Secret ----- #

#apiVersion: v1
#kind: Secret
#metadata:
#  labels:
#    k8s-app: kubernetes-dashboard
#  name: kubernetes-dashboard-certs
#  namespace: kube-system
#type: Opaque

#---
```

- edit dashboard service

```
kubectl apply -f kubernetes-dashboard.yaml
```

- check service is running

```
kubectl get service -n kube-system
```

- access chrome to dashboard

```
https://172.16.100.102:31097/
```

15.2 15.2 Creating Service Account and ClusterRoleBinding

- create POD serviceaccount and POD apiserver

```
kubectl create serviceaccount dashborad-admin -n kube-system
```

- create clusterrolebinding and cluster-admin

```
kubectl create clusterrolebinding dashborad-cluster-admin --
clusterrole=cluster-admin --serviceaccount=kube-system:dashborad-
admin
```

- create serviceaccount and serviceaccount


```
kubectl describe secret def-ns-admin
```

- Token `eyJhbGciOiJIUzI1NiIsInR5cCI6IkpzZW50b3R5IiwiaWF0IjoiMTUxMjM0NTY3In0=`

```
kubectl config set-credentials def-ns-admin --token=<TOKEN> --
→ kubeconfig=/tmp/test.conf
```

- `alias k8s='kubectl config set-context def-ns-admin@k8s-cluster --cluster=k8s-cluster --user=def-ns-admin --kubeconfig=/tmp/test.conf'` kaliarch `alias k8s='kubectl config set-context def-ns-admin@k8s-cluster --cluster=k8s-cluster --user=def-ns-admin --kubeconfig=/tmp/test.conf'`

```
kubectl config set-context def-ns-admin@k8s-cluster --cluster=k8s-
→ cluster --user=def-ns-admin --kubeconfig=/tmp/test.conf
```

- `alias k8s='kubectl config set-context def-ns-admin@k8s-cluster --cluster=k8s-cluster --user=def-ns-admin --kubeconfig=/tmp/test.conf'` kaliarch `alias k8s='kubectl config set-context def-ns-admin@k8s-cluster --cluster=k8s-cluster --user=def-ns-admin --kubeconfig=/tmp/test.conf'`

```
kubectl config use-context def-ns-admin@k8s-cluster --kubeconfig=/
→ tmp/test.conf
```

- `alias k8s='kubectl config use-context def-ns-admin@k8s-cluster --kubeconfig=/tmp/test.conf'` `alias k8s='kubectl config use-context def-ns-admin@k8s-cluster --kubeconfig=/tmp/test.conf'`

CHAPTER 16

16.1 16.1 éÄŽä£aælaqadN

K8S çŽĎç;ŠçzIJéÄŽä£aælaqadNäÉlçTs CNI æŒœäRçäyLçŽĎæRŠäzúæléäöđçŒriijNæRŠäzúéIJÄèçAäöđçŒriij
çŽŒäL'■æfTèçCætAèçNçŽĎæRŠäzúæIJL'riijŽflanneläÄAcalicoäÄAcaneläÄAkube-
router äÄç

- äçCä;TäLäèç;æRŠäzú

k8s äIJläRfäLçŽĎæUüäÄŽäijŽäŒriijŽ/etc/cni/net.d/ çŽŒä;TäyNärfæL'ç;ŠçzIJæRŠäzúçŽĎé■ç;ŒæŒŒ
äIJläLŽäzæUüäÄŽ k8s èrCçTléçŽäy lé■ç;ŒæŒŒGäzriijNçTsæRŠäzúæäzæ■ŒœçŽäy lé■ç;ŒæŒŒGäzúèçŽäçNä

16.1 16.1 éÄŽä£aælaqadN

1. äŒäZléU' éÄŽä£aælaqadNäRŒNäyÄäyI POD äEÄäð'ŽäyIäŒäZléU' çŽĎéÄŽä£aælaqadNä;ççTí
lo ç;Šä■äéÄŽä£a
2. PODéU' éÄŽä£aælaqadNäPOD IP çŽt' æŒœäyŒŒ POD IP éÄŽä£a
3. POD äyŒŒ ServiceriijŽPOD IP çŽt' æŒœäyŒŒ Cluster IP
4. Service äyŒŒÉçç;ð' äð'ŒéČläŒœäLüçñrçŽĎéÄŽä£aælaqadNäNingressäÄANodePortäÄALoadbacer

16.2 16.2 éÄŽä£aælaqadNäžTäśĆ

æUäèŒäšläyÄçg■ç;ŠçzIJæRŠäzúriijNäŒČäzñçTíäLrçŽĎäžTäśĆæŒzæäLéČ;æŒräžäyNäGäçg■riijŽ

1. èŽŽæNšç;ŠæäriijŽbrgriijNçTíçžrè;räzúäöđçŒräyÄäyIèŽŽæNšç;Šä■riijNäyÄçñräIJIPODäyLriijNäyÄçñr
2. äð'Žèüfäð'■çTíriijŽMacVLANriijNäšžäžŒ MAC çŽĎæŒzäijRäLŽäzž
VLAN riijNäyžæfRäyIèŽŽæNšæŒœäRçé■ç;ŒäyÄäyIçNñçñNçŽĎ MAC

āIJřāĪĀĭĭjŇāĭ;Łā; ŪāyĀāyĭŁL'ĲčŘĚĸ;Śā■āēL'Łēĭ;āđ'ŽāyĭāōzāZĭā;ŁčŤĭĭjŇēŁZāūāōzāZĭŁZt' æŌēā;ŁčŤĭ
MacVLAN èŁZēāŇēūĭēŁČĆZéĀZāŁāāĀĆ

3. çañāzūāzd' æ■ċĭjŽĸ;Śā■āēŤræŇAçañāzūāzd' æ■ċĭjŇSR-IOV ĭĭjĹā■Ťāāz-
IOèŽŽæŇšāŇŪĭĭjL' æŪzāĭjRĭĭjŇēŁZĸg■;Śā■āēŤræŇAçZt' æŌēāĪĲL'ĲčŘĚĸžgāĹnèŽŽæŇšāĠzād' Žā

16.3 16.3 K8S āŘ■çğřĸl'žéŮt'

K8S āŘ■çğřĸl'žéŮt' äyŌ POD ĸ;ŚçzĪJāŘ■çğřĸl'žéŮt' äy■āĪĹāyĀāyĭŁt' āžēĭĭjŇæL'Āāzēā■sā;ŁāĪĹāy■āŘŇ
K8S éZĚĸ;đ' āŘ■çğřĸl'žéŮt' āĒĒāĹZāžzĸZĎāy■āŘŇ PODĭĭjŇāzšāRřāzēēĀZēŁĸ;ŚçzĪJçZt' æŌēēĀZāŁāāĀĆ

ēĀŇçZōāL'■āžŤĸŤĭāĪJāāzŁčŽĎ flannel ĸ;ŚçzĪJæRŠāzūĭĭjŇæŸrāy■æŤræŇAēŁZĸg■āy■āŘŇéZĚĸ;đ' āŚ;

calico æŤræŇAāĪJřāĪĀāĹēĒ■ĭĭjŇāzšæŤræŇAāy■āŘŇéZĚĸ;đ' āŚ;āŘ■ĸl'žéŮt' çŽĎĸ;ŚçzĪJéŽŤççz■Ūç
BGP äyĹ'āsČĸ;ŚçzĪJēĭŇāRŠĭĭjŇæĀgēČ;ærŤ flannel āĭjžāĀĆ

āzšāRřāzēā;ŁčŤĭ flannel æĭēāĀZĸ;ŚçzĪJçōāçŘĒĭĭjŇāĒ■āŌL'ēčĒ calico
āžĒāžĒāĀZēZĚĸ;đ' āŚ;āŘ■ĸl'žéŮt' ĸ;ŚēūřēŽŤççz■ŪçŤĕĭĭjŇēŁZĸg■æRēĒ■æŪzæāĹāĀĆ

16.4 16.4 K8SçĭŚçzĪJæŇšæL'S

æL'ĀæĪJĹ' POD èŁđæŌēāĹĭĭjŇæĪJæĪJž cni0 æŌēāRčēŁZāyĭŁ;ŚçzĪJĭĭjŇcni0
æŌēāRčāRŠāĠžĸZĎæŁēæŪĠāĹrēĭ; flannel.1 èŁZāyĭāŌēāRčĭĭjŇēŁZāyĭāŌēāRčārĒæŁēæŪĠārĀēčĒāyžéŽg

- æšēĸĪJŇæĪJæĪJžĎæŌēāRč

```
1: lo: # æĪJŇāĪJřāžđčŌř
2: ens33: # äŸzæĪJžĸL'ĲčŘĚĸ;Śā■ā
3: docker0: # docker_
→éžŸēōđ' çŽĎæāēæŌēĸ;ŚçzĪJĭĭjŇāĪJĹ' k8s äy■æŪāçŤĭāRřāzēēāĹāēŽđ'
4: dummy0: #
5: kube-ipvs0: #
6: flannel.1: # flannel_
→èŽŽæŇšç;Śā■āĭĭjŇāRĀēčĒēŽgēĀšæŁēæŪĠ
7: cni0: # æL'ĀæĪJĹ' āōzāZĭāđ' ĎāžŌēŁZāyĭŁ;Śæāē
8: veth0c014b8b@if3: # āōzāZĭŁčŽĎĸ;Śā■āēŁđæŌēāĹr cni0
9: veth97c048e5@if3: # āōzāZĭŁčŽĎĸ;Śā■āēŁđæŌēāĹr cni0
11: vethd2f0bf2b@if3: # āōzāZĭŁčŽĎĸ;Śā■āēŁđæŌēāĹr cni0
12: veth648a500f@if3: # āōzāZĭŁčŽĎĸ;Śā■āēŁđæŌēāĹr cni0
```

- äyŇēĭ; bridge-utils āŇĒā;ŁčŤĭāŚ;āzd' brctl show cni0 æšēĸĪJŇ cni0 æŌēāRč

bridge	name	bridge id	STP	enabled	interfaces
cni0		8000.9a6ec95f8285	no		veth0c014b8b veth648a500f veth7a3f56b7 veth97c048e5 vethd2f0bf2b

16.5 16.5 flannel

flannel æÝřäýÄäýläýŞäýž kubernetes áóŽÁŁűçŽDäýL'ásĆç;ŞçzIJèğçâEşæŰzæaŁiijNäýžèçAçTlāžŌèğçâ

16.5.1 16.5.1 flannel áũěä;IjæIqâijR

- flannel.1 èŁŽäýlèŽŽæŃşç;Šâ■æTřæŃAđ'Žçg■äijäè;ŞæIqâijRiijŽVxLANāĀhost-gwāĀDirectroutingāĀAudp


æIqâijR	řäýŃçž■
VXLAN	■;ŁçTl VxLAN ä;IJäýžéŽgéAŞřAèçĚæŁæŰĜ
host-gw	äý■ä;ŁçTlāRāāŁäç;ŞçzIJiijNèĀŃæÝřäIjläýžæIJçŽDèŭřçTšèaļäý■āLZāžžāLřāĚūāžŰäýžæIJž subnet çŽDèŭřçTšæIaçŽōiijNæĀğèÇ;èŁCāè;iiijNçijžéŽūæÝřiijZæL'ĀæIJL' node èŁĆçĆžāŁĚēāzād'DāžŌāRŃNäýÄäýlāžŃāsĆç;ŞçzIJäý■āĀĆ
DirectRouting	ā;ŞäýžæIJžä;■āžŌāRŃNäýĀā■Rç;ŠæŰūāRřçTlçŽt'æŌèèŭřçTšiiijNäý■āIJlāŽdéĀĀLř VxLANāĀĆ
UDP	çŽt'æŌěä;ŁçTl UDP ā■RèōōiijNæĀğèÇ;āũō

16.5.2 16.5.2 VXLAN éĀŽäŁæŁĜçÍŃ

Flannel VXLAN áóđèťlāýŁæÝřäýĀçg■ āĀIJèçEçŽŰç;ŞçzIJ(overlay network)āĀI iijNāžšārsæÝřāřETCPæTřæ■ōāŃĚèçĚāIJlāRēäýĀçg■ç;ŞçzIJāŃĚéĜŃéIçèŁZēāŃēŭřçTšè;ŃāRŠāŠŃéĀŽāŁæĀ VPCāŠŃĜCEèŭřçTšç■L'æTřæ■ōè;ŃāRŠæŰžāijRāĀĆ

- flannel VXLAN éĀŽäŁæŁĜçÍŃ

āIJl K8S äýŁ POD äýŌ POD æÝřçŽt'æŌèéĀŽèŁĜāřzæŰžçŽD IP āIJřāIĀèŁZēāŃéĀŽāŁæçŽDriijNPOD āRŠāĜžçŽDæŁæŰĜçžRèŁĜ cni0 ç;ŠæāçāLřèŁ; flannel iijNflannel āřEæŁæŰĜāřAèçĚäýŁäýĀāšĆ VxLAN çŽDèçŰéČliijNāđ'ŰāsČāRlĚcŃāřAèçĚäýĀāšĆ UDP ā■RèōōçŽDèçŰéČliijNāRŠéĀAçžŽæIJñæIJçL'l'çRĚç;Šâ■āiijNæIJñæIJçL'l'çRĚç;Šâ■āāRlāřE flannel āRŠèŁĜæIèçŽDæŁæŰĜāđ'ŰāsČāřAèçĚäýŁ IP éçŰéČlāŠŃāžēāđ'lç;ŠäýgēçŰéČliijLMACiijL'çTšç;Š node èŁĆçĆžæTūāLřæŁæŰĜiijNāĚEäýāRŠçŌřæÝřäýĀäýl VxLAN çŽDāŃĚiijNæŃEæŌL' IP éçŰéČlēĀAçžŽ flannel āžTçTlçlNāžRiijNflannel æŃEæŌL VxLAN éçŰéČlāžūāřEāĚEēČlçŽDæTřæ■ōāRŠéĀAçžŽiijNcni0 ç;ŠæāçiijNcni0 æTūāLřāRŌè;ŃāRŠçžŽ PODāĀĆ



(continues on next page)

(continued from previous page)

node ç;šçžIJ	nodeç;šçžIJ	node ç;šçžIJ	VxLan	POD
→MAC	POD IP	data		
→äÿgëëŮéČÍMAC	IPëëŮéČÍ	UDP ëëŮéČÍ	ëëŮéČÍ	
→ëëŮéČÍ	ëëŮéČÍ	Payload		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+				
→-----+				

16.5.3 flannel éČĺç;šæŮžaijŘ

1. aIJÍ k8s éŽEç;đ' aŘřaŁlāL' ■■■ijNflannel çŽt' æŮëéČĺç;šāLřèŁCçĆzäÿŁiijNä;IJäÿzäÿÄäÿlāōŁæLd'èŁŽçÍN

äzzä;TäÿÄäÿléČĺç;šäžE kubelet çŽDèŁCçĆzéČ;āžTèrëéČĺç;š flannel
→iijNāžääÿž kubelet èëAāšāLl' flannel äÿž POD èöç;öç;šçžIJæŮëāŘč

2. ä;ŁçŤÍ kube-admin çŽt' æŮëāřE k8s èĞlāušçŽDçžDäzúāNĚæNñ flannel èŁŘëāNāIJÍ k8s äžNäÿŁçŽDéÍŽæĀA PODāĀĆ

āŁĚéāzéĚ■ç;öäÿžāĚšāžñ node èŁCçĆzç;šçžIJāŘ■çğřçl' žéŮt' çŽD
→PODiijNæL'Āäžè flannel POD æŮğāLúāžlāÿž DaemonSetāĀĆ

16.5.4 flannel éĚ■ç;öæŮĞäzú

- éĚ■ç;öæŮĞäzúéĀL'éāžāŘnāžL'

```
{
  "Network": "10.244.0.0/16",          // flannel ä;ŁçŤÍçŽD CIDR
  →æāijjāijŘçŽDç;šçžIJāIJřāiĀiijNçŤlāžŌäÿž POD éĚ■ç;öç;šçžIJāLšèČ;
  "SubnetLen": 24,                      // æŁL Network
  →ālĠāLēäÿžā■Řç;šäçŽāŘD node
  →èŁCçĆzä;ŁçŤlāŮüiijNä;ŁçŤlād'žéŁççŽDæŌl'çāāāLĠāLēiijNéžÿèöd'äÿž 24
  "SubnetMin": "10.244.10.0/24",       // çŤlāžŌāLēéĚ■çžŽ node
  →çŽDā■Řç;šèŁūāğNāIJřāiĀiijNāžŌèŁžäÿłç;šçžIJāijĀāğNāLēéĚ■ç;šçžIJ
  "SubnetMax": "10.244.255.0/24"       // çŤlāžŌāLēéĚ■çžŽ nide
  →çŽDā■Řç;šçžšāİSā;■ç;öiijNèŁžäÿlæÿræIJĀād'ğāLēéĚ■çŽDç;šèŮr
  "Backend": {                          // æNĞæŮ POD äÿŮ POD
  →èŮlèŁCçĆzéĀžāŁæŮūāĀžä;ŁçŤlçŽD flannel āŮëä;IJāLāāijŘ
    "Type": "vxlan",                     // āŮëä;IJāLāāijŘ
    "Directrouting": true                 //
  →æÿrāŘëä;ŁçŤlçŽt'æŮëèŮrçŤsāLāāijŘ
  }
}
```

- flannel æL'ŸçŌāāLř k8s äÿŁçŽDéĚ■ç;öæŮĞäzúüijNād'ĐäžŌ kube-flannel-cfg èŁŽäÿl configmap äÿ■āĀĆ


```
kubectl get configmap kube-flannel-cfg -n kube-system -o json
```

16.5.5 16.5.5 ãŁõæŤzãûëä;IJælaaijR

- ãŁõæŤz flannel ãûëä;IJælaaijRiiijNæûzãŁã Directrout-
ingiiijNèŁZäyŁæŞ■ä;IJãžŤèřèãIJlãŁŽãŁŽéČlç;šãõŇ k8s
éZEç;d'æUúãÄZãŁõæŤziiijNæÕlè■ŘãŁõæŤz

```
kubectl edit configmap kube-flannel-cfg -n kube-system
```

```
"Backend": {
  "Type": "vxlan",
  "Directrouting": true
}
```

- æŞëçIJNæIJnæIJžèûřçŤšèaÍ

```
ip route show
```

```
default via 172.16.100.254 dev ens33 proto static metric 100
10.244.1.0/24 via 10.244.1.0 dev ens33 # ãŁĚéazäÿž dev_
→çŁ'ł'çŘEç;Šã■æÕëãŘçiiijNãŘèãŁŽ Directrouting æšæIJL'èõç;õæŁŘãŁS
10.244.2.0/24 via 10.244.2.0 dev ens33 # ãŁĚéazäÿž dev_
→çŁ'ł'çŘEç;Šã■æÕëãŘçiiijNãŘèãŁŽ Directrouting æšæIJL'èõç;õæŁŘãŁS
172.16.100.0/24 dev ens33 proto kernel scope link src 172.16.100.
→101 metric 100
172.17.0.0/16 dev docker0 proto kernel scope link src 172.17.0.1
```

16.6 16.6 Calico

Calico ãŁZãžzãŠŇçõaçŘEâijÄäyŁæL'AãžşçŽDäyL'ãšČç;ŞçzIJ(äy■éIJÄèeA over-
lay)iiijNæřRäyŁãõžãŽlãijŽãŁEëĚ■äyÄäyŁãŘrèûřçŤšçŽD ipãÄČçŤšãžÕéÄŽãŁæUüäy■éIJÄèeAèğçãNĚãŠNãř

ãřRèğDælaéČlç;šæUüãŘräžèéÄŽèŁĜ bgp client çŽt' æÕëãžŞèAŤrijNãd' gèğDælaäyNãřféÄŽèŁĜæNĜãõ
BGP route reflector ælĚãõNæŁŘiiijNèŁŽæãüãŁĚërAæL' ÄæIJL'çŽDæŤřæ■õætAéĜRéČ;æŸféÄŽèŁĜ
IP èûřçŤšçŽDæŮžãijRãõNæŁŘãžŞèAŤçŽDãÄČ

Calico ãŞžãžÕ iptables èŁŸæRRä;ŽãžEäyřãřNèÄŇçAŤæt'zçŽDç;ŞçzIJ Pol-
icyiiijNãĚlĚrAéÄŽèŁĜãŘDäyŁèŁČçCzäyŁçŽD ACLs ælĚæRRä;Ž Workload
çŽDãd'ŽçğşæŁüéŽŤçzãÄAãõL'ãĚlçzDãžèãRLãĚüãžŮãŘrè;ŁæĂğéŽŘãŁLúç■L'ãŁşèČ;ãÄČ

æIJL'äyŁæŮřçŽDèazçŽŽõiiijŽcaneliiijNãõČéZEãŘLãžE flannel ãŠŇ calico çŽDäijŸçČzãÄČ

- æşlæĐŘ

Calico çŽDãL'■äy■æŤřæNãüëä;IJãIJl iptables äyŇçŽD kube-proxyiiijNäyNélcãžŇçz■
canal ç;ŞçzIJç■ŮçŤççŽDä;ŁçŤl

16.6.1 16.6.1 aóL'èčĚ canal

- äyÑèj;æyĚā■TæŮGäzūiijÑéIJĀèèAçfzácŽ

```
kubectl apply -f https://docs.projectcalico.org/v3.6/getting-
started/kubernetes/installation/hosted/canal/canal.yaml
```

16.6.2 16.6.2 æyĚā■TæŮŽäZL'

- æyĚā■TæäijäijRiijÑèfèègAijŽkubectl explain networkpolicy

```
egress                                <[]Object>      # āĜžčnžèġDāLžčŽDāržèsaāLŮèaÍ
  ports                              <[]Object>      # çŽōæāĜčnřāRčçŽDāržèsaāLŮèaÍ
    port                             <string>         #
    →æTřā■Ůā;ćaijRāLŮèĀĚæYřāS;āR■çŽDčnřāRč
    protocol                         # ā■Rèóó TCPāĀUDP
  to                                  <[]Object>      # çŽōæāĜāIJřāīĀāržèsaāLŮèaÍ
    ipBlock                          <Object>         # äyĀçžD IP āIJřāīĀ
      cidr                           <string>         # CIDR èaÍçd'žçŽD IP èNČāžt'
      except                         <[]string>        # æŌŠéŽd' CIDR
    →äy■çŽDæSŘäžžāIJřāīĀ
    namespaceSelector                <Object>         # āR■çġřçl'žéŮt'éĀL'æNl'āŽÍ
    podSelector                     <Object>         # POD
    →éĀL'æNl'āŽÍiijNçŽōæāĜāIJřāīĀāRřäžèäzSæYřäyĀçžD POD
ingress                              <[]Object>      # āĚèčnžèġDāLžčŽDāržèsaāLŮèaÍ
  from                               <[]Object>      # æžRāIJřāīĀāržèsaāLŮèaÍ
    ipBlock                          <Object>         # äyĀçžD IP āIJřāīĀ
      cidr                           <string>         # CIDR èaÍçd'žçŽD IP èNČāžt'
      except                         <[]string>        # æŌŠéŽd' CIDR
    →äy■çŽDæSŘäžžāIJřāīĀ
    namespaceSelector                <Object>         # āR■çġřçl'žéŮt'éĀL'æNl'āŽÍ
    podSelector                     <Object>         # POD
    →éĀL'æNl'āŽÍiijNæžRāIJřāīĀāzžSāRřäžèæYřäyĀçžD POD
    ports                            <[]Object>      # POD
    →èĠāũsçŽDčnřāRčiiijNèaÍçd'žæŌġāLŮèĠāũsçŽDčnřāRčæYřāRēāRřäžèèčnèøŁéŮöiijNçŽDārž
    port                             #
    →æTřā■Ůā;ćaijRāLŮèĀĚæYřāS;āR■çŽDčnřāRč
    protocol                         # ā■Rèóó TCPāĀUDP
podSelector                          <Object>         # POD
    →éĀL'æNl'āŽÍāEşāōžèġDāLžāžTçTīāIJlāşlāžž POD äyĻ
policyTypes                          <[]string>       # āRřäžèæYř "Ingress", "Egress"
    →", æLŮèĀĚ "Ingress,Egress"
    →iijNèaÍçd'žæT;èaŊæzaèũşèŁZäžžèġDāLžèøŁéŮö
```

16.6.3 16.6.3 policyTypes

- éeŮāĚĹāōŽäZL' āR■çġřçl'žéŮt'

```
kubectl create namespace dev
kubectl create namespace prod
```

- `āĬĬāyĎ'āyĭāŚ;āŘ■čĭ'žēŮt'āĽĖāĽnāĽZāzžāyĀāyĭ` POD

```
apiVersion: v1
kind: Pod
metadata:
  name: pod1
  namespace: dev
  labels:
    app: myapp
spec:
  containers:
  - name: myapp
    image: ikubernetes/myapp:v1
```

```
---
apiVersion: v1
kind: Pod
metadata:
  name: pod1
  namespace: prod
  labels:
    app: myapp
spec:
  containers:
  - name: myapp
    image: ikubernetes/myapp:v1
```

```
kubectl apply -f pod-a.yaml -n dev
```

- `æŇŠčziæĽ'ĀæĬĬ' dev čĭ'žēŮt'čŽĎæĽæŮĜ`

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: deny-all-ingress
  namespace: dev
spec:
  podSelector: {} # {}
  → čĭ'žčŽĎĖĀĽ'æŇĭ'āžĭèāĭčĎ'žēĀĽ'æŇĭ'āĖĭéČĭ
  policyTypes:
  - Ingress # æŇĜæŸŎ Ingress èĝĎāĽZčŤSæŤĬiijŇāŇzéĖ■
  → Ingress āřĖèčnæŤ;èāŇiijŇāēČæĎĬJæšāāōžāžĽ' Ingress
  → āĽZāy■èČ;āŇzéĖ■æĽ'ĀæĬĬ' iijŇāi jžæŇŠčziāĖĭéČĭ
  # policyTypes æšæĬĬ' Egress
  → èāĭčĎ'žāy■æŎĝāĽŮ Egress iijŇézŸèōĎ'äyžāĖĀèōy
```

- `āĬĬæŇĜāōŽāŚ;āŘ■čĭ'žēŮt'āžŤčŤĭèĝĎāĽZæŮĜāzŮ`

```
kubectl apply -f deny-all-ingress.yaml -n dev
```

- `kubectl get networkpolicy -n dev`

```
kubectl get networkpolicy -n dev
```

- `kubectl get pods -n dev -o wide`

```
kubectl get pods -n dev -o wide
```

```
curl 10.244.1.2
```

- `kubectl get pods -n dev -o wide`

```
kubectl get pods -n dev -o wide
```

```
curl 10.244.2.2
```

- `kubectl get pods -n dev -o wide`

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-80-ingress
  namespace: dev
spec:
  podSelector:
    matchLabels:
      app: myapp
  ingress:
    - from:
      - ipBlock:
          cidr: 10.244.0.0/16
        except:
          - 10.244.1.2/32
      - podSelector:
          matchLabels:
            app: myapp
    ports:
      - port: 80
      - port: 443
      protocol: TCP
  policyTypes:
    - Ingress
```

- æšćIJNèĝDǎLŽ

```
kubect1 get networkpolicy -n dev
```

- æŇŠčziǎĜžæǎLčŽǾL'ĂæIJL'èrũæśĆ

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: deny-all-egress
  namespace: prod
spec:
  podSelector: {} # {}
  ↳ çl'žčŽǾL'æŇl'ǎžl'èǎlčd'žéǎL'æŇl'ǎĚléčl'
  policyTypes:
    - Egress # æŇĜæŸŌ Egress èĝDǎLŽčŤSæŤLii jŇǎŇzéĚ
    ↳ Egress ǎřĚèćnǎŤč;èǎŇii jŇǎĕĆǎđIJæšǎǎŏžǎžL' Egress
    ↳ ǎlŽǎÿ;ǎŇzéĚæL'ĂæIJL'ii jŇæŇŠčziǎĚléčl'
    # policyTypes æšæIJL' Ingress
    ↳ èǎlčd'žǎÿæŌĝǎlŭ Egress iijŇézŸèŏd'ǎÿžǎĚAèŏÿ
```

CHAPTER 17

■AäyČ èřČāžęç■ÚçȚě

Master äyzèeAæYřèfRèaŇéZÉç;d'çŽDæŎğĀLūāzşéİcçzDäzűçŽDiiŇNapiserverãĀAscheduleraĀAcontro
èfYä;İetŰäyŎ etcd èfZæăüçŽDă■YăCİèŁCçCzãĀC

kubeadm éČlç;şçŽDÉZÉç;d'äijZârE Master çŽDæŎğĀLūçzDäzűéČ;èfRèaŇäyžéİZæĀA
POD äžEiiŇNäzŎæIJñet'İæİèèšiiŇNèfZäzZçzDäzűârşæYřèfRèaŇâIJİ
Master èŁCçCzäyŞäyžéZÉç;d'æRŘä;ZæIJ■âŁaçŽDèfZçİŇiiŇNMaster
æYřäy■et'şet'çèfRèaŇăüëä;IJet'şè;çŽDăĀC

node èŁCçCzæYřet'şet'çèfRèaŇăüëä;IJet'şè;ç ŽDçŽyâĒşèŁCçCziiŇNçTİæLūâRİéIJĀèeAârEèfR
Master ârsâRřäžëäžEiiŇNçTİæLūâŰæIJĀâĒşâfCèfRèaŇâIJİâŞtäyİ node èŁCçCzäyLiiŇN-
Master æȚt'ârĬläžEæL'ĀæIJL' node äyžäyĀäyİèŽZæŇşçŽDètDæžRæšăăĀC

17.1 17.1 PODăĬZăžžæȚAçİŇ

çTİæLūâĬZăžžçŽDäzžâŁæIJçZĬâžTèrèèfRèaŇâIJİâŞtäyİ node
èŁCçCzäyLiiŇNæYřçTş Master èŁCçCzäyŁçŽD scheduler âĒşâŏŽçŽDiiŇNèĀŇ scheduler
äžşæYřâĒAèŏyçTİæLūâŏŽăžL'âŏCçŽDăüëä;IJçL'žæĀğçŽDiiŇNéžYèŏd'æČĒâĒtäyŇiiŇNæĬSăžŇæşqæIJL'âŏŽ
scheduler èřČāžęăZİăĀC

â;ŞæĬSăžŇă;ŁçTİ kubectİ describe pods myapp æşęçIJŇ POD
ăfææAřæŰăăZiiŇNäijZæIJL'äyĀäyİ Events â■Űæŏtäy■iiŇNæIJL'âĒşăžŎèřČāžęçzŞædIJçŽDçŽyâĒşâfææAřă

Scheduler äijZăžŎäijŰad'ŽçŽD node èŁCçCzäy■æŇSéĀL'âĞžçŇæĀRĬ
POD èfRèaŇèeAæşCçŽDèŁCçCziiŇNçDūâRŎârEéĀL'âŏŽçŽD node
èŁCçCzăfææAřèŏrâ;ȚâIJİ etcd äy■iiŇNkubetlet âğŇçZĬ waitch çİĀ apiserver
äyŁçŽDâĒşăžŎæIJñèŁCçCzçŽDăfææAřăRŸâŇŰiiŇNkubetlet âřsäijZăŎž apiserver
èŎăâRŰâĒşăžŎârŸâŇŰăfææAřçŽDèĒ■ç;ŏæyĒă■ȚiiŇNæăžæ■ŏéĒ■ç;ŏæyĒă■ȚçŽDăŏŽăžL'âŎZăĬZăžž
PODăĀC

17.2 17.2 Service

Service is a Kubernetes abstraction that defines a logical set of Pods that can be discovered and accessed by other Pods. Service is a logical set of Pods that can be discovered and accessed by other Pods. Service is a logical set of Pods that can be discovered and accessed by other Pods.

Service is a Kubernetes abstraction that defines a logical set of Pods that can be discovered and accessed by other Pods. Service is a logical set of Pods that can be discovered and accessed by other Pods. Service is a logical set of Pods that can be discovered and accessed by other Pods.

17.3 17.3 etcd

- etcd is a distributed key-value store that is used by Kubernetes to store cluster state.
- etcd is a distributed key-value store that is used by Kubernetes to store cluster state.

17.4 17.4 Scheduler

- The scheduler is responsible for scheduling Pods onto nodes. It takes into account various constraints and preferences to ensure that Pods are scheduled onto suitable nodes.
 - The scheduler is responsible for scheduling Pods onto nodes. It takes into account various constraints and preferences to ensure that Pods are scheduled onto suitable nodes.
 - The scheduler is responsible for scheduling Pods onto nodes. It takes into account various constraints and preferences to ensure that Pods are scheduled onto suitable nodes.
- POD is a container that runs on a node. It is managed by the scheduler.

```
nodeSelector # selector to select nodes for scheduling
nodeSelector # selector to select nodes for scheduling
nodeSelector # selector to select nodes for scheduling
```

- POD is a container that runs on a node. It is managed by the scheduler.

nodeSelector is a label selector that is used to select nodes for scheduling.

POD is a container that runs on a node. It is managed by the scheduler. POD is a container that runs on a node. It is managed by the scheduler.

POD is a container that runs on a node. It is managed by the scheduler. POD is a container that runs on a node. It is managed by the scheduler.

Taints are used to mark nodes as unsuitable for certain Pods.

Tolerations `iiJLæšaçĆzãóžãf■iiJL'iiJZäYÄäYl` POD `èÇ;ãd'šãóžãf■` node
`äYŁçŽDæšaçĆziiJNæÇCædIJèfRèaÑèfGçlNäY■` node `ãGžçŒræŪrçŽDæšaçĆziiJNéCčázL`
 POD `ãRfäzë`

`él'séÄRPODiiJŽnode` `çZŽäYÄäYlèŽRãóŽçŽDæŪúéŪr'iiJNèol'` POD
`ççzãijAèfZäYlèLÇçCzãÄC`

17.5 17.4 éCĐÉÄL'ãŽăçt'ă

`äYNéIççŽDèCĐÉÄL'æIäzüzéIJÄèAæzæùšæL'ÄæIJL'çŽDèCĐÉÄL'æIäzüzæL'■ãRfäzëéÄŽèfGèCĐÉÄL'`

1. CheckNodeConditionPred

```
æčÄæSëèLÇçCzæYrãRææ■čäYy
```

2. GeneralPredicates

ãRç■ŪçãPçTl	
Host-Name	<code>æčÄæšëäYzæIJzãR■çğræYrãY■æYr</code> <code>pod.spec.hostname æNĠãóŽçŽD</code> <code>nodeName</code>
Pod-Fit-HostPorts	<code>æčÄæšë</code> <code>Pod æEÆæfRäYÄäYlãóžãZl</code> <code>pods.spec.containers.ports.hostPort</code> <code>æYÈã■TæYrãRæãùšècñãEüãóČãóžãZl■äçTl</code> <code>iiJNæÇCædIJæIJL'æL'ÄéIJÄçŽD</code> <code>HostPort äY■æzæùšéIJÄsĆiiJNéCčázL</code> <code>Pod äY■èÇ;èrČäzæãLrèfZäYlæYzæIJzäYl</code>
MatchNodeSelector	<code>æčÄæšë</code> <code>POD æóžãZlæYlæãóžãZl'äZÈ</code> <code>pods.spec.nodeSelector æšëçIJN</code> <code>node æãGç■æYrãRæèÇ;ãd'šãNzéÈ■</code>
Pod-Fit-Resources	<code>æčÄæšë</code> <code>node æYrãRææIJL'èüšãd'šçŽDètDæžRèfRèaÑæ■d'</code> <code>POD çŽDãšzæIJnèAæšC</code>

3. NoDiskConflictiiJLéZŸèód'æšææIJL'ãRçTlriJL

```
æčÄæSë pod æóžãZl'çŽDã■YãCíæYrãRæãIJl node èLÇçCzäYlæ;fçTlãÄC
```

4. PodToleratesNodeTaints

```
æčÄæSëèLÇçCzçŽDæšaçĆz nodes.spec.taints æYrãRææYr POD_
→æšaçĆzãóžãf■æYÈã■TäY■ pods.spec.tolerations çŽDã■RèŽÈ
```

5. PodToleratesNodeNoExecuteTaints

```
æčÄæSë pod æYrãRæãóžãf■èLÇçCzäYlæIJL' NoExecute æšaçĆzãÄCNoExecute_
→èfZäYlæšaçĆzæYrãTæDŔæÄIãšCãÄæÇCædIJäYÄäYl pod_
→äYlèfRèaÑãIJlæYÄäYlæšææIJL' æšaçĆzçŽDèLÇçCzäYlæRŒiiJNèfZäYlèèLÇçCzãRlçzãZãäYlæš
→NoExecute_
→èãlçd'žèfZäYlæŪrãLäæšaçĆzçŽDèLÇçCzäYlæjZél'séÄRãEüäYlæ■čãIJlèfRèaÑçŽD_
→podiiJZäYlæLä NoExecute äY■äijZél'séÄRèLÇçCzäYlèfRèaÑçŽD_
→podiiJNèãlçd'žæŒæRŪæŪcãLŔãžNãódiiJNèfZäYrèZŸèód'ç■ŪçTlæÄC
```


(continued from previous page)

6. CheckNodeLabelPresenceiijLézYèód'æšæIJL'âRřçTliijL'

```
æčĀæSëèŁĆćZăyŁæŃĠăôŽăăĠç;çŽĐăYăIJlăĀġiijŃăęĆăđIJèŁĆćZăIJL'podăŃĠăôŽçŽĐăăĠç;ç
```

7. CheckServiceAffinityiijLézYèód'æšæIJL'âRřçTliijL'

```
ăyĀăyŁ Service äyŃăRřăžăIJL'ăđ'ŽăyŁ PODiijŃăRŤăęĆăfŽăžŽ POD
→éČ;èĤRëăŃăIJl 1ăĀĀ2ăĀĀ3 æIJžăŽlăyŁiijŃăĀŃăšăIJL'èĤRëăŃăIJl
→4ăĀĀ5ăĀĀ6 æIJžăŽlăyŁiijŃăĀĀžŁCheckServiceAffinity
→ăřśăăġčđ'žăŮřăĀăĀĒëçŽĐ POD éČ;éŽĚăy■èĤRëăŃăIJl 1ăĀĀ2ăĀĀ3
→æIJžăŽlăyŁiijŃăĤŽăăŮéŽĚăy■ăë;ăđ'ĐăYřăyĀăyŁ Service äyŃ POD
→ăžŃéŮt'ăĚĚéČléĀŽăĤăçŽĐăŤŁçŌĠăRŸénYăžĚăĀĈ
```

8. MaxEBSVolumeCount

```
çăŌăĤİăŮšăŃĆè;;çŽĐăžŽél'ŋéĀŁ EBS
→ă■YăĆİă■ăy■ëŮĚëĤĚëç;çŽĐăIJĀăđ'ġăĀiijijŃéžYèód'39
```

9. MaxGCEPDVolumeCount

```
çăŌăĤİăŮšăŃĆè;;çŽĐGCEă■YăĆİă■ăy■ëŮĚëĤĚëç;çŽĐăIJĀăđ'ġăĀiijijŃéžYèód'16
```

10 MaxAzureDiskVolumeCount

```
çăŌăĤİăŮšăŃĆè;;çŽĐAzureă■YăĆİă■ăy■ëŮĚëĤĚëç;çŽĐăIJĀăđ'ġăĀiijijŃéžYèód'16
```

11. CheckVolumeBinding

```
æčĀæSëèŁĆćZăyŁçŽĐ PVC æYřăRëëćŋăĤŋçŽĐ POD çžŚăŌŽăžĚ
```

12. NoVolumeZoneConflict

```
æčĀæSëçžžăŌŽçŽĐ zone (æIJžăĤĤ)
→éŽŘăĤăĤ'■ăRŘăyŃiijŃăčĀæSëăęĆăđIJăIJlăđ'ăyžæIJžăyŁéČİç;š POD
→æYřăRëă■YăIJlă■ăĤĚšçĤă
```

13. CheckNodeMemoryPressure

```
æčĀæSë node èŁĆćZăyŁăĚĚă■YăYřăRëă■YăIJlăŌŃăĤŽ
```

14. CheckNodeDiskPressure

```
æčĀæSëçČĀçŽY IO æYřăRëăŌŃăĤŽèĤĠăđ'ġ
```

15. CheckNodePIDPressure

```
æčĀæSë node èŁĆćZăyŁ PID èŤĐăžŘăYřăRëă■YăIJlăŌŃăĤŽ
```

16. MatchInterPodAffinity

```
æČǺæŠě Pod æŸřǺŘęæzæűşăžšǺŠŇæĀğæĹŪèĀĔǺǺ■ăžšǺŠŇæĀğ
```

17.6 17.5 äijŸéĀĹ'ǺĜǺæŢř

ǺĪĴǺřǺŸǺĹèĹČĴČzǺĹ'ğęǺŇäijŸéĀĹ'ǺĜǺæŢřǺijŇǺřĒçzŞǺđĴǺřǺŸǺäijŸéĀĹ'ǺĜǺæŢřçŸŸǺĹäǺijŇǺǺŪǺĹĒ

1. least_requested.go

```
éĀĹ' æŇĹ' æűĹèĀŪæĴĴǺřǺčŽĐèĹČĴČzǺijĴǺăžæ■ōçĹ' žéŪşǺřŢçŌĜěřĐäijř_
→cpu (æĀžǺōžéĜŖ-sum (Ǻűşă; ĲçŢĹ) *10/æĀžǺōžéĜŖ) iijĴ'
```

2. balanced_resource_allocation.go

```
ǺĴĜęǺǺęĐǺžǺčŽĐǺ; ĲçŢĹæűzǺijŖǺijŇǺǺĴçd' žăžę cpu_
→ǺŠŇǺĒĒǺ■ŸǺ■ǺçŢĹçŌĜçŽĐçŸŸèĴŖçĹŇǺžęǺ; ĴǺŸžěřĐäijřǺǺĜǺĜĒiijŇǺžŇǺĒĒǺ■ǺçŢĹèűĴǺŌěèĴŖ
```

3. node_prefer_avoid_pods.go

```
çĴŴŇèĹČĴČzǺŸřǺŘęæĴĴĴ' æşĴèğçǺǺǺǺǺř "scheduler.alpha.kubernetes.io/
→preferAvoidPods" _
→ǺĀčǺşǺæĴĴĴ' èĴžǺŸǺşĴèğçǺǺǺǺǺřǺijŇǺěřŢ' æŸŌèĴžǺŸǺĹèĹČĴČzǺŸřéĀčǺǺĴĴèĴŖǺŇǺĴǺŸǺ_
→POD çŽĐǺĀč
```

4. taint_toleration.go

```
ǺřĒ pods.spec.tolerations äŸŌ nodes.spec.taints_
→ǺĴŪǺǺĴǺǺžéŖžǺǺŇǺŇzéĒ■ǺžęæčǺæŠěiijŇǺŇzéĒ■çŽĐǺĴǺçŽŌěűĴǺđ' žǺijŇǺǺŪǺĹĒèűĴǺ; ŌǺĀč
```

5. selector_spreading.go

```
æŠěæĴ' çǺ; ŞǺĴ' ■ POD ǺřžǺşǺǺřžǺžŢçŽĐ_
→serviceǺǺstatefulsetǺǺreplīcatset_
→ç■Ĵ' æĴ' ĀǺŇžéĒ■çŽĐǺǺĜç■; éĀĴ' æŇĹ' ǺžĴǺijŇǺĴĴèĹČĴČzǺŸŸèĴŖǺŇçŽĐǺŸęæĴĴĴ' èĴžǺşǺşǺĜç■
→POD ęűĴǺřǺşǺ; ŪǺĴĒèűĴĒǺŸǺijŇǺĴžǺşǺűçŽĐ POD äijŸéĀĴ' èčŇéĀĴ' ǺĜžǺĀč_
→èĴžǺřǺşǺŸřǺěřŢ' æĴŖǺžŇǺęǺǺĴǺǺŇǺŸǺǺŸǺǺǺĜç■; éĀĴ' æŇĹ' ǺžĴǺŸŸǺĴŖǺŇçŽĐ_
→POD æŢčǺijǺ (spreading) ǺĴŖǺđ' žǺŸǺĹèĹČĴČzǺŸŸǺĀč
```

6. interpod_affinity_test.go

```
éĀ■ǺŌĒ POD_
→ǺřžǺşǺǺžšǺŠŇæĀğçŽĐǺĴǺçŽŌiijŇǺžűǺǺřĒéččǺžžèç; Ǻđ' ŞǺŇžéĒ■ǺĴŖèĹČĴČzǺĴčĒççŽŸǺĴǺiijŇ
```

7. most_requested.go

```
ǺǺĴçd' žǺř; ǺŖŖèç; çŽĐǺĴǺŸǺǺŸǺĹèĹČĴČzççŽĐęĐǺžǺŖǺĒĴçŢĹǺŌŇiijŇǺĴžǺŸǺǺŠŇ_
→least_requested çŸŸǺř■iijŇǺžŇǺĒĒǺŸǺ■èç; ǺŖŇæűűǺ; ĲçŢĹǺĀč
```

8. node_label.go

```
æǎžæ■óèŁĆĆzæŸráŘęæŃæIJL' æǎĞç■; iijŃäŸ■ǎĚşǎŁČæǎĞç■;æŸräzǎžŁiijŃæİëèrĎäijřǎŁEǎŤřǎ
```

9. image_locality.go

```
èǎÍçd' žæǎžæ■óæžǎèŸşǎ;şǎL'■ POD┐  
→ǎřžèşǎéIJǎşĆçŽĎǎŸşæIJL' éŤIJçŽĎä;şçğřǎd' ğǎřŘǎžŃǎŠŃæİëéǎL' æŃl' èŁĆĆzçŽĎǎǎĆ
```

10. node_affinity.go

```
æǎžæ■ó POD ǎřžèşǎäŸ■çŽĎ┐  
→nodeselector iijŃǎřžèŁĆĆzèŁžèǎŃǎŃzéĚ■ǎžęǎčǎşĚiijŃèč;ǎd' şǎŁŘǎŁşǎŃzéĚ■çŽĎǎŤřéĞŘ
```

17.7 17.6 éǎL'æŃl'ǎĞǎŤř

ǎ;şéǎŽèŁĞǎijŸéǎL'çŽĎèŁĆĆzæIJL'ǎd'ŽǎŸŋijŃéČčǎžŁǎžŎäŸ■éŽŘæIJzéǎL'æŃl'ǎŸǎǎŘř

CHAPTER 18

Pod 18.1

Pod 18.1

18.1 Pod 18.1

- Pod 18.1

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-schedule-demo
  namespace: default
  labels:
    app: myapp
spec:
  containers:
  - name: myapp
    image: ikubernetes/myapp:v1
  nodeSelector:
    gpu: ok
```

- Pod 18.1

```
kubectl label nodes node2 gpu=ok --overwrite
```

- `kubectl label nodes node2 gpu=ok --overwrite`

```
kubectl get pods -o wide
```

18.2 18.2 `kubectl explain pods.spec.affinity.nodeAffinity`

`kubectl explain pods.spec.affinity.nodeAffinity`

- `kubectl explain pods.spec.affinity.nodeAffinity`

```
nodeAffinity                                <Object>                                #_
→ POD affinity node                       <[]Object>                                #_
  preferredDuringSchedulingIgnoredDuringExecution
  → preferredDuringSchedulingIgnoredDuringExecution
  preference                                <Object>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  matchExpressions                          <[]Object>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  key                                       <string>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  operator                                  <string>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  values                                    <[]string>                               #_
  → preferredDuringSchedulingIgnoredDuringExecution
  matchFields                              <[]Object>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  key                                       <string>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  operator                                  <string>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  values                                    <[]string>                               #_
  → preferredDuringSchedulingIgnoredDuringExecution
  weight                                    <integer>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  requiredDuringSchedulingIgnoredDuringExecution <Object>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  nodeSelectorTerms                         <[]Object>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  matchExpressions                          <[]Object>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  key                                       <string>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  operator                                  <string>                                #_
  → preferredDuringSchedulingIgnoredDuringExecution
  values                                    <[]string>                               #_
  → preferredDuringSchedulingIgnoredDuringExecution
```

(continues on next page)

(continued from previous page)

```

    matchFields      <[]Object>                                #_
    ↪æSěæL'¿ā■Ůæōt
        key          <string>                                #_
    ↪æăĜç■¿
        operator     <string>                                #_
    ↪æŞ■ă;IJiijŽæřřè¿Č
        values       <[]string>                              #_
    ↪ăĀij

```

• çd'žä¿NéĚ■ç¿õ

```

apiVersion: v1
kind: Pod
metadata:
  name: pod-nodeaffinity1-demo
  namespace: default
  labels:
    app: myapp
spec:
  containers:
  - name: myapp
    image: ikubernetes/myapp:v1
  affinity:
    nodeAffinity:
      requiredDuringSchedulingIgnoredDuringExecution:      #_
    ↪çañăžšăŠNæĀĜèçAæśĆiijNăÿ■æzæűşăřś Pending
      nodeSelectorTerms:
      - matchExpressions:
        - key: zone
          operator: In
          values:
          - foo
          - bar
---
apiVersion: v1
kind: Pod
metadata:
  name: pod-nodeaffinity2-demo
  namespace: default
  labels:
    app: myapp
spec:
  containers:
  - name: myapp
    image: ikubernetes/myapp:v1
  affinity:
    nodeAffinity:
      preferredDuringSchedulingIgnoredDuringExecution:      #_
    ↪è¿řăžšăŠNæĀĜèçAæśĆiijNăÿ■æzæűşăžśSăŘřăžčăřzăČŚ

```

(continues on next page)

(continued from previous page)

	labelSelector	<Object>	#_
↪æǻĜç■¿éĂĹ'æŃl'ăŽĺăržèśaqđĹÛèqÍ			
	matchExpressions	<[]Object>	#_
↪æǻĜç■¿éĂĹ'æŃl'ăŽĺăržèśaqïijÑěĂĹ' POD æǻĜç■¿			
	key	<string>	#_
↪æǻĜç■¿			
	operator	<string>	#_
↪æŞ■ä;IJii jžærřè¿Č			
	values	<[]string>	#_
↪ăÄij			
	matchLabels	<map[string]string>	#_
↪éŽEǻŘĹæǻĜç■¿éĂĹ'æŃl'ăŽĺ			
	namespaces	<[]string>	#_
↪ăŘ■çğřçł'žéŮt'čŽďĎĹÛèqÍ			
	topologyKey	<string>	#_
↪ăžšăŠňăĽd'æŮ■æİaqăžű			
	weight	<integer>	#_
↪æİćéĞ■ 1 - 100			
	requiredDuringSchedulingIgnoredDuringExecution	<[]Object>	#_
↪çañæǻĜăžšăŠňæǻĜïijNăy■æžaëüşăĹŻ Pending			
	labelSelector	<Object>	#_
↪æǻĜç■¿éĂĹ'æŃl'ăŽĺăržèśaqđĹÛèqÍ			
	matchExpressions	<[]Object>	#_
↪æǻĜç■¿éĂĹ'æŃl'ăŽĺăržèśaqïijÑěĂĹ' POD æǻĜç■¿			
	key	<string>	#_
↪æǻĜç■¿			
	operator	<string>	#_
↪æŞ■ä;IJii jžærřè¿Č			
	values	<[]string>	#_
↪ăÄij			
	matchLabels	<map[string]string>	#_
↪éŽEǻŘĹæǻĜç■¿éĂĹ'æŃl'ăŽĺ			
	namespaces	<[]string>	#_
↪ăŘ■çğřçł'žéŮt'čŽďĎĹÛèqÍ			
	topologyKey	<string>	#_
↪ăžšăŠňăĽd'æŮ■æİaqăžű			

- çd'žä;NéĚ■ç;ö

```
apiVersion: v1
kind: Pod
metadata:
  name: pod1
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
```

(continues on next page)

- | NAME | READY | STATUS | RESTARTS | AGE | IP | NODE | |
|------------|-------|---------|-----------|-------|------------|-------|--|
| ↪NOMINATED | | NODE | READINESS | GATES | | | |
| pod1 | 1/1 | Running | 0 | 3m33s | 10.244.2.4 | node3 | |
| ↪<none> | | <none> | | | | | |
| pod2 | 1/1 | Running | 0 | 3m33s | 10.244.2.5 | node3 | |
| ↪<none> | | <none> | | | | | |

```

podAntiAffinity                                <Object>                                #_
→ POD árzãĚüäzŮ POD čŽĎãĚ■äžšãŠÑæĚ
  preferredDuringSchedulingIgnoredDuringExecution <[]Object> #_
→ è;řæĚğãĚ■äžšãŠÑæĚğiiĴÑãř;éĚŘæzæèŮşäžšãŠÑæĚ
    podAffinityTerm                            <Object>                                #_
→ ãĚ■äžšãŠÑčŽĎ POD árzèsa
      labelSelector                            <Object>                                #_
→ æăĚç■čéĚĹ'æŇĹ'ăŽĹãřžèsaăĹŮèaĹ
        matchExpressions                      <[]Object>                                #_
→ æăĚç■čéĚĹ'æŇĹ'ăŽĹãřžèsaġiiĴÑéĚĹ' POD æăĚç■č
          key                                <string>                                #_
→ æăĚç■č
          operator                            <string>                                #_
→ æş■ă;IJiiĴžæřĤèčĚ
          values                              <[]string>                                #_
→ ăĚij
            matchLabels                       <map[string]string>                        #_
→ éŽĚãĚĹæăĚç■čéĚĹ'æŇĹ'ăŽĹ
              namespaces                     <[]string>                                #_
→ ãĚ■çğřçĹ'žéŮt'čŽĎãĹŮèaĹ
                topologyKey                  <string>                                #_
→ äžšãŠÑãĹd'æŮ■æĹaäžŮ
                weight                        <integer>                                #_
→ æĹĚéĚ■ 1 - 100
            requiredDuringSchedulingIgnoredDuringExecution <[]Object> #_
→ çañæĚğãĚ■äžšãŠÑæĚğiiĴÑäÿ■æzæèŮşãĹŽ Pending
              labelSelector                  <Object>                                #_
→ æăĚç■čéĚĹ'æŇĹ'ăŽĹãřžèsaăĹŮèaĹ
                matchExpressions              <[]Object>                                #_
→ æăĚç■čéĚĹ'æŇĹ'ăŽĹãřžèsaġiiĴÑéĚĹ' POD æăĚç■č
                  key                        <string>                                #_
→ æăĚç■č
                  operator                    <string>                                #_
→ æş■ă;IJiiĴžæřĤèčĚ
                  values                      <[]string>                                #_
→ ăĚij
                    matchLabels               <map[string]string>                        #_
→ éŽĚãĚĹæăĚç■čéĚĹ'æŇĹ'ăŽĹ
                      namespaces              <[]string>                                #_
→ ãĚ■çğřçĹ'žéŮt'čŽĎãĹŮèaĹ
                        topologyKey           <string>                                #_
→ äžšãŠÑãĹd'æŮ■æĹaäžŮ

```

- *éĚ■ç;öæÿĚã■Ĥ*

```

apiVersion: v1
kind: Pod
metadata:
  name: pod3
  namespace: default

```

(continues on next page)

(continued from previous page)

```

labels:
  app: myapp
  tier: frontend
spec:
  containers:
  - name: myapp
    image: ikubernetes/myapp:v1

---
apiVersion: v1
kind: Pod
metadata:
  name: pod4
  namespace: default
  labels:
    app: db
    tier: db
spec:
  containers:
  - name: busybox
    image: busybox:latest
    imagePullPolicy: IfNotPresent
    command:
    - "sh"
    - "-c"
    - "sleep 3600"
  affinity:
    podAntiAffinity:
      requiredDuringSchedulingIgnoredDuringExecution: #
→ çañăžšăŠÑăĀğēēAæśĆiijŇăÿ■æzæűşçŽĎ Pending
    - labelSelector:
        matchExpressions:
        - key: app
          operator: In
          values:
          - myapp
        topologyKey: kubernetes.io/hostname #
→ âĤ■ăžšăŠÑăĀğçŽĎä; İæ■ōăÿžăĤŇăÿĀăÿlăÿzæIJžăĤ■

```

18.5 18.5 node æśaçĆż

æśaçĆżăĤćŤlăIJl node äÿŁçŽĎéŤôăĀijăśđæĀğiiJLn-
odes.spec.taintsīijL'īijŇăőČçŽĎä;IJçŤlæŸræŇŠçziäÿ■ēČ;ăőžăſ■ēſŽăžZæśaçĆżçŽĎ POD
ēſŖëąŇçŽĎīijŇăŽăæ■d'ēIJĀēēAăIJl POD äÿLăőŽăZl'ăőžăſ■ăžēīijŁpods.spec.tolerationsīijL'īijŇăőČăžşæŸſ
POD âĤŕăžăőžăſ■çŽĎæśaçĆżăLŮëălăĀĆ

äÿĀăÿl POD ēČ;äÿ■ēČ;ēſŖëąŇăIJlăÿĀăÿlēŁĆçĆžăÿLīijŇăŕşæŸſ pods.spec.tolerations

āLŪēāīāy■æŸřāŘēāÑĚāÑñāžĚ nodes.spec.taints āy■çŽDæŦřæ■ōāĂĆ

- node æśāçĆžæyĚā■ŦæāijāijŦīijÑèřēğAīijŽkubectl explain node.spec.taints

```
taints          <[]Object>          # æśāçĆžāřžèśāāLŪēāĪ
effect          <string>             # ā;Ŧ POD
→āy■ēČ;āōžāē■ēŁžāyŁæśāçĆžçŽDæŦŪāāŽīijÑèřēĀēĠĠāRŪçŽDēāÑāyžīijÑāžŚāřśæŸřæŦŦēāy■ā
→POD
  NoSchedule    #
→ā;śāŦ■ēřČāžēēŁĠçĹŦīijÑā;ĒæŸřāŪŦçžŦèřČāžēāōÑæĹŦ POD æŦāā;śāŦ■
  PreferNoSchedule #
→ā;śāŦ■ēřČāžēēŁĠçĹŦīijÑāřĪēŦŦēl' séĀŦèřČāžēāŪŦçžŦāōÑæĹŦçŽDä;Ēāy■āōžāē■æŦřæśāçĆžç
→POD
  NoExecute     #
→æŦřāčđçŽDæśāçĆžīijÑā;śāŦ■æŦřçŽDēřČāžēēŁĠçĹŦīijÑāyŦāijžāŁžēl' séĀŦèřČāžēāŪŦçžŦāō
→POD
  key           <string>             # ēŦŦ
  timeAdded     <string>             #
  value         <string>             # āĀij
```

- çŽŽ node æL'ŦāyŁæśāçĆžīijÑēŦŦāyž node-type āĀijāyž productionīijÑæśāçĆžāĹā;Ī

```
kubectl taint node node2 node-type=production:NoSchedule
```

- āĹāēŽd' node āyŁçŽDāyĀāyŁæśāçĆž

```
kubectl taint node node2 node-type-
```

- æŦŦēřŦæyĚā■Ŧ

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp-deploy
  namespace: default
spec:
  replicas: 4
  selector:
    matchLabels:
      app: myapp
      release: canary
  template:
    metadata:
      labels:
        app: myapp
        release: canary
    spec:
      containers:
        - name: myapp
          image: ikubernetes/myapp:v2
          ports:
```

(continues on next page)

(continued from previous page)

```
- name: http
  containerPort: 80
```

- `kubectl get pods -o wide` displays the PODs on node2

NAME	IP	NODE	NOMINATED	READY	STATUS	RESTARTS	AGE	
				1/1	Running	0	9s	
myapp-deploy-675558bfc5-4x5cf	10.244.2.13	node3	<none>	1/1	Running	0	9s	
myapp-deploy-675558bfc5-58f2s	10.244.2.10	node3	<none>	1/1	Running	0	9s	
myapp-deploy-675558bfc5-gz4kv	10.244.2.12	node3	<none>	1/1	Running	0	9s	
myapp-deploy-675558bfc5-hlxdd	10.244.2.11	node3	<none>	1/1	Running	0	9s	

- `kubectl taint node node3` adds the taint to node3

```
kubectl taint node node3 node-type=dev:NoExecute
```

- `kubectl get pods -o wide` displays the PODs on node3

NAME	IP	NODE	NOMINATED	READY	STATUS	RESTARTS	AGE	
				0/1	Pending	0	10s	
myapp-deploy-675558bfc5-22wpj	<none>	<none>	<none>	0/1	Pending	0	14s	
myapp-deploy-675558bfc5-lctv5	<none>	<none>	<none>	0/1	Pending	0	15s	
myapp-deploy-675558bfc5-m5qdh	<none>	<none>	<none>	0/1	Pending	0	14s	
myapp-deploy-675558bfc5-z8c4q	<none>	<none>	<none>	0/1	Pending	0		

18.6 18.6 POD Tolerations

- `kubectl explain pods.spec.tolerations`

```
tolerations          <[]Object>          # array of tolerations
effect               <string>          # effect of the taint
-äýŁçŽĎæšÇžél' séĀŘčŮčTěijNäýžçl' žèáĺçd' žăőžăŋăžžă;Tél' séĀŘčŮčTě
  NoSchedule         # effect of the taint on the pod
-NoSchedule
  PreferNoSchedule   # effect of the taint on the pod
-PreferNoSchedule
```

(continues on next page)

(continued from previous page)

```

NoExecute # èĈ;ăôžăĤ node æśaçĆżĹĐ
→NoExecute
key <string> # æśaçĆżĹĐěŤó
operator <string> # Exists
→æśaçĆżăŸăIJlăÿçôăăžĂăžĹăĂijĭijŃEqual æśaçĆżĹĐăĂijăĤĚăžçL'ăĂij
tolerationSeconds <integer> #
→ăôžăĤæŮŮéŮt'ĭijŃăşăçĈăđIJèćnéł'séĂŕĭijŃăŕŕăžěçL'ăđ'žăžĚăĤèĥŕĭijŃézŸèôd'
→0 çġŖĭijŃNoExecute ä;ĤçŤĹ
value <string> # æśaçĆżĹĐăĂij

```

- çĹŽ node2 ĂĂnode3 ĂĹĚĂĹăĤL'ŞæśaçĆż

```

kubect1 taint node node2 node-type=production:NoSchedule
kubect1 taint node node3 node-type=dev:NoExecute

```

- ăôžăĤL POD æŸĚăŤæŮĠăžŮŕĭijŃăôžăĤ node äŸĹăŸăIJĹ node-type ăĂijăŸ dev çĹĐæśaçĆżăĂăŖŮŮèćnéł'séĂŕăĂĈ

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp-deploy
  namespace: default
spec:
  replicas: 4
  selector:
    matchLabels:
      app: myapp
      release: canary
  template:
    metadata:
      labels:
        app: myapp
        release: canary
    spec:
      containers:
        - name: myapp
          image: ikubernetes/myapp:v2
          ports:
            - name: http
              containerPort: 80
      tolerations:
        - key: node-type
          operator: Equal
          value: dev
          effect: NoExecute

```

- æŖçĹJŃçŹşăđIJĭijŃkubect1 get pods -o wideŕĭijŃĚĤŕĚăŃăIJĹĚĠăŮşăôžăĤçĹĐæśaçĆżĹĐĹĈĈăŸăĹ

NAME	READY	STATUS	RESTARTS	AGE	
→IP	NOMINATED NODE	READINESS	GATES		
myapp-deploy-97578cf74-5v2r6	1/1	Running	0	6m22s	
→10.244.2.16	node3	<none>			
myapp-deploy-97578cf74-gbfj7	1/1	Running	0	6m22s	
→10.244.2.14	node3	<none>			
myapp-deploy-97578cf74-l4lbv	1/1	Running	0	6m22s	
→10.244.2.15	node3	<none>			
myapp-deploy-97578cf74-zvn8f	1/1	Running	0	6m20s	
→10.244.2.17	node3	<none>			

- äyžèŁĆćĆzácđāŁæŮřčŽĎæśaçĆziiŇèőç;óél'sçež POD

```
kubectl taint node node3 disk=hdd:NoExecute --overwrite
```

- æšěçIJŇčžŠæđIJiijŇkubectl get pods -o wideiijŇPOD
äy■èČ;ăőžâŁæŮřčŽĎæśaçĆziiŇčžŠæđIJèćnél'séĂŘ

NAME	READY	STATUS	RESTARTS	AGE	
→IP	NOMINATED NODE	READINESS	GATES		
myapp-deploy-97578cf74-84bfz	0/1	Pending	0	6s	
→<none>	<none>	<none>			
myapp-deploy-97578cf74-fxk2d	0/1	Pending	0	5s	
→<none>	<none>	<none>			
myapp-deploy-97578cf74-jp99j	0/1	Pending	0	6s	
→<none>	<none>	<none>			
myapp-deploy-97578cf74-vdkbx	0/1	Pending	0	6s	
→<none>	<none>	<none>			

CHAPTER 19

å■Aäzi åóžåŽìèťDæžŘéŽŘåŁú

èťúăġNăĀij requests æIJĀä;ŎăĤléŽIJ

çžŁçžŜăĀij limits çăñéŽŘåŁú

- CPU

```
1 éćŮ CPU = 1000 millicores
0.5 éćŮ CPU = 500 m
```

- åĖĚå■Ÿ

```
EiăĀĀPiăĀĀTiăĀĀGiăĀĀMiăĀĀKi
```

19.1 19.1 èťDæžŘéŽŘåŁú

- æŸĚå■ŤæăijăijŘiijNèřèġĀiijŽkubectl explain pods.spec.containers.resources

resources	<Object>	# èťDæžŘéŽŘåŁú
limits	<map[string]string>	# èťDæžŘæIJĀénŸéŽŘåŁú
cpu	<string>	# å■Ťă;■ m
memory	<string>	# å■Ťă;■ GiăĀĀMi
requests	<map[string]string>	# èťDæžŘæIJĀä;ŎèĕAæśĆ
cpu	<string>	# å■Ťă;■ m
memory	<string>	# å■Ťă;■ GiăĀĀMi

- æŸĚå■Ťçď'žăĤNriijNnode èŁĆÇĆžçŽĎ CPU äŸž 12 æăŸăĤČrijNcpu limits èŎĥç;ŏăŸž 1000m äžşăřşæŸřăĚĀèŏŸ


```

apiVersion: v1
kind: Pod
metadata:
  name: pod-resources-demo
  namespace: default
  labels:
    app: myapp
    tier: frontend
spec:
  containers:
  - name: nginx
    image: ikubernetes/stress-ng
    command:
      - "/usr/bin/stress-ng"
      # - "-m 1" # äžěā■ŦçžŁçłŃăŎŇætŃăĚĚă■Ÿ
      - "-c 1" # äžěā■ŦçžŁçłŃăŎŇætŃCPU
      - "--metrics-brief"
    ports:
      - name: http
        containerPort: 80
      - name: https
        containerPort: 443
    resources:
      requests:
        cpu: 1000m #
        ↪äŏČăĚşăŏžăIJléčĎéĀL' éŸŭæŏŧæŭŸæsřăŞłăžŽăÿzæIJž
        memory: 512Mi
      limits:
        cpu: 1000m # èāÍçd'žéŽŘăĹŭăŏžăžłä;ŁçŦÍ node
        ↪èŁČçČžçŽĎăÿĂéćŮ
        ↪CPUiijŇăŮăèŏžăd'žăřSèŁžçłŃiijŇăŏČăžŇæIJĀăd'žăŘłèč;ă■ăçŦÍ node
        ↪èŁČçČžçŽĎăŘ CPU
        memory: 512Mi

```

• æşçIJŇçžŞæđIJ

```

Mem: 855392K used, 139916K free, 10188K shrd, 796K buff, 350368K
↪cached
CPU0:  0% usr  0% sys  0% nic  99% idle  0% io  0% irq  0%
↪sirq
CPU1: 100% usr  0% sys  0% nic  0% idle  0% io  0% irq  0%
↪sirq # ä■ăæžăăžĚăÿĂéćŮ CPU
CPU2:  0% usr  0% sys  0% nic  99% idle  0% io  0% irq  0%
↪sirq
CPU3:  0% usr  0% sys  0% nic  99% idle  0% io  0% irq  0%
↪sirq
CPU4:  0% usr  0% sys  0% nic  99% idle  0% io  0% irq  0%
↪sirq
CPU5:  0% usr  0% sys  0% nic  99% idle  0% io  0% irq  0%
↪sirq

```

(continues on next page)

(continued from previous page)

```

CPU6:   0% usr    0% sys    0% nic   99% idle    0% io    0% irq    0%
↪sirq
CPU7:   0% usr    0% sys    0% nic   99% idle    0% io    0% irq    0%
↪sirq
CPU8:   0% usr    0% sys    0% nic   99% idle    0% io    0% irq    0%
↪sirq
CPU9:   0% usr    0% sys    0% nic   99% idle    0% io    0% irq    0%
↪sirq
CPU10:  0% usr    0% sys    0% nic   99% idle    0% io    0% irq    0%
↪sirq
CPU11:  0% usr    0% sys    0% nic   99% idle    0% io    0% irq    0%
↪sirq
Load average: 0.84 0.50 0.40 3/485 11
  PID  PPID  USER      STAT  VSZ  %VSZ  CPU  %CPU  COMMAND
    6    1  root       R      6888   1%    1    8% {stress-ng-cpu} /usr/
↪bin/stress-ng -c 1 --metrics-brief
    1    0  root       S      6244   1%   10    0% /usr/bin/stress-ng -c
↪1 --metrics-brief
    7    0  root       R      1504   0%   11    0% top

```

19.2 19.2 qos è'léĜŔçõaçŔĖ

- GuranteedW

```

æŕŔÄÿłăôžăŽĺăŔŇæŮűèőç;őăžĚ CPU ăŠŇăĚăĚă■ŸçŽĎ requests ăŠŇ
↪limitsiijŇèĂŇăÿŦ
    cpu.limits = cpu.requests
    memory.limits = memory.requests
éĆčăžĺăôČăŕĚăĭjŸăĚĺèćńèŕČăžę

```

- Burstable

```

èĜşăŕşŖæIJL' äÿĂäÿłăôžăŽĺèőç;ő CPU æĹŮăĚăĚă■ŸèťĎăžŔçŽĎ requests ăśđæĂĝ
éĆčăžĺăôČăŕĚăĚűæIJL' äÿ■ç■L' äijŸăĚĺçžĝ

```

- BestEffort

```

æşăæIJL' äžžă;ŦäÿĂäÿłăôžăŽĺèőç;őăžĚ requests æĹŮ limits ăśđæĂĝ
éĆčăžĺăôČăŕĚăŕłæIJL' æIJĂä;ŎäijŸăĚĺçžĝiijŇă;şèťĎăžŔăÿ■ăđ'şçŦĺçŽĎăűűăĂŽiijŇèĚăÿłă
↪Burstable ăŠŇ Guranteed

```

- oom ç■ŮçŦě

```

æIJĂăĚĺăĭĂă■ă■ăçŦĺéĜŔăŠŇéIJĂæsĆéĜŔçŽĎærŦă;Ňăđ'ğçŽĎ

```

CHAPTER 20

āžŅā■A HeapSterçŽŠæŌģīīJLāžšāijČäy■īīJL'

```
kubectl top æÝř k8s āEĚç;ōæšēcIJŇ POD çŽŠæŌģāfæAřçŽDāŠ;äzd'īīJŅāŌČæÝřäzŌ
HeapSter äy■āRŪā;ŪæTřæ■ōīīJŅēĀŇ HeapSter æÝřèŁRèqŅāIJĪ K8S
éŽEç;đ'çžğāŁhçŽDçŽŠæŌģè;řäzŭāĀĆ
```

```
kubectl æIJL'äyÄäyĹāEĚç;ōæRŠäzŭīīJŅāRŋ cAdvisor āŌČçTĹæIèæTŭéZE node
èŁČçČzāŠŅèŁČçČzäyŁçŽD POD äyŁçŽDètDæžRä;ŁçTĹéGRīīJŅHeapSter āRřäzèæTŭéZE
cAdvisor āIJĹæřRäyĹèŁČçČzäyŁéGĚZEçŽDæTřæ■ōīīJŅèŁZäžZæTřæ■ŌæČšèeAæŅAçz■ā■YāČīīJŅéČčāzĹ
influx DB çŽDæTřæ■ŌāžŠäy■īīJŅçDŭāRŌāRřäzèä;ŁçTĹ Grafana éĚç;ŏ influx DB
äyžæTřæ■ŌæžRīīJŅçDŭāRŌāsTçd'žāĀĆ
```

- heapster āŌYæŪzéažçŽŌ

```
https://github.com/kubernetes-retired/heapster
```

20.1 20.1 āŌL'èčĚ influx DB

```
influx DB æÝřäyÄäyĹæŪŭāžRæTřæ■ŌāžŠīīJŅāŌČéIJĀèeAäyÄäyĹæŅAžžĒā■YāČĹæIèæĹā■YæTřæ■ōīīJ
emptyDir āŁŌæTžäyžāĒŭæIJL'æŅAžžĒā■YāČĹèČ;āŁŽçŽDā■YāČĹā■ŭāĀĆ
```

- heapster/deploy/kube-config/influxdb/influxdb.yaml

```
apiVersion: apps/v1 # æ■d'ād'DäŁŌæTžäyž apps/
→v1īīJŅçTšäžŌāŁŌæTžäžEæ■d'ād'DæL'ÄäžèèŁYéIJĀèeAäŁŌæTžæāĢç■çéĀL'æŅĹ'āžĹ
kind: Deployment
metadata:
  name: monitoring-influxdb
  namespace: kube-system
spec:
```

(continues on next page)

(continued from previous page)

```

replicas: 1
selector: # æůžăĽăăĞç■¿
  matchLabels: # æůžăĽăăĞç■¿
    task: monitoring # æůžăĽăăĞç■¿
    k8s-app: influxdb # æůžăĽăăĞç■¿
template:
  metadata:
    labels:
      task: monitoring
      k8s-app: influxdb
  spec:
    containers:
      - name: influxdb
        image: k8s.gcr.io/heapster-influxdb-amd64:v1.5.2
        volumeMounts:
          - mountPath: /data
            name: influxdb-storage
        volumes:
          - name: influxdb-storage
            emptyDir: {} #
→æ■d'ăd'ĎăžĤëřăăŁôæĤžăÿžăĚűæI JL'æŇăžĚă■ŸăĆíèĈ;ăĽŽçŽĎ
---
apiVersion: v1
kind: Service
metadata:
  labels:
    task: monitoring
    # For use as a Cluster add-on (https://github.com/kubernetes/
→kubernetes/tree/master/cluster/addons)
    # If you are NOT using this as an addon, you should comment out
→this line.
    kubernetes.io/cluster-service: 'true'
    kubernetes.io/name: monitoring-influxdb
  name: monitoring-influxdb
  namespace: kube-system
spec:
  ports:
    - port: 8086
      targetPort: 8086
  selector:
    k8s-app: influxdb

```

- æšĉIJŇ service äŠŇ pod

```

$ kubectl get svc -n kube-system
NAME                                TYPE                CLUSTER-IP          EXTERNAL-IP
→PORT(S)                          AGE
kube-dns                          ClusterIP           10.96.0.10          <none>
→53/UDP, 53/TCP, 9153/TCP          9d

```

(continues on next page)

(continued from previous page)

kubernetes-dashboard	NodePort	10.109.5.194	<none>	└
→ 443:30894/TCP	3d14h			
monitoring-influxdb	ClusterIP	10.104.173.236	<none>	└
→ 8086/TCP	18s			
\$ kubectl get pod -n kube-system				
NAME		READY	STATUS	RESTARTS └
→ AGE				
.....				
monitoring-influxdb-866db5f944-d7rkd		1/1	Running	0 └
→ 69s				

20.2 20.2 ǎŎL'ěčĚ HeapSter

- éĕŮǎĚĹǎŎL'ěčĚ rbac ċŽĎċŤĹǎĹüijŇheapster/deploy/kube-config/rbac/heapster-rbac.yaml

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: heapster
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: system:heapster
subjects:
- kind: ServiceAccount
  name: heapster
  namespace: kube-system
```

- ǎŎL'ěčĚ heapsterüijŇheapster/deploy/kube-config/influxdb/heapster.yaml

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: heapster
  namespace: kube-system
---
apiVersion: apps/v1 # ǎŎL'ěčĚ d'ǎd'ĎċŽĎǎŷž apps/
→ vliijŇǎŎL'ěčĚ d'ǎd'ĎǎžǎǎŎŎēŷŷēIJǎēēǎŷǎǎ selector
kind: Deployment
metadata:
  name: heapster
  namespace: kube-system
spec:
  replicas: 1
  selector: # ǎžǎǎŷžǎŎL'ěčĚ api└
→ ċŤ'ĹǎIJŇǎŤ'ǎžǎēēIJǎēēǎŷǎǎ
```

(continues on next page)

(continued from previous page)

```

matchLabels:                                # æŭžăĹă
  task: monitoring                           # æŭžăĹă
  k8s-app: heapster                         # æŭžăĹă
template:
  metadata:
    labels:
      task: monitoring
      k8s-app: heapster
  spec:
    serviceAccountName: heapster
    containers:
      - name: heapster
        image: k8s.gcr.io/heapster-amd64:v1.5.4
        imagePullPolicy: IfNotPresent
        command:
          - /heapster
          - --source=kubernetes:https://kubernetes.default
          - --sink=influxdb:http://monitoring-influxdb.kube-system.
→svc:8086
---
apiVersion: v1
kind: Service
metadata:
  labels:
    task: monitoring
    kubernetes.io/cluster-service: 'true'
    kubernetes.io/name: Heapster
  name: heapster
  namespace: kube-system
spec:
  ports:
    - port: 80
      targetPort: 8082
  selector:
    k8s-app: heapster

```

- æšćĹŃčžŠæđĹĵĴkubectl get pods -n kube-system

NAME	READY	STATUS	RESTARTS	
→ AGE				
...				
heapster-5d4bf58946-6dmgf	1/1	Running	0	
→ 113s				
monitoring-influxdb-866db5f944-d7rkd	1/1	Running	0	
→ 23m				

(continued from previous page)

```

    hostPath:
      path: /etc/ssl/certs                                # æşĺæĎŘéĚ■ç;ő ssl èřAăžę
    - name: grafana-storage                               # æşĺæĎŘéĚ■ç;őæŇAăžĚă■ŸăĆĺ
      emptyDir: {}
---
apiVersion: v1
kind: Service
metadata:
  labels:
    kubernetes.io/cluster-service: 'true'
    kubernetes.io/name: monitoring-grafana
  name: monitoring-grafana
  namespace: kube-system
spec:
  type: NodePort                                          #
  →âĕĆæđIJéIJĀëĕAéĀŽèĚĜăd'ŮéĈĺèőĚéŮőăŘřăžěăL'şaijĀiijŇăŘřăžěăĚşéŮ■
  →heapster çŽĎæ■d'éaz
  ports:
    - port: 80
      targetPort: 3000
  selector:
    k8s-app: grafana

```


CHAPTER 21

21.1 21.1 æäyå£CæŃGæäGætAært'çž£

21.1 21.1 æäyå£CæŃGæäGætAært'çž£

çŤś kubeletāĀmetrics-server äžěāRŁçŤś apiserver æŘŘä;ŽçŽĐ api
çžĐæLŘiijŽäyžèeA CPUçŤŕeōaä;£çŤlçŤŌĜāĀāEĒā■ŸāōđæŮüä;£çŤlçŤŌĜāĀPOD
ètĐæžŘā■āçŤlçŤŌĜāRŁāōžāŽlçŽĐçčAçŽŸā■āçŤlçŤŌĜāĀC

- metrics-serveriijLæŮŕäyĀäzççŽĐètĐæžŘæŃGæäGèŮāRŮæŮžäijRiijL

āōCæŸŕäyĀäyĭ apiserver iijŃāōCäzĒäzĒçŤlāžŌæIJ■āLāžŌæäyā£CæŃGæäGæIJ■āLāçŽĐiijŃāōCäy■æ
k8s çžĐçžĐæLŘēČlāLEiijŃäzĒäzĒæŸŕæLŸçōāāIJl k8s äžŃäyL PODāĀC

k8s çžĐ apiserver āšŃ metrics-server çžĐ apiserver
āL■çŋŕāžŤŕēāLāyĀäyĭäzççŘEæIJ■āLāžŤlīijŃāōČāŕsæŸŕäyĀäyĭæAžāRŁāŽlīijŃæLŁæĪēēĜlād'Žäyĭäy■āRŁ
apiserver èAžāRŁæLŘäyĀäyĭāĀCāōČāŕsæŸŕ kube-aggregatoriijŃçžŘē£ĜāōČèAžāRŁāRŌçžĐ
api æLŠäzLāŕEēĀŽē£Ĝ /apis/metrics.k8s.io/v1/beta1 æĪēēŮāRŮāĀC

21.2 21.2çžŽSæŌğætAært'çž£

çŤlāžŌäzŌçšçžçžæŤŭéZEāRĐçĝ■æŃGæäGæŤŕæ■ōāžŭæŘŘä;ŽçžLçŋŕçŤlæLŭāĀā■ŸāČlçšçžçžäžěāR
HPAiijŃāōČāŃĒāŕŋæäyā£CæŃGæäGāšŃēĪđæäyā£CæŃGæäGīiijŃēĪđæäyā£CæŃGæäGäy■ēČ;ēčŋ
k8s æLĀçŘEēĝčīiijŃk8s-prometheus-adapter āŕsæŸŕē;ŋæ■cäyž k8s
æLĀçŘEēĝcæāijāijRçžĐäyĀäyĭæRŠäzŭ

- prometheus

CNCŦäyŃçžĐçŋŋāžŃād'ĝéāççžŽōiijŃæŤŭéZEāRĐçĝ■çžŤ'āžççžĐæŃGæäGīiijŃ

āōCæŤŭéZEçžĐä£æAŕiijŃæĪēāEšāōŽæŸŕāRçē£ZEāŃ HPAiijLēĜlāLāijyçijl'iijL
çžĐäyĀäyĭæāĜāĜE

- ăǫŸæŮzäzŞăžŞiijÑè£ŽéĞŇæĹŚä;£çŤlçñňäyĂäyŧ

- `tree/master/deploy/iijNäŁœTž metrics-server-deployment.yaml æŮĜäzŮ`

```
$ kubectl apply -f ./
```

- æſëçIJN POD åŠŇ Service çŽDåŘráŁíæČĚåEt

```
$ kubectl get pods -n kube-system
$ kubectl get svc -n kube-system
```

- æſſecIJN API äy■æŸráŘeå■ŸaIJliijŇmetrics.k8s.io/v1beta1

```
$ kubectl api-versions
```

- éĀžēfĜætĠNērTæŌěāRčèŌuāRŮçŽSæŌğæTṛæ■ōiijŊkubectl proxy -port
8080iijŊkubectl top äžšāRräžěæ■čāvyä;ŁçŤlāžE

```
$ curl http://127.0.0.1:8080/apis/metrics.k8s.io/v1beta1
$ kubectl top nodes
```

21.4 21.4 aóL'èčĚ prometheus

- aũëä;IJãŎšçŘĚ

```
- prometheus éĀžèfĜ pull metrils æŇĜäzd' äžŎærĚäÿł Jobs/exporters_
→æŇL' äŔŮæŦŕæ■Ŏ
- äĚŮäzŮčŽĎ short-lived jobs äžšäŔŕäzèéĀžèfĜäŔš pushgateway_
→äÿzäŁłäŔšéĀæŦŕæ■ŎiiijŇčŦš prometheus ècňäŁłæŎëäŦŮ
- prometheus_
→èĜłèžnáŏđčŎŕäžĚäÿÄäÿłæŮúéŮŦ' äžŔäłŮæŦŕæ■ŎäžšiiijŇäijžäŕĚä;ŮäłŕčžĎæŦŕæ■ŎäŦŸäĆíäłŕä
- äIJł k8s éIJÄèĚÄä;fçŦł service discovery_
→æłëäŔšçŎŕäIJ■äŁääŔŮä;ŮéIJÄèĚÄçŽšæŎĝçŽĎçŽŏäĝ
- äŔŕäžèä;fçŦł apiclientäĀAwebuiäĀAGrafanaäĀAæłëäŕĚ prometheus_
→äÿ■çŽĎæŦŕæ■ŎäšŦçđ' žäĜžæłë
- ä;šéIJÄèĚÄæŁèè■ĚçŽĎäŮüäĀžèfŸäijžæŎłéĀÄçžŽ alertmanager_
→èfžäÿłçžĎäžúçŦšèfžäÿłçžĎäžüæłëäŔšéĀæŁèè■Ě
```

- éČłç;šæŮĜäzŮ

```
https://github.com/kubernetes/kubernetes/tree/master/cluster/addons/
→prometheus
https://github.com/kubernetes/k8s-prom
```

21.5 21.5 HPAäŠ;äzd'èqŇæŮžäijŔ

- äŁŽäžž POD äšŇ service

```
kubectl run myapp --image=ikubernetes/myapp:v1 --replicas=1 --
→requests='cpu=50m',memory='256Mi' --limits='cpu=50m,memory=256Mi'_
→--labels='app=myapp' --expose --port=80
```

- äŁŽäžž HPA æŎĝäŁüäŽł

```
kubectl autoscale deployment myapp --min=1 --max=8 --cpu-percent=60
```

- æšĚçIJŇ HPA æŎĝäŁüäŽłiiijŇkubectl get hpa

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS
→AGE					
myapp	Deployment/myapp	0%/60%	1	8	1
→17s					

- äijÄäĝŇäŎŇäŁžæŦŇŕŦ

```
ab -c 100 -n 5000000 http://172.16.100.102:32749/index.html
```

- æŦŇŕŦçžšæđIJiiijŇèĜłäŁłæŁ'äŏžçŦšæŦł

\$ kubectl get hpa -w						
NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	
↪ AGE						
myapp	Deployment/myapp	0%/60%	1	8	1	
↪ 7m35s						
myapp	Deployment/myapp	34%/60%	1	8	1	
↪ 9m58s						
myapp	Deployment/myapp	102%/60%	1	8	1	
↪ 11m						
myapp	Deployment/myapp	102%/60%	1	8	2	
↪ 11m						
myapp	Deployment/myapp	96%/60%	1	8	2	
↪ 12m						
myapp	Deployment/myapp	96%/60%	1	8	4	
↪ 12m						
myapp	Deployment/myapp	31%/60%	1	8	4	
↪ 13m						
myapp	Deployment/myapp	26%/60%	1	8	4	
↪ 14m						
myapp	Deployment/myapp	0%/60%	1	8	4	
↪ 15m						
myapp	Deployment/myapp	0%/60%	1	8	4	
↪ 17m						
myapp	Deployment/myapp	0%/60%	1	8	3	
↪ 18m						
\$ kubectl get pods						
NAME	READY	STATUS	RESTARTS	AGE		
myapp-64bf6764c5-45qwj	0/1	Terminating	0	7m1s		
myapp-64bf6764c5-72crv	1/1	Running	0	20m		
myapp-64bf6764c5-gmz6c	1/1	Running	0	8m1s		

21.6 21.6 HPAæÿĚā■Ț

- æÿĚā■ȚāŃŽāZL'èrèègAijȚ kubectl explain hpa.spec

maxReplicas	<integer>	#
↪ èĜłāŁłäijȚijl'çŽĎ POD æȚřéĜŘäÿŁéŽŘ		
minReplicas	<integer>	#
↪ èĜłāŁłäijȚijl'çŽĎ POD æȚřéĜŘäÿŃéŽŘ		
scaleTargetRef	<Object>	#
↪ āĚŭāzŮçŽĎäijȚijl'æŃĜæāĜ		
apiVersion	<string>	# æŃĜæāĜçšzāđŃ
↪ çL'ŁæIJŋ		
kind	<string>	# æŃĜæāĜçšzāđŃ
name	<string>	# āŘrçŤÍæŃĜæāĜ
targetCPUUtilizationPercentage	<integer>	# æāžæ■ŃçŽŃŃæāĜ
↪ āžšāİĜ CPU āŁł'çŤÍçŎĜéŸŁāĀijèrĎäijřèĜłāŁłäijȚijl'		

- çd'žä;ÑæÿĚā■TrijŇăőČăőđçŎřăžĚărž myapp èŁZăył deployment æŎğăLŭăZlăyŇçŽĐ
POD èŁZăyŇăĚăLlăL'łăőž

```

apiVersion: autoscaling/v2beta1
kind: HorizontalPodAutoscaler
metadata:
  name: myapp-hpa-v2
spec:
  scaleTargetRef:
    apiVersion: apps/v1
    kind: Deployment
    name: myapp
  minReplicas: 1
  maxReplicas: 10
  metrics:
  - type: Resource
    resource:
      name: cpu
      targetAverageUtilization: 55
  - type: Resource
    resource:
      name: memory
      targetAverageValue: 50Mi

```

CHAPTER 22

äžŃă■AäžŃ K8SăŃĚçőaçŘĚăŽÍ

Helm	æYr	Deis	aijAaRŠçŽDäyÄäylçTlāžŎ	Kubernetes
azTçTlçŽDāNĖçoaçRĖaũeāEũijNāyžèeAçTlāiēçoaçRĖ			ChartsāĀCæIJLçĆžçszāijijāžŎ	
Ubuntu äy■çŽD APT æLŪ	CentOS äy■çŽD	YUMāĀC		

Helm Chart æYřçTlæleāŕAèĈĖ Kubernetes āŎșçTšăžTçTlćlNăžRçŽDăyĂçșzālŬ
YAML æŨĞazũāĀCārRăzeāIljā;ăēĆlç;șăžTçTlćZDæUūāĂZēGlăōZăzl'ăžTçTlćlNăžRçŽDăyĂăžZ
MetadataiijNăžeä;£ăžŌăžTçTlćlNăžRçŽDălEărSăĀĊ

Helm

ærzāžŌażTčTlāRŚayČēĀĒèĀÑelĀijNāRfrazēēĀŽēŁG
æLŠāNĒażTčTlāAAçōaçRĒażTčTlājIetŪāĒšçszāĀAçōaçRĒażTčTłcLŁæIJnāzūāRŚayČażTčTlālRè;ražūaz.

árzāžŌä;ƒçTlèĀĖèĀÑelĀiijNä;ƒçTl Helm āRŌäy■çTlèIJAèçAçijŪāEZāđ'■āIČçŽDāžTçTlélČlç;šæŪGā
Kubernetes äyLæššæL';āĀĀāōL'èçĒāĀā■GçžgāĀāŽđæžŽāĀā■yè;|āžTçTlçlNāžRāĀČ

22.1 22.1 aſżçąĂæęĆąŁt

- Helm

Helm æYřäYÄäyĹaŚ;āzd'èaÑäyNçŽĐāōcæŁuçnřaũeăĔũăĂĆäyžèeAçŦĹăŹŎ Kubernetes
 āžŦçŦĹćĹNāžŦŦ Chart çŽĐāĹZāžžăĂAaŁŦŠăŦĔăĂAaŦŠăyČăžěaŦŦĹăĹZăžzăŠŦŦçoaçŦŦæĹJŦăĹJŦăŦŦŦŦĹJćĹNçŽ
 Chart äžŠăžŠăĂĆ

- Tiller

Tiller	æYf	Helm	çŽDæIJ■āŁaçnfrijŃēČlçj;šāIJl	Kubernetes
éŽEçj;d'äy■āĀĆTiller		çŤlāžŌæŌēæŤŭ	Helm	çŽDērūæšĆrijŃāžūæāžæ■ō
Chart	çŤšæLŔ	Kubernetes	çŽDēČlçj;šæŬGāzūiijLHelm	çğřāyž Re-
leaseiijL'rijŃçDūāŔŌæŔŔāžd'çŽŽ		Kubernetes	āLŽāžžāžŤçŤlāĀĆTiller	ēŤYæŔŔäçJāžE
Release	çŽDā■GçžgāĀāāLāēZd'āĀāžDæžŽç■L'äyĀçšžāLŬāšēČjāĀĆ			

- Chart

Helm çŽĐèířázúãÑĚíijÑéĜĜçŤÍ TAR æāijāijRāĀĆşşzāijijžŮ APT çŽĐ DEB
ãÑĚãĹŮèĀĚ YUM çŽĐ RPM ãÑĚíijÑãĚúãÑĚãRñázĒäyĀçžĐãŏŽázĹ Kubernetes
èĹĐæžŘçŽyãĚşçŽĐ YAML æŮĜäzúãĀĆ

- Repoistory

Helm çŽĐèířázúázŞázŞíijÑRepository æIJñèt'íäyĹæŸřäyĀäyĹ
Web æIJ■ãĹãŽíijÑèřèæIJ■ãĹãŽíäĹíã■ŸäžĒäyĀçşşzãĹŮçŽĐ Chart
èířázúãÑĚäzèä;ŽçŤÍæĹüäyÑè;íijÑázúäyŤæŘŘä;ŽázĒäyĀäyĹèř Repository çŽĐ Chart
ãÑĚçŽĐäyĒã■ŤæŮĜäzúãzèä;ŽæşèèçíijÑHelm áRřázèãRñæŮŮçŏaçĒĒäđ'ŽäyĹäy■ãRñçŽĐ
RepositoryãĀĆ

- Release

ä;ĹçŤÍ helm install áŞ;äzđ'ãIJĹ Kubernetes éŽĒç;đ'äy■éĆíçşçŽĐ Chart çğřäyž Re-
leaseãĀĆChart äyŮ Release çŽĐãĚşçşşzşşzāijijžŮŌéĹcãRŞáržèşäy■çŽĐçşşzäyŮãŏđä;ŹçŽĐãĚşçşşzãĀĆ

22.2 22.2 Helm áüěä;IJãŮşçŘĒ

- Chart Install èĹĜćÍŹ

1. Helm äžŮãÑĜãŏŽçŽĐçŽŏã;ŤæĹŮèĀĚ TAR æŮĜäzúäy■èğçæđŘãĜž Chart
→çžşæđĐäĹæAřãĀĆ
2. Helm áRĒæÑĜãŏŽçŽĐ Chart çžşæđĐãŞŹ Values äĹæAřéĀžèĹĜ gRPC
→äijäéĀşçžŹ TillerãĀĆ
3. Tiller æžã■Ů Chart áŞŹ Values çŤŞæĹŘäyĀäyĹ ReleaseãĀĆ
4. Tiller áRĒ Release áRŞéĀAçžŹ Kubernetes çŤĹäžŮçŤŞæĹŘ ReleaseãĀĆ

- Chart Update èĹĜćÍŹ

1. Helm äžŮãÑĜãŏŽçŽĐçŽŏã;ŤæĹŮèĀĚ TAR æŮĜäzúäy■èğçæđŘãĜž Chart
→çžşæđĐäĹæAřãĀĆ
2. Helm áRĒéIJĀèèAæžt'æŮřçŽĐ Release çŽĐãR■çğřãĀAChart çžşæđĐãŞŹ
→Values äĹæAřäijäéĀşçžŹ TillerãĀĆ
3. Tiller çŤŞæĹŘ Release äžúæžt'æŮřæÑĜãŏŽãR■çğřçŽĐ Release çŽĐ
→HistoryãĀĆ
4. Tiller áRĒ Release áRŞéĀAçžŹ Kubernetes çŤĹäžŮæžt'æŮř ReleaseãĀĆ

- Chart Rollback èĹĜćÍŹ

1. Helm áRĒèèAäžđæžžçŽĐ Release çŽĐãR■çğřäijäéĀşçžŹ TillerãĀĆ
2. Tiller æžã■Ů Release çŽĐãR■çğřæşèæĹ'; HistoryãĀĆ
3. Tiller äžŮ History äy■èŮãRŮäyĹäyĀäyĹ ReleaseãĀĆ
4. Tiller áRĒäyĹäyĀäyĹ Release áRŞéĀAçžŹ Kubernetes
→çŤĹäžŮæžĹæ■cã;ŞãĹ'■ ReleaseãĀĆ

- Chart áđ'ĐçŘĒä;ĪèŮèřt'æŸŮ

```
Tiller aIjIád'DčŘĚ Chart æUüiijŇčŽt'æŎěářĚ Chart
→ äžěärŁăĚüä;ĬètŮčŽĎæL'ĂæIJL' Charts āŘŁăžüäÿžäÿĂäÿĬ
→ ReleaseiijŇăŘŇæUüaijăěĂšçžŽ KubernetesăĂĈ
ăŽăæ■d' Tiller
→ äžüäÿ■èt'Sèt'ččŏaçŘĚă;ĬètŮăžŇéŮt'čŽĎăŘăŁĬéăžăžŘăĂĈChart
→ äÿ■čŽĎăžŤčŤĬéIJĂèēAēĈ;ăd'SèĜĬéăŇăd'DčŘĚă;ĬètŮăĚšçšžăĂĈ
```

22.3 22.3 éČĭčĭš Helm

ăŎŸæŮž github āIJřăĬĂiijŽ<https://github.com/helm/helm>

- äÿŇē;ăžŇēŁăŁŮčL'ĬæIJŇiijŇēĝčăŎŇăžŮăŎL'èĈĚ helm

```
$ wget https://storage.googleapis.com/kubernetes-helm/helm-v2.13.1-
→ linux-amd64.tar.gz
$ tar xf helm-v2.13.1-linux-amd64.tar.gz
$ mv helm /usr/local/bin/
```

- āĬĬăĝŇăŇŮ tiller æUüăĂŽăijŽèĜĬăĬĬérzăŘŮ ~/.kube
čŽŏă;ŤiijŇæL'ĂăžēĬIJĂèēAçăŏăĬ config æŮĜăžŮă■ŸăIJĬăžŮēŏd'èřAæĬŘăĬš
- tiller éĚ■čĭ;ŏ rbacĭijŇæŮřăžž rbac-config.yamlĭijŇăžŮăžŤčŤĬ

```
https://github.com/helm/helm/blob/master/docs/rbac.md #
→ āIJĬèĚŽăÿĬéăŧéĬčăÿ■æL'čăĬř rbac-config.yaml
```

```
$ kubectl apply -f tiller-rbac.yaml
```

- āĬĬăĝŇăŇŮ tiller æUüăĂŽăijŽèĜĬăĬĬérzăŘŮ ~/.kube
čŽŏă;ŤiijŇæL'ĂăžēĬIJĂèēAçăŏăĬ config æŮĜăžŮă■ŸăIJĬăžŮēŏd'èřAæĬŘăĬš

```
$ helm init --service-account tiller
```

- æŮžăĬă incubator æžŘ

```
$ helm repo add incubator https://aliacs-app-catalog.oss-cn-
→ hangzhou.aliyuncs.com/charts-incubator/
$ helm repo update
```

- āŎL'èĈĚăŎŇæĬŘiijŇæšĉIJŇčL'ĬæIJŇ

```
$ helm version
```

```
Client: &version.Version{SemVer:"v2.13.1", GitCommit:
→ "618447cbf203d147601b4b9bd7f8c37a5d39fbb4", GitTreeState:"clean"}
Server: &version.Version{SemVer:"v2.9.1", GitCommit:
→ "20adb27c7c5868466912eebdf6664e7390ebe710", GitTreeState:"clean"}
```


- helm åŏÿæŨzâŔŕçŦŧçŽĐ chart äzŞăžŞ


`http://hub.kubeapps.com/`

- ǎŝ;äzd'ǎŝžæIŋä;£çŦí

```
completion #
→äÿžæŃĜăŏŽčŽĎshellçŤšæĹŘèĜĭăĹĹăŎŃæĹŘèĎŽæIJŋiijĹbashæĹŮzshiiJĹ'
create # ăĹŽăžžăÿĂăÿĭăĹŮæIJĹ'çžžăŏžăŘ■çĝřçŽĎŮř chart
delete # äžŎ Kubernetes ăĹăéžď'æŃĜăŏžăŘ■çĝřçŽĎ release
dependency # çŏaçŤŤ chart çŽĎăçîèťŮăĹşçşz
fetch # äžŎă■ŸăĹĭăžšăÿŃë;; chart
→ăžúiiijĹăĹŕéĂĹ' iijĹ'ăŕĖăĹŮëğçăŎŃçijĹ'ăĹŕæIJŋăIJŕçŽŏă;Ťăÿ■
get # äÿŃë;;ăÿĂăÿĭăš;ăŘ■ release
help # ăĹŮăĜžæĹ'ĂæIJĹ'ăÿŏăĹĹ'ăĹăæĂŕ
history # èŎŮăĹŮ release ăŎĖăĹŕš
home # æŸçď'ž HELM_HOME çŽĎă;■ç;ŏ
init # ăIJĹăŏçăĹŮçŋŕăšŃæIJ■ăĹăăžĭăÿĹăĹĭăğŃăŃŮHelm
inspect # æčĂæšĕ chart èŕççžĖăĹăæĂŕ
install # ăŎĹ'èčĚ chart ă■Ÿăăč
lint # ăŕž chart èĹžèăŃèŕ■æşŤæčĂæšĕ
list # releases ăĹŮëăĹ
package # ăŕĖ chart çŽŏă;ŤăĹ'šăŃĚăĹĹ chart æăčăăĹ
plugin # æûžăĹăăĹŮëăĹăĹŮăĹăéžď' helm æŖšăžú
repo # æûžăĹăăĹŮëăĹăĹăéžď'æžt'æŮŕăšŃçt'ćăijŤ chart ă■ŸăĹĭăžš
reset # äžŎéžĖççď'ăÿ■ă■ÿë;; Tiller
rollback # ăŕĖçĹ'ĹæIJŋăžďæžžăĹŕăžăĹ'■çŽĎçĹ'ĹæIJŋ
search # ăIJĹ chart ă■ŸăĹĭăžšăÿ■æŖIJçt'ćăĖşéŤŏă■Ů
serve # ăŖŕăĹĹæIJŋăIJŕhttpç;šçžIJæIJ■ăĹăăžĭ
status # æŸçď'žæŃĜăŏŽ release çŽĎçĹŮæĂă
template # æIJŋăIJŕæÿşæššăĹăæĭĹ
test # æŤŃèŕŤăÿĂăÿĭ release
upgrade # ă■ĜçžğăÿĂăÿĭ release
verify # éĭŃèŕĂççžăŏžèŮŕăçĎăÿĹçŽĎ chart
→æŸŕăŖëăŮşçç;ăŘ■ăÿŤæIJĹ'æŤĹ
version # æĹ'šă■ŕăŏçăĹŮçŋŕ/æIJ■ăĹăăžĭçĹ'ĹæIJŋăĹăæĂŕ
dep # ăĹĖăďŖ Chart ăžúăÿŃë;;ăçîèťŮ
```

- æŃĜǻőŽ values.yaml éČĺç;šäyÄäył chart

```
helm install --name els1 -f values.yaml stable/elasticsearch
```

-  GçžgäyÄäyı chart

```
helm upgrade --set mysqlRootPassword=passwd db-mysql stable/mysql
```

- $\mathfrak{aZdæzZ\ddot{a}y\ddot{A}äy\mathfrak{l}}$ chart

```
helm rollback db-mysql 1
```

- åŁăÉŽd'äŸĂäŸł release

```
helm delete --purge db-mysql
```

- **Chart** is a package of Kubernetes manifests that can be used to install and manage a set of related Kubernetes resources.

```
helm install/upgrade xxx --dry-run --debug
```

22.4 Chart

```
myapp/
├── charts
│   └── charts.yaml
├── templates
│   ├── deployment.yaml
│   ├── _helpers.tpl
│   ├── ingress.yaml
│   ├── NOTES.txt
│   ├── service.yaml
│   ├── tests
│   └── test-connection.yaml
└── values.yaml
```

The `charts` directory contains a sub-chart named `charts`. The `charts.yaml` file defines the chart's metadata. The `templates` directory contains the Kubernetes manifests for the chart. The `values.yaml` file contains the default values for the chart's parameters.

22.5 Helm + Ceph éČlčš EFK

Helm is a package manager for Kubernetes. It allows you to install and manage a set of related Kubernetes resources (charts) in a consistent and repeatable way. Ceph is a distributed storage system that can be used to store data in a scalable and reliable way. EFK (Elasticsearch, Fluentd, Kibana) is a logging stack that can be used to collect, store, and analyze logs.

The `charts` directory contains a sub-chart named `charts`. The `charts.yaml` file defines the chart's metadata. The `templates` directory contains the Kubernetes manifests for the chart. The `values.yaml` file contains the default values for the chart's parameters.

helm install `ÉŸzāđēfGčlŇäijZäyNē;ēTijāČRāRřēČ;äijZæřTē;ČæĚčāČ`

helm `éGŇéÍcæIJL'ā;Ĺād'ZāRřäzēāōZāLŮčZDēāzčZōrijNēfZéGŇæĹSāřsāyāōZāLŮāzErijNāRāčæĹS`

22.6 22.6 Storage Class

```

---
apiVersion: v1
kind: Secret
metadata:
  name: ceph-admin-secret
  namespace: kube-system
type: "kubernetes.io/rbd"
data:
  # ceph auth get-key client.admin | base64
  key: QVFER3U5TmMQNXQ4SlJBAAhHMGltdXZlNFZkUXAvN2tTZ1BENGc9PQ==

---
apiVersion: v1
kind: Secret
metadata:
  name: ceph-secret
  namespace: kube-system
type: "kubernetes.io/rbd"
data:
  # ceph auth get-key client.kube | base64
  key: QVFCcUM5VmNWVDdQCCCCWR1NUxFNfVKeTAiazdUWVhOa3N2UWc9PQ==

---
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ceph-rbd
provisioner: ceph.com/rbd
reclaimPolicy: Retain
parameters:
  monitors: 172.16.100.9:6789
  pool: kube
  adminId: admin
  adminSecretName: ceph-admin-secret
  adminSecretNamespace: kube-system
  userId: kube
  userSecretName: ceph-secret
  userSecretNamespace: kube-system
  fsType: ext4
  imageFormat: "2"
  imageFeatures: "layering"

```

22.7 Helm Elasticsearch

- `helm fetch stable/elasticsearch`

```
helm fetch stable/elasticsearch
```

- `values.yaml` `storageClass` `ceph-rbd`

```
storageClass: "ceph-rbd"
```

- `helm install --name elasticsearch -f values.yaml stable/elasticsearch`

```
helm install --name elasticsearch -f values.yaml stable/elasticsearch
```

- `kubectl get pods`

```
$ kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
elasticsearch-client-55696f5bdd-qczbf	1/1	Running	1	78m
elasticsearch-client-55696f5bdd-tdwdc	1/1	Running	1	78m
elasticsearch-data-0	1/1	Running	1	78m
elasticsearch-data-1	1/1	Running	1	56m
elasticsearch-master-0	1/1	Running	1	78m
elasticsearch-master-1	1/1	Running	1	53m
elasticsearch-master-2	1/1	Running	1	52m
rbd-provisioner-9b8ffbcc-nxdjd	1/1	Running	2	81m

- `helm status elasticsearch`

```
$ helm status elasticsearch
```

```
LAST DEPLOYED: Sun May 12 16:28:56 2019
NAMESPACE: default
STATUS: DEPLOYED
```

```
RESOURCES:
==> v1/ConfigMap
```

NAME	DATA	AGE
elasticsearch	4	88m
elasticsearch-test	1	88m

(continues on next page)

(continued from previous page)

```

==> v1/Pod(related)
NAME                                READY   STATUS    RESTARTS   AGE
els1-elasticsearch-client-55696f5bdd-qczbf  1/1     Running   0           88m
els1-elasticsearch-client-55696f5bdd-tdwdc  1/1     Running   0           88m
els1-elasticsearch-data-0                  1/1     Running   0           88m
els1-elasticsearch-data-1                  1/1     Running   0           66m
els1-elasticsearch-master-0                1/1     Running   0           88m
els1-elasticsearch-master-1                1/1     Running   0           63m
els1-elasticsearch-master-2                1/1     Running   0           62m

==> v1/Service
NAME                                TYPE        CLUSTER-IP    EXTERNAL-IP
els1-elasticsearch-client           ClusterIP    10.98.197.185 <none>
els1-elasticsearch-discovery        ClusterIP    None           <none>

==> v1/ServiceAccount
NAME                                SECRETS   AGE
els1-elasticsearch-client           1         88m
els1-elasticsearch-data             1         88m
els1-elasticsearch-master           1         88m

==> v1beta1/Deployment
NAME                                READY   UP-TO-DATE   AVAILABLE   AGE
els1-elasticsearch-client           2/2     2             2           88m

==> v1beta1/StatefulSet
NAME                                READY   AGE
els1-elasticsearch-data             2/2     88m
els1-elasticsearch-master           3/3     88m

NOTES:
The elasticsearch cluster has been installed.

Elasticsearch can be accessed:

* Within your cluster, at the following DNS name at port 9200:

```

(continues on next page)

(continued from previous page)

```

    els1-elasticsearch-client.default.svc

* From outside the cluster, run these commands in the same shell:

    export POD_NAME=$(kubectl get pods --namespace default -l
    ↪ "app=elasticsearch,component=client,release=els1" -o jsonpath="{.
    ↪ items[0].metadata.name}")
    echo "Visit http://127.0.0.1:9200 to use Elasticsearch"
    kubectl port-forward --namespace default $POD_NAME 9200:9200

```

- `kubectl run cirros1 --rm -it --image=cirros -- /bin/sh`

```

$ kubectl run cirros1 --rm -it --image=cirros -- /bin/sh

/ # nslookup els1-elasticsearch-client.default.svc
Server:      10.96.0.10
Address 1: 10.96.0.10 kube-dns.kube-system.svc.cluster.local

Name:      els1-elasticsearch-client.default.svc
Address 1: 10.98.197.185 els1-elasticsearch-client.default.svc.
↪ cluster.local
/ # curl els1-elasticsearch-client.default.svc.cluster.local:9200/_
↪ cat/nodes
10.244.2.28 7 96 2 0.85 0.26 0.16 di - els1-elasticsearch-data-0
10.244.1.37 7 83 1 0.04 0.06 0.11 di - els1-elasticsearch-data-1
10.244.2.25 19 96 2 0.85 0.26 0.16 i - els1-elasticsearch-client-
↪ 55696f5bdd-tdwdc
10.244.2.27 28 96 2 0.85 0.26 0.16 mi * els1-elasticsearch-master-2
10.244.1.39 19 83 1 0.04 0.06 0.11 i - els1-elasticsearch-client-
↪ 55696f5bdd-qczbf
10.244.2.29 21 96 2 0.85 0.26 0.16 mi - els1-elasticsearch-master-1
10.244.1.38 23 83 1 0.04 0.06 0.11 mi - els1-elasticsearch-master-0

```

22.8 Helm fluentd-elasticsearch

- `helm repo add kiwigrd https://kiwigrd.github.io`

```
helm repo add kiwigrd https://kiwigrd.github.io
```

- `helm fetch kiwigrd/fluentd-elasticsearch`

```
helm fetch kiwigrd/fluentd-elasticsearch
```

- `els1-elasticsearch-client.default.svc.cluster.local:9200`

```
els1-elasticsearch-client.default.svc.cluster.local:9200
```

- `values.yaml` file for `elasticsearch` service

```
elasticsearch:
  host: 'els1-elasticsearch-client.default.svc.cluster.local'
  port: 9200
```

- `elasticsearch` pod is scheduled on the `Master` node.

```
tolerations:
- key: node-role.kubernetes.io/master
  operator: Exists
  effect: NoSchedule
```

- `elasticsearch` pod is scheduled on the `Master` node.

```
podAnnotations:
  prometheus.io/scrape: "true"
  prometheus.io/port: "24231"

service:
  type: ClusterIP
  ports:
  - name: "monitor-agent"
    port: 24231
```

- `elasticsearch` pod is scheduled on the `Master` node.

```
helm install --name fluel -f values.yaml kiwigrid/fluentd-
elasticsearch
```

- `elasticsearch` pod is scheduled on the `Master` node.

```
[root@master fluentd-elasticsearch]# helm status fluel
LAST DEPLOYED: Sun May 12 18:13:12 2019
NAMESPACE: default
STATUS: DEPLOYED

RESOURCES:
==> v1/ClusterRole
NAME                                AGE
fluel-fluentd-elasticsearch        17m

==> v1/ClusterRoleBinding
NAME                                AGE
fluel-fluentd-elasticsearch        17m

==> v1/ConfigMap
NAME                                DATA  AGE
fluel-fluentd-elasticsearch        6      17m

==> v1/DaemonSet
```

(continues on next page)

(continued from previous page)

NAME	DESIRED	CURRENT	READY	UP-TO-DATE	
↪AVAILABLE	NODE SELECTOR	AGE			
fluentd-elasticsearch	3	3	3	3	3
↪	<none>	17m			
==> v1/Pod(related)					
NAME	READY	STATUS	RESTARTS	AGE	
fluentd-elasticsearch-p49fc	1/1	Running	1	17m	
fluentd-elasticsearch-q5b9k	1/1	Running	0	17m	
fluentd-elasticsearch-swftv	1/1	Running	0	17m	
==> v1/Service					
NAME	TYPE	CLUSTER-IP	EXTERNAL-IP		
↪PORT(S)	AGE				
fluentd-elasticsearch	ClusterIP	10.106.106.209	<none>		
↪24231/TCP	17m				
==> v1/ServiceAccount					
NAME	SECRETS	AGE			
fluentd-elasticsearch	1	17m			
NOTES:					
1. To verify that Fluentd has started, run:					
<pre>kubectl --namespace=default get pods -l "app.kubernetes.io/ ↪name=fluentd-elasticsearch,app.kubernetes.io/instance=fluentd"</pre>					
THIS APPLICATION CAPTURES ALL CONSOLE OUTPUT AND FORWARDS IT TO ↪					
↪elasticsearch . Anything that might be identifying,					
including things like IP addresses, container images, and object ↪					
↪names will NOT be anonymized.					
2. Get the application URL by running these commands:					
<pre>export POD_NAME=\$(kubectl get pods --namespace default -l "app. ↪kubernetes.io/name=fluentd-elasticsearch,app.kubernetes.io/ ↪instance=fluentd" -o jsonpath="{.items[0].metadata.name}") echo "Visit http://127.0.0.1:8080 to use your application" kubectl port-forward \$POD_NAME 8080:80</pre>					

- æŸřăŘęçŤşæĹŖăžEçť'ćăijŤijŇçŽť æŌěă;ŕçŤlěőŕéŮő elasticsearch çŽĐ RESTfull API æŌěăŖčăĂĆ

```
$ kubectl run cirros1 --rm -it --image=cirros -- /bin/sh
/ # curl elsl-elasticsearch-client.default.svc.cluster.local:9200/_
↪cat/indices
green open logstash-2019.05.10 a2b-GyKsSLOZPqGKbCpyJw 5 1 158 0 ↪
↪84.2kb 460b
green open logstash-2019.05.09 CwYylNhdRf-A5UELhrzHow 5 1 71418 0 ↪
↪34.3mb 17.4mb
```

(continues on next page)

(continued from previous page)

```
green open logstash-2019.05.12 5qRFpV46RGG_bWC4xbsyVA 5 1 34496 0
→26.1mb 13.2mb
```

22.9 Helm kibana

- `stable/kibana`

```
helm fetch stable/kibana
```

- `values.yaml` `elasticsearch` `elasticsearch`

```
elasticsearch.hosts: http://els1-elasticsearch-client.default.svc.
→cluster.local:920
```

- `service` `type: NodePort`

```
service:
  type: NodePort
```

- `helm install --name kib1 -f values.yaml stable/kibana`

```
helm install --name kib1 -f values.yaml stable/kibana
```

- `service` `type: NodePort`

```
$ kubectl get svc
NAME                                TYPE                CLUSTER-IP      EXTERNAL-IP
→IP    PORT(S)          AGE
els1-elasticsearch-client          ClusterIP           10.98.197.185    <none>
→      9200/TCP        4h51m
els1-elasticsearch-discovery       ClusterIP           None             <none>
→      9300/TCP        4h51m
flu1-fluentd-elasticsearch         ClusterIP           10.101.97.11     <none>
→      24231/TCP      157m
kib1-kibana                        NodePort            10.103.7.215     <none>
→      443:31537/TCP  6m50s
kubernetes                         ClusterIP           10.96.0.1        <none>
→      443/TCP        3d4h
```

- `service` `type: NodePort` `172.16.100.6:31537`

```
172.16.100.6:31537
```

CHAPTER 23

Introducing ETCD

23.1 ETCD

ETCD is a distributed key-value store that provides a reliable way of storing data that can be accessed by a large number of clients. It is designed to be highly available and consistent, and it is used by many other systems, including Kubernetes.

23.1.1 ETCD Architecture

ETCD is built on top of the Raft consensus algorithm, which ensures that all nodes in the cluster agree on the state of the key-value store. The architecture consists of a cluster of nodes, each of which can act as a leader or a follower. The leader is responsible for managing the cluster and for committing new data to the store. The followers replicate the data from the leader and can also act as leaders if the current leader fails.

23.1.2 ETCD Features

2013-06

Initial Commit

- CoreOS contribution

2014-06¹

etcd v0.2

- Kubernetes v0.4
- 10x community

2015-02

etcd v2.0 First Please Release

- Raft consistency protocol
- 1000s of writes/s

2017-01

etcd v3.1

- New API
- Fast linearized read
- gRPC proxy

2018-11

CNCF Incubation

- 30+ of projects use etcd
- 400+ contributors
- 9 maintainers of 8 companies

2019

etcd v3.4

- learner member
- fully concurrent read
- performance enhancement

23.1.5.2 leader

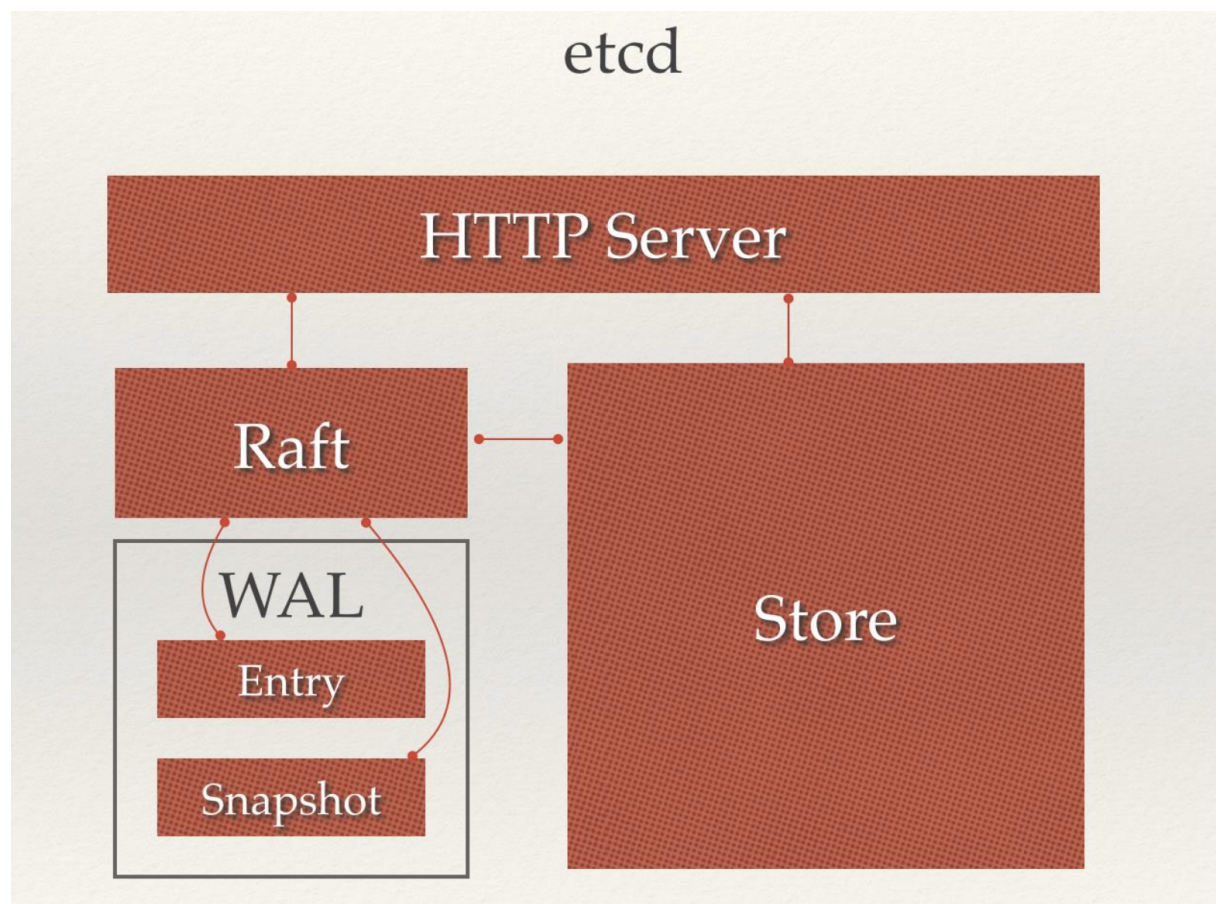
etcd uses Raft for consensus. The leader is responsible for proposing and committing updates to the etcd database. It also handles client requests for read-only data. The leader is elected by the other etcd nodes in the cluster. The election process is based on the Raft algorithm, which ensures that there is only one leader at any given time.

23.1.5.3 etcd

etcd uses Raft for consensus. The leader is responsible for proposing and committing updates to the etcd database. It also handles client requests for read-only data. The leader is elected by the other etcd nodes in the cluster. The election process is based on the Raft algorithm, which ensures that there is only one leader at any given time. The etcd database is distributed across all nodes in the cluster. Each node has a local copy of the database, and they all replicate updates from the leader. This ensures that the database is highly available and consistent.

23.2 ETCD

23.2.1 etcd



23.2.2 23.2.2 ædūædDègčædŘ

äzŎ etcd çŽDædūædDāZ; äy■æĹSāznāRřazēcIJNāĹřiiJNetcd
äyžèeAāĹEāyžāZŽāyĹéČĹāĹEāĀĆ

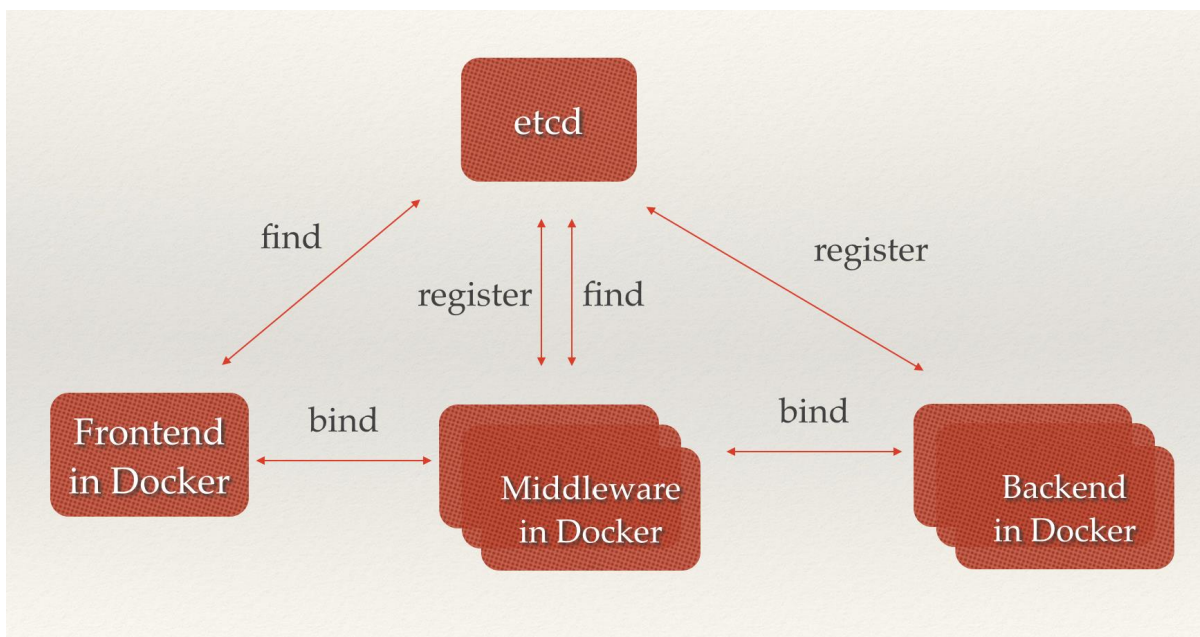
- HTTP ServeriiJŽčTĹāžŎād' DčRĚčTĹāĹūāRŠéĀAçŽD API èrūæsČäzēāRĹāĒūāōČ etcd èĹČçČžčŽDāRŊæ■ēäyŎāĹČēūšāĹqæAřèrūæsČāĀĆ
- StoreiiJŽčTĹāžŎād' DčRĚ etcd æTřæNāçŽDāRĎčšzāĹšèČ; çŽDāžNāĹāiiJNāNĚæNñæTřæ■ōčt' cāijTāĀ etcd āřžčTĹāĹūāRŘā; ŽčŽDād' gād' ŽæTř API āĹšèČ; çŽDāĒūā; ŠāōđčŎřāĀĆ
- RaftiiJŽRaft āijžäyĀèGt' æĀgčōŬæsTçŽDāĒūā; ŠāōđčŎřiiJNæŸr etcd çŽDæäyāĹČāĀĆ
- WALiiJŽWrite Ahead LogiiJĹécDāĒŽāijRæŬēāĹŬiiJL'iiJNæŸr etcd çŽDæTřæ■ōā■ŸāČĹāŬzāijRāĀĆéZd' āžEāIJĹāĒEā■Ÿäy■ā■ŸæIJL' æĹĀĹæTřæ■ōčŽDçĹūāĀAžēā āřséĀŽēĹG WAL èĹZēāNæNāžĒāNŬā■ŸāČĹāĀĆWAL äy■iiJNæĹĀĹæIJL' çŽDæTřæ■ōæRŘāžd' āĹ■ēČ; āijžāžNāĒĹēōřā; TřæŬēāĹŬāĀĆ
 - Snapshot æŸrāyžāžĒEēŸsæ■cæTřæ■ōēĹGād' ŽēĀNēĹZēāNçŽDçĹūāĀAāĹñçĒgiiJŽ
 - Entry èāĹčd' žā■ŸāČĹçŽDāĒūā; ŠæŬēāĹŬāĒEāōžāĀĆ

éĀŽāyŸiiJNāyĀāyĹçTĹāĹūçŽDèrūæsČāRŠéĀAēĹGæĹēriJNāijŽčžRçTš HTTP Server è;ñāRŠçžŽ Store èĹZēāNāĒūā; ŠçŽDāžNāĹāād' DčRĚiiJNāçČædIJæŬĹāRĹāĹēĹČçčžčŽDāĹōæTžiiJNāĹZāž Raft æĹāāĹŬēĹZēāNçĹūāĀAçŽDāRŸæŽt' āĀAæŬēāĹŬçŽDēōřā; TřiiJNçDūāRŎāĒē■āRŊæ■ēčžZāĹñçŽD etcd èĹČçČžäzēçāōēōd' æTřæ■ōæRŘāžd' iiJNæIJĀāRŎēĹZēāNæTřæ■ōčŽDæRŘāžd' iiJNāĒē■āñāRŊæ■ēāĀĆ

23.3 23.3 āžTčTĹāĹJžæŽr

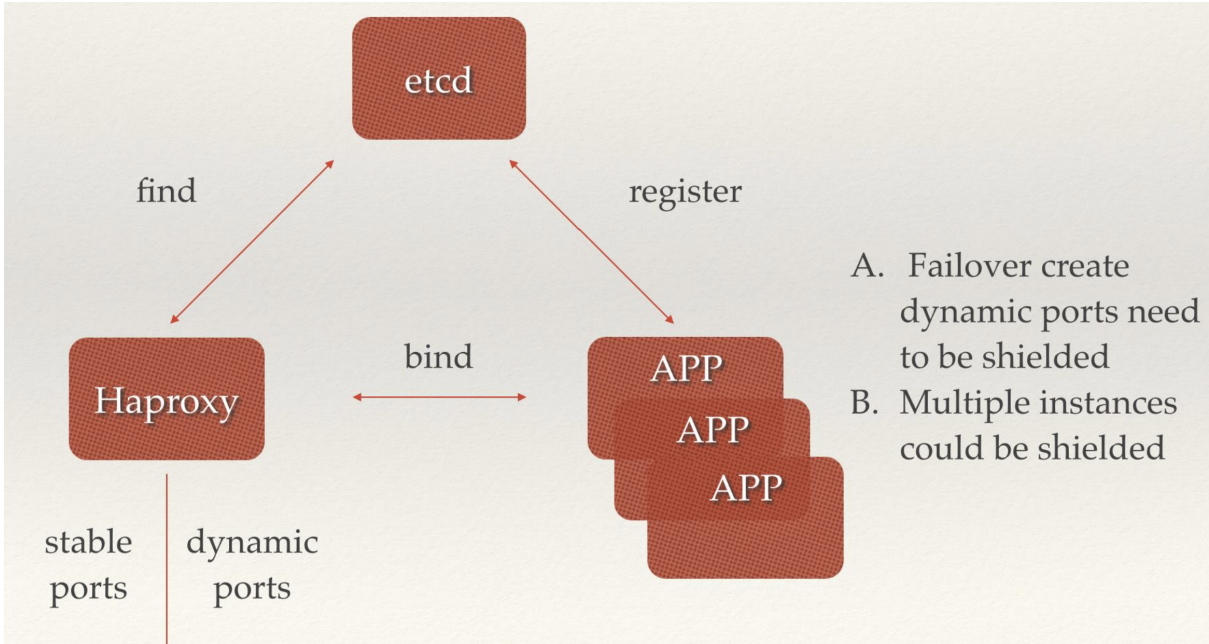
23.3.1 23.3.1 æĹJ■āĹæšĹāĒNāRŠçŎř

- āĹ■āRŎçñrāyŽāĹæšĹāĒNāRŠçŎř



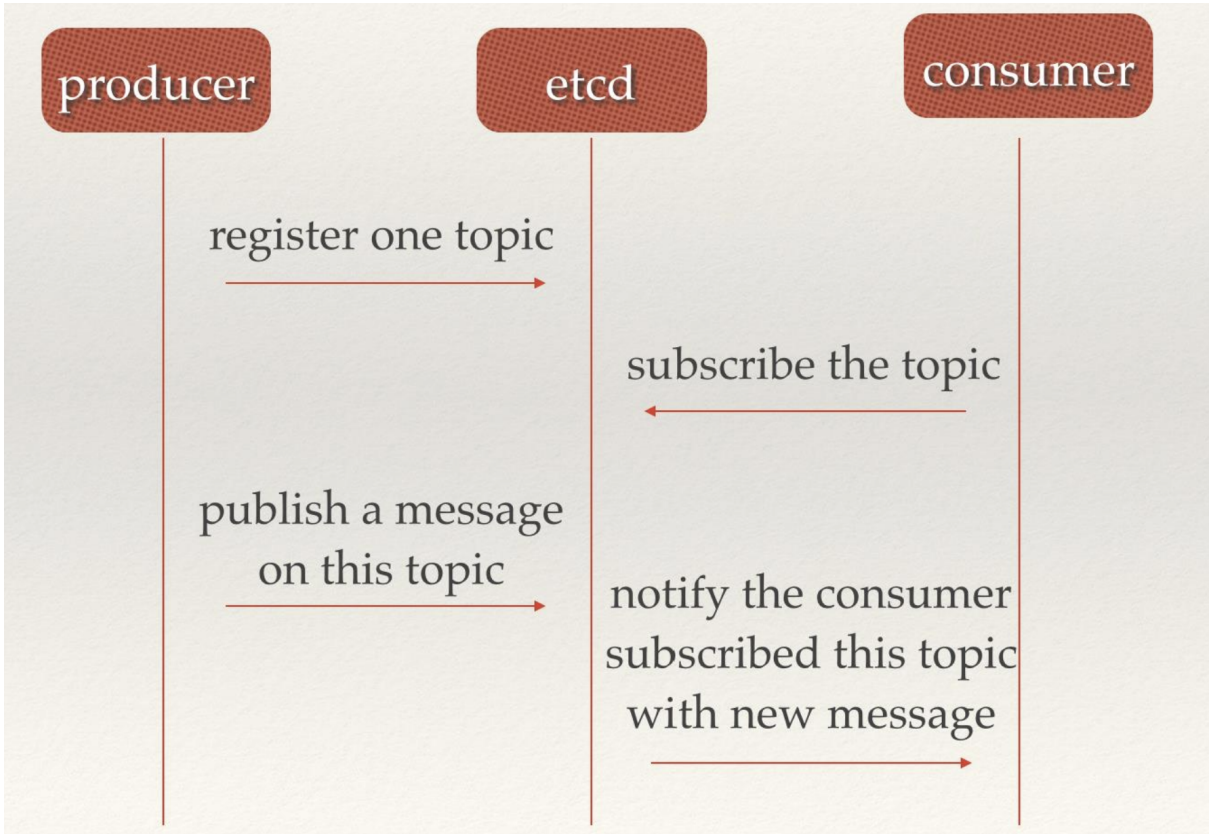
äy■éÜt' äzüâũşçzRâRŎçnræIJ■âŁaâIJletcdäy■æşláEŇijŇâL'■çnrâŠŇäy■éÜt' äzüâRřäzěâŁLè; zæİŁçŽDžäz

- âd'ŽçzDâRŎçnræIJ■âŁaâŽÍæşláEŇâRŠçŎř



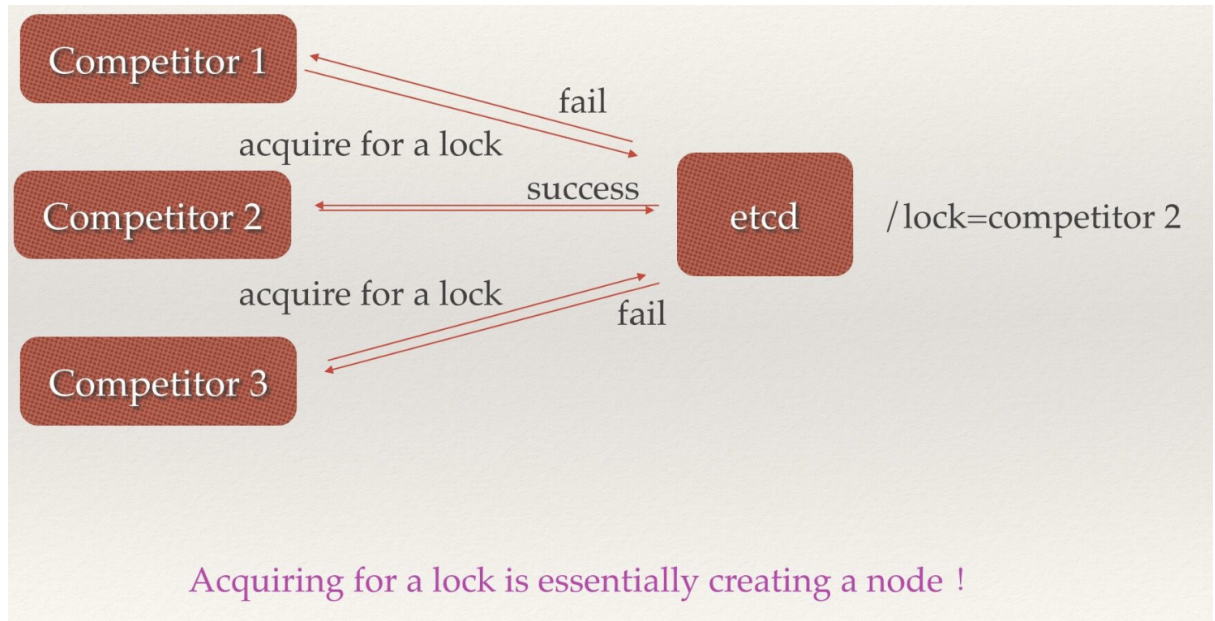
âRŎçnrâd'ŽäyŁæŬăçŁŭæĂAçŽyâRŇâL'řæIJŇçŽDappâRřäzěâRŇäžŇæşláEŇâL'řetcdäy■ijŇâL'■çnrâRřäz

23.3.2 23.3.2 æŭŁæAřâRŠäyČäyŎëöcéŸĚ



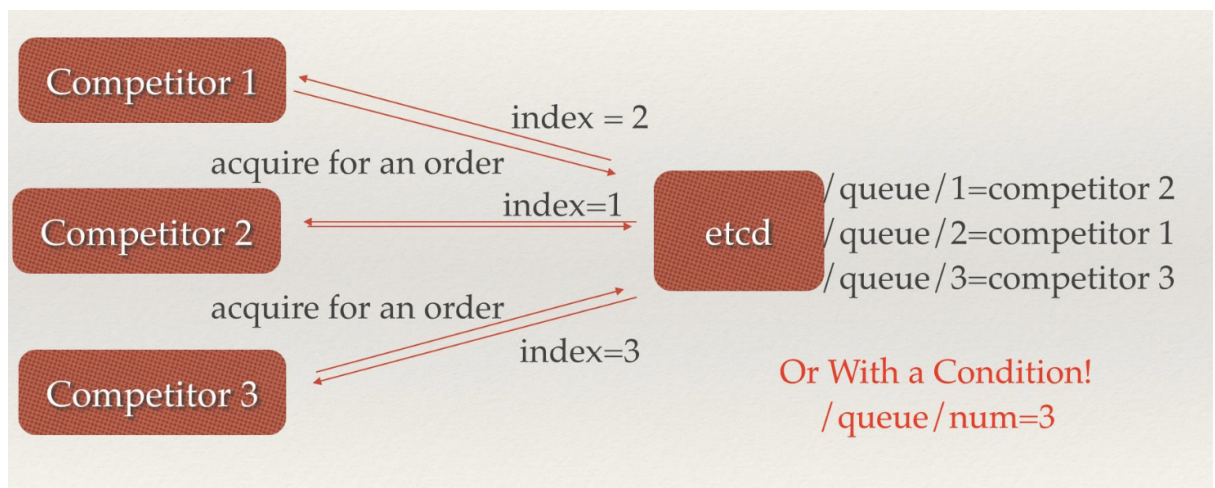
- `setcd watch` `etcd` `lock` `competitor 2`
- `etcd` `lock` `competitor 2` `lock=competitor 2`
- `etcd` `lock` `competitor 2` `lock=competitor 2`

23.3.5 23.3.5 `etcd` `lock` `competitor 2`



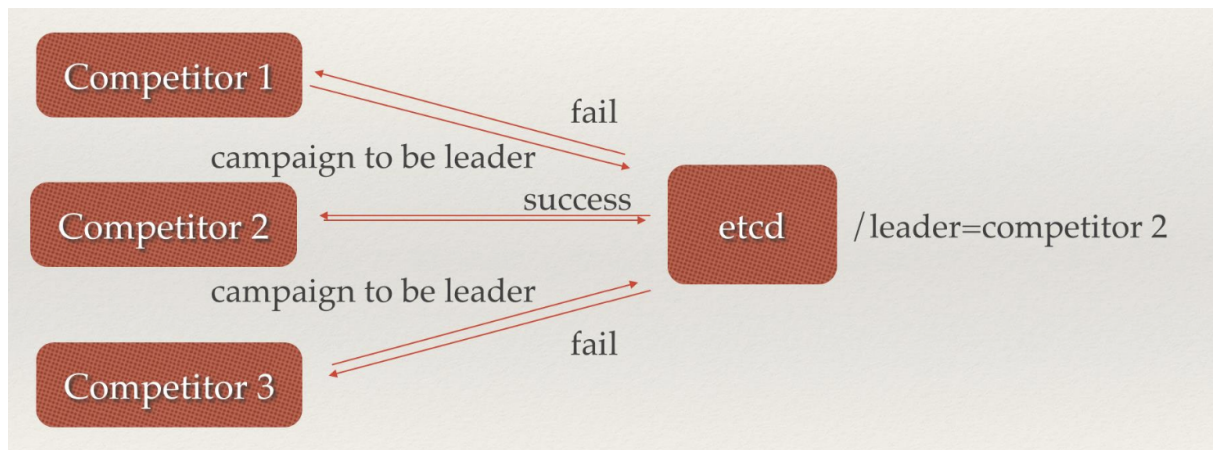
`etcd` `lock` `competitor 2` `lock=competitor 2`

23.3.6 23.3.6 `etcd` `lock` `competitor 2`



`etcd` `lock` `competitor 2` `lock=competitor 2`

23.3.7 23.3.7 éŽĚçŁd'āŘŁĆŁŚä;äÿŒLeaderéĀL'äÿŁ



etcdāŘřäzëæāžæ■ōraftçŒŮæşŤāIJĹād'ŽäÿŁnodeèŁĆçĆzæĹëéĀL'äÿŁ;āĜžleader

23.4 23.4 āŒL'èĉĚéČĭş

23.4.1 23.4.1 ā■ŤæIJžāŒL'èĉĚ

āŘřäzëä;ŁçŤĹäžŇëŁZāŁŮæĹŮæžŘçăĀäÿŇè;ĭ;āŒL'èĉĚĭjŇä;ĒæŸřā■śāŒşéIJĀèçĀèĜĹāŮśāĒZéĚ■ç;ŏæŮĜă

```

hostnamectl set-hostname etcd-1
wget http://dl.fedoraproject.org/pub/epel/epel-release-latest-7.
↳noarch.rpm
rpm -ivh epel-release-latest-7.noarch.rpm
# yum äžŞäžŞäÿ■çŽĎetcdçŁ'ĹæIJŇäÿž3.3.
↳11iijŇăçĀēIJéIJĀèçĀæIJĀŮřçŁ'ĹæIJŇçŽĎetcdāŘřäzëèŁZèqŇäžŇèŁZāŁŮāŒL'èĉĚ
yum -y install etcd
systemctl enable etcd
  
```

āŘřäzëæşçIJŇyumaŒL'èĉĚçŽĎetcdçŽĎæIJŁæŤĹéĚ■ç;ŏæŮĜăžŮiijŇăäžæ■ŏèĜĹāŮşçŽĎéIJĀæśĀĹëä

- etcd éžŸèŒd'ārĒæŤřæ■ŏā■ŸæŤĭ;āĹřā;ŞāŁ■ēŮřāĭĎçŽĎ default.etcd/ çŽŒā;ŤäÿŇ
- āIJĹ http://localhost:2380 āŞŇéŽĚçŁd'äÿ■āĒŮäžŮèŁĆçĆzéĀŽăġ
- āIJĹ http://localhost:2379 æŘŘăĭŽ HTTP API
æIJ■āĹiijŇăĭŽăŒçæĹŮçŇřäžd'äžŞ
- èřèèŁĆçĆzéçŽĎāŘ■çğřéžŸèŒd'äÿž default
 - heartbeat äÿž 100msiijŇăŘŒéĹçäijŽèřt'æŸŒèŁZäÿĹéĚ■ç;ŏçŽĎä;IJçŤĹ
- election äÿž 1000msiijŇăŘŒéĹçäijŽèřt'æŸŒèŁZäÿĹéĚ■ç;ŏçŽĎä;IJçŤĹ
- snapshot count äÿž 10000iijŇăŘŒéĹçäijŽèřt'æŸŒèŁZäÿĹéĚ■ç;ŏçŽĎä;IJçŤĹ
- éŽĚçŁd'āŞŇæřĀäÿĹèŁĆçĆzéČ;äijŽçŤşæĹŘäÿĀäÿŁ uuid

- 安装 etcd 服务

```
[root@VM_0_8_centos tmp]# grep -Ev "^#|^$" /etc/etcd/etcd.conf
ETCD_DATA_DIR="/var/lib/etcd/default.etcd"
ETCD_LISTEN_CLIENT_URLS="http://localhost:2379"
ETCD_NAME="default"
ETCD_ADVERTISE_CLIENT_URLS="http://localhost:2379"
[root@VM_0_8_centos tmp]# systemctl status etcd
```

23.4.2 23.4.2 安装 etcd 服务

23.4.2.1 安装 etcd 服务

ID/名称	监控	状态	可用区	实例类型	实例配置	主IP地址	实例计费模式	网络计费模式	所属项目	操作
搜索“所属项目:默认项目”，找到 3 条结果 返回原列表										
<input type="checkbox"/> ins-2s58otg xuel-tmp-node3		运行中	广州四区	标准型S4	1核 1GB 1Mbps 系统盘: 高性能云硬盘 网络: Default-VPC	106.53.65.233 (公) 172.16.0.14 (内)	按量计费 2019-12-03 14:36:20创建	按流量计费	默认项目	登录 更多
<input type="checkbox"/> ins-m58rtap xuel-tmp-node2		运行中	广州四区	标准型S4	1核 1GB 1Mbps 系统盘: 高性能云硬盘 网络: Default-VPC	106.52.249.36 (公) 172.16.0.17 (内)	按量计费 2019-12-03 14:36:18创建	按流量计费	默认项目	登录 更多
<input type="checkbox"/> ins-qjvi70s xuel-tmp-node1		运行中	广州四区	标准型S4	1核 1GB 1Mbps 系统盘: 高性能云硬盘 网络: Default-VPC	129.204.172.112 (公) 172.16.0.8 (内)	按量计费 2019-12-03 14:36:13创建	按流量计费	默认项目	登录 更多

共 3 项

每页显示行 20

1 / 1 页

安装 etcd 服务	操作系统	IP地址	etcd 服务
etcd-0-8	CentOS 7.3	172.16.0.8	etcd
etcd-0-17	CentOS 7.3	172.16.0.17	etcd
etcd-0-14	CentOS 7.3	172.16.0.14	etcd

23.4.2.2 HOSTS 配置

安装 etcd 服务

```
cat >> /etc/hosts << EOF
172.16.0.8 etcd-0-8
172.16.0.14 etcd-0-14
172.16.0.17 etcd-0-17
EOF
```

23.4.2.3 ETCD 安装

安装 etcd 服务

(continued from previous page)

```

#
#[Clustering]
ETCD_INITIAL_ADVERTISE_PEER_URLS="http://172.16.0.8:2380"
→ #
→ èrëèLĆĆZāRñäijt'çŽŚāRñāIJřāĪĀiijNèŁŽäÿłāĀijäijŽāŚŁèrL'éŽĚç;d'äÿ■āĚūāzŮèLĆĆZ
ETCD_ADVERTISE_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
→ 8:2379" #
→ ārzād'ŮāĚñāŚŁçŽĎèrëèLĆĆZāōcæŁûçnrçŽŚāRñāIJřāĪĀiijNèŁŽäÿłāĀijäijŽāŚŁèrL'éŽĚç;d'
#ETCD_DISCOVERY=""
#ETCD_DISCOVERY_FALLBACK="proxy"
#ETCD_DISCOVERY_PROXY=""
#ETCD_DISCOVERY_SRV=""
ETCD_INITIAL_CLUSTER="etcd-0-8=http://172.16.0.8:2380,etcd-0-
→ 17=http://172.16.0.17:2380,etcd-0-14=http://172.16.0.14:2380"
→ #
→ éŽĚç;d'äÿ■æL'ĀæIJL'èLĆĆZçŽĎāfā
ETCD_INITIAL_CLUSTER_TOKEN="etcd-token" #
→ āĹŽāzžéŽĚç;d'çŽĎ
→ tokeniijNèŁŽäÿłāĀijjærĤäÿłéŽĚç;d'āĹĪæNĀāĤřäÿĀāĀCèŁŽæūçŽĎèrĪiijNāçCæĪJä;āèçAéĜ
→ uuidiijŽāŘēāĹŽāiijŽāři jeĜt'ād'ŽäÿłéŽĚç;d'āžNéŮt'çŽĎāĚšçĹĀiijNéĀāĹRæIJĹçSěçŽĎéĪ
ETCD_INITIAL_CLUSTER_STATE="new"
#ETCD_STRICT_RECONFIG_CHECK="true"
→ # æŮřāzžéŽĚç;d'çŽĎæŮūāĀŽiijNèŁŽäÿłāĀijäÿž
→ newiijŽāĀĜāçCāũšçžŘā■ŸāIJĹçŽĎéŽĚç;d'iijNèŁŽäÿłāĀijäÿž existing
#ETCD_ENABLE_V2="true"
#
#[Proxy]
#ETCD_PROXY="off"
#ETCD_PROXY_FAILURE_WAIT="5000"
#ETCD_PROXY_REFRESH_INTERVAL="30000"
#ETCD_PROXY_DIAL_TIMEOUT="1000"
#ETCD_PROXY_WRITE_TIMEOUT="5000"
#ETCD_PROXY_READ_TIMEOUT="0"
#
#[Security]
#ETCD_CERT_FILE=""
#ETCD_KEY_FILE=""
#ETCD_CLIENT_CERT_AUTH="false"
#ETCD_TRUSTED_CA_FILE=""
#ETCD_AUTO_TLS="false"
#ETCD_PEER_CERT_FILE=""
#ETCD_PEER_KEY_FILE=""
#ETCD_PEER_CLIENT_CERT_AUTH="false"
#ETCD_PEER_TRUSTED_CA_FILE=""
#ETCD_PEER_AUTO_TLS="false"
#
#[Logging]
#ETCD_DEBUG="false"
#ETCD_LOG_PACKAGE_LEVELS=""

```

(continues on next page)

(continued from previous page)

```
#ETCD_LOG_OUTPUT="default"
#
#[Unsafe]
#ETCD_FORCE_NEW_CLUSTER="false"
#
#[Version]
#ETCD_VERSION="false"
#ETCD_AUTO_COMPACTION_RETENTION="0"
#
#[Profiling]
#ETCD_ENABLE_PPROF="false"
#ETCD_METRICS="basic"
#
#[Auth]
#ETCD_AUTH_TOKEN="simple"
```

etcd-0-8

```
[root@etcd-server ~]# hostnamectl set-hostname etcd-0-8
[root@etcd-0-8 ~]# egrep "^#|^$" /etc/etcd/etcd.conf -v
ETCD_DATA_DIR="/data/app/etcd/"
ETCD_LISTEN_PEER_URLS="http://172.16.0.8:2380"
ETCD_LISTEN_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
→8:2379"
ETCD_NAME="etcd-0-8"
ETCD_INITIAL_ADVERTISE_PEER_URLS="http://172.16.0.8:2380"
ETCD_ADVERTISE_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
→8:2379"
ETCD_INITIAL_CLUSTER="etcd-0-8=http://172.16.0.8:2380,etcd-0-
→17=http://172.16.0.17:2380,etcd-0-14=http://172.16.0.14:2380"
ETCD_INITIAL_CLUSTER_TOKEN="etcd-token"
ETCD_INITIAL_CLUSTER_STATE="new"
```

etcd-0-14

```
[root@etcd-server ~]# hostnamectl set-hostname etcd-0-14
[root@etcd-server ~]# mkdir -p /data/app/etcd/
[root@etcd-0.14 ~]# egrep "^#|^$" /etc/etcd/etcd.conf -v
ETCD_DATA_DIR="/data/app/etcd/"
ETCD_LISTEN_PEER_URLS="http://172.16.0.14:2380"
ETCD_LISTEN_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
→14:2379"
ETCD_NAME="etcd-0-14"
ETCD_INITIAL_ADVERTISE_PEER_URLS="http://172.16.0.14:2380"
ETCD_ADVERTISE_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
→14:2379"
ETCD_INITIAL_CLUSTER="etcd-0-8=http://172.16.0.8:2380,etcd-0-
→17=http://172.16.0.17:2380,etcd-0-14=http://172.16.0.14:2380"
ETCD_INITIAL_CLUSTER_TOKEN="etcd-token"
ETCD_INITIAL_CLUSTER_STATE="new"
```

- etcd-0-7

```
[root@etcd-server ~]# hostnamectl set-hostname etcd-0-17
[root@etcd-server ~]# mkdir -p /data/app/etcd/
[root@etcd-0-17 ~]# egrep "^#|^$" /etc/etcd/etcd.conf -v
ETCD_DATA_DIR="/data/app/etcd/"
ETCD_LISTEN_PEER_URLS="http://172.16.0.17:2380"
ETCD_LISTEN_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
17:2379"
ETCD_NAME="etcd-0-17"
ETCD_INITIAL_ADVERTISE_PEER_URLS="http://172.16.0.17:2380"
ETCD_ADVERTISE_CLIENT_URLS="http://127.0.0.1:2379,http://172.16.0.
17:2379"
ETCD_INITIAL_CLUSTER="etcd-0-8=http://172.16.0.8:2380,etcd-0-
17=http://172.16.0.17:2380,etcd-0-14=http://172.16.0.14:2380"
ETCD_INITIAL_CLUSTER_TOKEN="etcd-token"
ETCD_INITIAL_CLUSTER_STATE="new"
```

- etcd-0-8

```
systemctl start etcd
```

23.4.2.5 etcd-0-8

- etcd-0-8

```
[root@etcd-0-8 default.etcd]# systemctl status etcd
âĖĖ etcd.service - Etcd Server
   Loaded: loaded (/usr/lib/systemd/system/etcd.service; enabled;
   vendor preset: disabled)
   Active: active (running) since âĖĖ 2019-12-03 15:55:28 CST; 8s
   ago
   Main PID: 24510 (etcd)
   CGroup: /system.slice/etcd.service
           âĖĖâĖâĖ24510 /usr/bin/etcd --name=etcd-0-8 --data-dir=/
           data/app/etcd/ --listen-client-urls=http://172.16.0.8:2379

12âĖĖâĖ 03 15:55:28 etcd-0-8 etcd[24510]: set the initial cluster
   version to 3.0
12âĖĖâĖ 03 15:55:28 etcd-0-8 etcd[24510]: enabled capabilities for
   version 3.0
12âĖĖâĖ 03 15:55:30 etcd-0-8 etcd[24510]: peer 56e0b6dad4c53d42
   became active
12âĖĖâĖ 03 15:55:30 etcd-0-8 etcd[24510]: established a TCP
   streaming connection with peer 56e0b6dad4c53d42 (stream Message
   reader)
12âĖĖâĖ 03 15:55:30 etcd-0-8 etcd[24510]: established a TCP
   streaming connection with peer 56e0b6dad4c53d42 (stream Message
   writer)
```

(continues on next page)

23.5 23.5 cŏÅa■Tä;£çTÍ

23.5.1 23.5.1 ácdâŁă

- set

æŊĜăŏŽæšŘäyĤéTŏçŽDăĀijăĂCă;ŊăçC:

```
$ etcdctl set /testdir/testkey "Hello world"
Hello world
```

æŤræŊAçŽDěĂL'éązăŊĚæŊñijŽ

```
--ttl '0'
→èrěéTŏăĀijçŽDěúĚæŮúæŮúéŮt' (ă■Tă;■ăÿžçĝš) iijŊăÿ■éĚ■ç;ŏ (ézÿèŏd' äÿž0) âĹžæŕÿäÿ■èúĚæ
--swap-with-value value
→èŊĚèrěéTŏçŎŕăĀIJĹçŽDăĀijăŸŕvalueiijŊăĹžèĹžèaŊĚŏç;ŏæš■ă;IJ
--swap-with-index '0'
→èŊĚèrěéTŏçŎŕăĀIJĹçŽDçt' căijŤăĀijăŸŕæŊĜăŏŽçt' căijŤiijŊăĹžèĹžèaŊĚŏç;ŏæš■ă;IJ
```

- mk

æçCădIJçzŽăŏŽçŽDěTŏäÿ■ă■ŸăIJĹijŊăĹŽăĹŽăzžäÿĂäÿĤæŮŕçŽDěTŏăĀijăĂCă;ŊăçC:

```
$ etcdctl mk /testdir/testkey "Hello world"
Hello world
```

ă;šéTŏă■ŸăIJĹçŽDăŮúăĂŽiijŊăĹ'ĝèaŊĚŕéăš;ăzd'ăijŽæĹéTŽiijŊă;ŊăçC:

```
$ etcdctl mk /testdir/testkey "Hello world"
Error: 105: Key already exists (/testdir/testkey) [8]
```

æŤræŊAçŽDěĂL'éązăÿž:

```
--ttl '0' èúĚæŮúæŮúéŮt' (ă■Tă;■ăÿžçĝš iijL' iijŊăÿ■éĚ■ç;ŏ (ézÿèŏd' äÿž
→0) âĂCăĹžæŕÿäÿ■èúĚæŮú
```

- mkdir

æçCădIJçzŽăŏŽçŽDěTŏçŽŏă;Ťäÿ■ă■ŸăIJĹijŊăĹŽăĹŽăzžäÿĂäÿĤæŮŕçŽDěTŏçŽŏă;ŤăĂCă;ŊăçCiiijŽ

```
$ etcdctl mkdir testdir2
```

ă;šéTŏçŽŏă;Ťă■ŸăIJĹçŽDăŮúăĂŽiijŊăĹ'ĝèaŊĚŕéăš;ăzd'ăijŽæĹéTŽiijŊă;ŊăçCiiijŽ

```
$ etcdctl mkdir testdir2
Error: 105: Key already exists (/testdir2) [9]
```

æŤræŊAçŽDěĂL'éązăÿžiiijŽ


```
--ttl '0'
→èúĚæŮúæŮúéŮt' (â■Ṫä;■äÿžçğš) iijNäÿ■éĚ■ç;ó(ézÿèód' äÿž0) âĹZæřÿäÿ■èúĚæŮúāĀĆ
```

- setdir

âĹZâzzäÿÄäÿléŤóçZōâ;ṪāĀĆæĈCæđIJçZōâ;Ṫäÿ■â■ÿâIJlârşâĹZâzziiijNäçCæđIJçZōâ;Ṫâ■ÿâIJlæŽt' æŮŮ

```
$ etcdctl setdir testdir3
```

æŤræNÄçŽĐéĀĹ'éąäÿž:

```
--ttl '0'
→èúĚæŮúæŮúéŮt' (â■Ṫä;■äÿžçğš) iijNäÿ■éĚ■ç;ó(ézÿèód' äÿž0) âĹZæřÿäÿ■èúĚæŮúāĀĆ
```

23.5.2 23.5.2 âĹäéŽd'

- rm

âĹäéŽd' æšRäÿléŤōâĀijāĀĆäĹNäçĆ:

```
$ etcdctl rm /testdir/testkey
PrevNode.Value: Hello
```

â;šĚŤōäÿ■â■ÿâIJlæŮŮiiijNâĹZâijŽæĹéŤŽāĀĆäĹNäçĆ:

```
$ etcdctl rm /testdir/testkey
Error: 100: Key not found (/testdir/testkey) [7]
```

æŤræNÄçŽĐéĀĹ'éąäÿžiiijŽ

```
--dir æĈCæđIJéŤōæŸrăÿłçl' žçZōâ;ṪæĹŮèĀĚéŤōâĀijârżâĹZâĹäéŽd'
--recursive âĹäéŽd' çZōâ;ṪâšNæL' ÄæIJL' â■ŘéŤō
--with-value æĈĀæšĚçŎřæIJL' çŽĐâĀijæŸrăŘęâŇzéĚ■
--with-index '0' æĈĀæšĚçŎřæIJL' çŽĐindexæŸrăŘęâŇzéĚ■
```

- rmdir

âĹäéŽd' äÿÄäÿłçl' žçZōâ;ṪiiijNæĹŮèĀĚéŤōâĀijârżāĀĆ

```
$ etcdctl setdir dir1
$ etcdctl rmdir dir1
```

èNěçZōâ;Ṫäÿ■çl' žiiijNâijŽæĹéŤŽ:

```
$ etcdctl set /dir/testkey hi
hi
$ etcdctl rmdir /dir
Error: 108: Directory not empty (/dir) [17]
```

23.5.3 23.5.3 æŽt'æŰř

- update

ǎĭŞéŤōā■ŶǎIJǎŰŭiijŇǎŽt'æŰřǎĀijǎĚĚǎōzǎǎĆǎĭŇǎĕĆriijŽ

```
$ etcdctl update /testdir/testkey "Hello"
Hello
```

ǎĭŞéŤōäy■ǎ■ŶǎIJǎŰŭiijŇǎĹŽǎijŽǎĚĚŤŽǎǎĆǎĭŇǎĕĆ:

```
$ etcdctl update /testdir/testkey2 "Hello"
Error: 100: Key not found (/testdir/testkey2) [6]
```

ǎŤřǎŇǎĕŽĎĚǎĹ'ĕǎžǎyž:

```
--ttl '0'  ěŭĚǎŰŭǎŰŭéŰt' (ǎ■Ťǎ;■ǎyžċğš) ĩijŇǎy■ĕĚ■ċ;ǎ (ĕžŶĕǎd' äyž_
→0) ǎĹŽǎřyǎy■ĕŭĚǎŰŭǎǎĆ
```

- updatedir

ǎŽt'æŰřǎyǎǎyǎŭšċzŤǎ■ŶǎIJǎŽĎċŽǎĭŤǎǎĆ

```
$ etcdctl updatedir testdir2
```

ǎŤřǎŇǎĕŽĎĚǎĹ'ĕǎžǎyž:

```
--ttl '0' _
→ĕŭĚǎŰŭǎŰŭéŰt' (ǎ■Ťǎ;■ǎyžċğš) ĩijŇǎy■ĕĚ■ċ;ǎ (ĕžŶĕǎd' äyž0) ǎĹŽǎřyǎy■ĕŭĚǎŰŭǎǎĆ
```

23.5.4 23.5.4 æŞĕĕĕř

- get

ĕŎŭǎŤŰǎŇǎĠǎōŽĕŤōċŽĎǎĀijǎǎĆǎĭŇǎĕĆriijŽ

```
$ etcdctl get /testdir/testkey
Hello world
```

ǎĭŞéŤōäy■ǎ■ŶǎIJǎŰŭiijŇǎĹŽǎijŽǎĚĚŤŽǎǎĆǎĭŇǎĕĆriijŽ

```
$ etcdctl get /testdir/testkey2
Error: 100: Key not found (/testdir/testkey2) [5]
```

ǎŤřǎŇǎĕŽĎĚǎĹ'ĕǎžǎyž:

```
--sort ǎřžċžŞǎdIJĕĹŽĕǎŇǎŎŖǎžŤ
--consistent_
→ǎřĚĕŕŭǎšĆǎŤŤċžŽǎyžĕĹĊċĊziiijŇǎĹĚĕŕǎĕŎŭǎŤŰǎĚĚǎōzċžĎǎyǎĕĠt' ǎǎġǎǎĆ
```

- ls

ä¿ÑăĕĆïjŽ

```
$ etcdctl ls
/testdir
/testdir2
/dir
```

æŦræŦçŽĐéĂL'éązãŦĚæŦň:

```
--sort řřĚčŠăĜžčžŞæđIǰăŎŠăžŘ
--recursive┐
↳ăĖĈæđIǰčžŎă;ǰăŸNăIǰL'ă■ŘčžŎă;ǰiǰjŇăIǰžĚăŠă;ŞĚčŞăĜžăĚŸăŸ■čžĐăĚĚăŎž
-p řřžžžŎĚčŞăĜžăŸžčžŎă;ǰiǰjŇăIǰIǰăIǰĴăŎŮžăĴă/ĚĚžĚăŇăŇžăĴĚ
```

23.5.5 23.5.5 watch

- watch

çZŠætŊäyÄäyleŤōāĀijçZĐāRŸāNŮiijŊäyÄæUëēŤōāĀijaŔŚçŤŧæZŧ æŮriijŊāŕšäijŽèŁŞāGžæIJæŮŕçZŧ
äŁŊāęĆ:çŤŭæŁūāZŧ æŮŕtestkeyēŤōāĀijäyžHello watchāĂĆ

```
$ etcdctl get /testdir/testkey
Hello world
$ etcdctl set /testdir/testkey "Hello watch"
Hello watch
$ etcdctl watch testdir/testkey
Hello watch
```

æT̃ræÑAçŽD'éǻzǻÑĚæNň:

```
--forever  äÿÄçŽt' çŽŚætŃçŽt' äĹŕçĹĹæĹuæŃĹ' CTRL+CéÄÄgŽ
--after-index '0' äĹĹæŃgĹăŹindexäžŃäĹ' ■äÿÄçŽt' çŽŚætŃ
--recursive èĹŤăžđæĹ' ÄæĹĹ' çŽĐéŤŉäĹi jašŃä■RéŤŉäĹi j
```

- exec-watch

çZŚætNäyÄäyleTōāĀijçZDāRŸāNŪijNäyÄæUēéTōāĀijāRŚçTſæZt æŪrijNārsæL'gēaŊçzZāōZāŚjāzç
äçNāçĀijçZçTlæLūæZt æŪrtestkeyéTōāĀijāĀĆ

```
$ etcdctl exec-watch testdir/testkey -- sh -c 'ls'
config Documentation etcd etcdctl README-etcdctl.md README.md
↪ READMEv2-etcdctl.md
```

æŦræŦAçŽĐéĀL'éazăŦĚæŦň:

```
--after-index '0' âIJIæŃĠăŏŽ index äzŃâL'■äÿĂçŽt'çŽŚætŃ
--recursive èfTâžđæL'ĂæIJL'çŽĐéŤŏăĂijăŠŃă■ŘéŤŏăĂij
```

23.5.6 23.5.6 âd'Ġăž;

âd'Ġăž;etcdçŽĐæŤræ■ŏăĂĆ

```
$ etcdctl backup --data-dir /var/lib/etcd --backup-dir /home/etcd_
↳ backup
```

æŤræŃAçŽĐéĂL'éąžăŃĚæŃň:

```
--data-dir etcdçŽĐæŤræ■ŏçŽŏă;Ť
--backup-dir âd'Ġăž;ăĹræŃĠăŏŽèŭrăĹĐ
```

23.5.7 23.5.7 member

éĂŽèfĠlistăĂAaddăĂAremoveăŚ;ăžd'ăĹŮăĠžăĂAæŭzăĹăăĂAăĹăéŽd'etcdăŏđă;ŃăĹretcdéŽEç;
æšëçIJŃéŽEç;d'äÿ■ă■ŸăIJçŽĐèĹĆçĆž

```
$ etcdctl member list
8e9e05c52164694d: name=dev-master-01 peerURLs=http://localhost:2380
↳ clientURLs=http://localhost:2379 isLeader=true
```

ăĹăéŽd'éŽEç;d'äÿ■ă■ŸăIJçŽĐèĹĆçĆž

```
$ etcdctl member remove 8e9e05c52164694d
Removed member 8e9e05c52164694d from cluster
```

ăŘŚéŽEç;d'äÿ■æŮřăĹăèĹĆçĆž

```
$ etcdctl member add etcd3 http://192.168.1.100:2380
Added member named etcd3 with ID 8e9e05c52164694d to cluster
```

23.6 23.6 çd'žă;Ń

```
# èŏç;ŏäÿĂäÿİkeyăĂij
[root@etcd-0-8 ~]# etcdctl set /msg "hello k8s"
hello k8s

# èŮăăŤŮkeyçŽĐăĂij
[root@etcd-0-8 ~]# etcdctl get /msg
hello k8s
```

(continues on next page)

(continued from previous page)

```

# èŒûâŔŨkeyâĀijçŽĎëręçzĒäŁæĀŕ
[root@etcd-0-8 ~]# etcdctl -o extended get /msg
Key: /msg
Created-Index: 12
Modified-Index: 12
TTL: 0
Index: 12

hello k8s

# èŒûâŔŨäÿ■â■ŸâIJÍçŽĎkeyâŽđæŁééŦŽ
[root@etcd-0-8 ~]# etcdctl get /xxxz
Error: 100: Key not found (/xxxz) [12]

# èŒçç;ŒkeyçŽĎttlīijŇèŁĜæIJSâŔŒăijŽèćnèĜłâŁÍâŁăéŽd'
[root@etcd-0-8 ~]# etcdctl set /testkey "tmp key test" --ttl 5
tmp key test
[root@etcd-0-8 ~]# etcdctl get /testkey
Error: 100: Key not found (/testkey) [14]

# key æŽŁæ■ćæ$■ă;IJ
[root@etcd-0-8 ~]# etcdctl get /msg
hello k8s
[root@etcd-0-8 ~]# etcdctl set --swap-with-value "hello k8s" /msg
↪ "goodbye"
goodbye
[root@etcd-0-8 ~]# etcdctl get /msg
goodbye

# mk äžĚâ;Şkeyäÿ■â■ŸâIJÍæŨŮâŁŽâžž(setâržâŔŇäÿĀäÿłkeyäijŽèęEçŽŨ)
[root@etcd-0-8 ~]# etcdctl get /msg
goodbye
[root@etcd-0-8 ~]# etcdctl mk /msg "mktest"
Error: 105: Key already exists (/msg) [18]
[root@etcd-0-8 ~]# etcdctl mk /msg1 "mktest"
mktest

# âŁŽâžžèĜłæŒŖăžŔçŽĎkey
[root@etcd-0-8 ~]# etcdctl mk --in-order /queue s1
s1
[root@etcd-0-8 ~]# etcdctl mk --in-order /queue s2
s2
[root@etcd-0-8 ~]# etcdctl ls --sort /queue
/queue/000000000000000000021
/queue/000000000000000000022
[root@etcd-0-8 ~]# etcdctl get /queue/000000000000000000021
s1

```

(continues on next page)

(continued from previous page)

```

# æŽt'æŮřkeyâĀij
[root@etcd-0-8 ~]# etcdctl update /msg1 "update test"
update test
[root@etcd-0-8 ~]# etcdctl get /msg1
update test

# æŽt'æŮřkeyčŽĎttlâĀĀĀij
[root@etcd-0-8 ~]# etcdctl update --ttl 5 /msg "aaa"
aaa

# âĹžâžžčŽŮâ;Ŧ
[root@etcd-0-8 ~]# etcdctl mkdir /testdir

# âĹăéžd'çl'žčŽŮâ;Ŧ
[root@etcd-0-8 ~]# etcdctl mkdir /test1
[root@etcd-0-8 ~]# etcdctl rmdir /test1

# âĹăéžd'éïdçl'žčŽŮâ;Ŧ
[root@etcd-0-8 ~]# etcdctl get /testdir
/testdir: is a directory
[root@etcd-0-8 ~]#
[root@etcd-0-8 ~]# etcdctl rm --recursive /testdir

# âĹŮâĜžčŽŮâ;ŦâĒĒâŮž
[root@etcd-0-8 ~]# etcdctl ls /
/tmp
/msg1
/queue
[root@etcd-0-8 ~]# etcdctl ls /tmp
/tmp/a
/tmp/b

# éĂšâ;šâĹŮâĜžčŽŮâ;ŦčŽĎâĒĒâŮž
[root@etcd-0-8 ~]# etcdctl ls --recursive /
/msg1
/queue
/queue/000000000000000000021
/queue/000000000000000000022
/tmp
/tmp/b
/tmp/a

# çŽšâĤňkeyiijŇâ;škeyâĤšçŦšæŦžâĤŦçŽĎâŮâĀžæL'šâ■řâĜžâĤŦâŇŮ
[root@etcd-0-8 ~]# etcdctl watch /msg1
xxx

[root@VM_0_17_centos ~]# etcdctl update /msg1 "xxx"
xxx

```

(continues on next page)

(continued from previous page)

```

# çŻŚăŘňăŞŘăÿiçŻôă;ŦiijŇă;ŞçŻôă;Ŧăÿ■ăżżă;Ŧ node_
→æŦzăŘŸçŻĎăŮăăĂŹiijŇéČ;ăijžæL'Şă■řăĜžæİě
[root@etcd-0-8 ~]# etcdctl watch --recursive /
[update] /msg1
xxx

[root@VM_0_17_centos ~]# etcdctl update /msg1 "xxx"
xxx

# äÿĂçŽt'çŻŚăŘňiijŇéžd'éİđ `CTL + C` árijèĜt'éĂĂăĜžçŻŚăŘň
[root@etcd-0-8 ~]# etcdctl watch --forever /

# çŻŚăŘňçŻôă;ŦiijŇă;ŞăŘŚçŦŞăŘŸăŇŮăŮăæL'ğèaŇăÿĂæİaăŞ;ăzd'
[root@etcd-0-8 ~]# etcdctl exec-watch --recursive / -- sh -c "echo_
→change"
change

# backup
[root@etcd-0-14 ~]# etcdctl backup --data-dir /data/app/etcd --
→backup-dir /root/etcd_backup
2019-12-04 10:25:16.113237 I | ignoring EntryConfChange raft entry
2019-12-04 10:25:16.113268 I | ignoring EntryConfChange raft entry
2019-12-04 10:25:16.113272 I | ignoring EntryConfChange raft entry
2019-12-04 10:25:16.113293 I | ignoring member attribute update on /
→0/members/2d2e457c6a1a76cb/attributes
2019-12-04 10:25:16.113299 I | ignoring member attribute update on /
→0/members/d2d2e9fc758e6790/attributes
2019-12-04 10:25:16.113305 I | ignoring member attribute update on /
→0/members/56e0b6dad4c53d42/attributes
2019-12-04 10:25:16.113310 I | ignoring member attribute update on /
→0/members/56e0b6dad4c53d42/attributes
2019-12-04 10:25:16.113314 I | ignoring member attribute update on /
→0/members/2d2e457c6a1a76cb/attributes
2019-12-04 10:25:16.113319 I | ignoring member attribute update on /
→0/members/d2d2e9fc758e6790/attributes
2019-12-04 10:25:16.113384 I | ignoring member attribute update on /
→0/members/56e0b6dad4c53d42/attributes

# ä;£çŦŦlv3çL'ŦæIŦň
[root@etcd-0-14 ~]# export ETCDCTL_API=3
[root@etcd-0-14 ~]# etcdctl --endpoints="http://172.16.0.8:2379,
→http://172.16.0.14:2379,http://172.16.0.17:2379" snapshot save_
→mysnapshot.db
Snapshot saved at mysnapshot.db
[root@etcd-0-14 ~]# etcdctl snapshot status mysnapshot.db -w json
{"hash":928285884,"revision":0,"totalKey":5,"totalSize":20480}

```

- áŦŦçĚĂčŚ;æŦě

- <https://github.com/etcd-io/etcd>
- <https://www.yuque.com/lurunhao/nl81zh/hb8sie>
- <https://www.hi-linux.com/posts/40915.html>
- <https://cizixs.com/2016/08/02/intro-to-etcd/>
- Etcd Raftä;ŁçŤlăĚěéŮlăŔŁăŦŝçŔĚèğčæđŔ
- <https://juejin.im/post/5dabc50ef265da5b591b761a>
- <https://www.infoq.cn/article/coreos-analyse-etcd/>


```
#!/bin/bash
# auth:kaliarch

clear_evicted_pod() {
    ns=$1
    kubectl delete pods -n ${ns} $(kubectl get pods -n ${ns} | grep
    ↪Evicted |awk '{print $1}')
}

clear_crash_pod() {
    ns=$1
    kubectl delete pods -n ${ns} $(kubectl get pods -n ${ns} | grep
    ↪CrashLoopBackOff |awk '{print $1}')
}

clear_exited_container() {
    docker rm `docker ps -a | grep Exited |awk '{print $1}`
}

echo "1.clear existed pod"
echo "2.clear crash pod"
echo "3.clear exited container"
read -p "Please input num:" num

case ${num} in
"1")
    read -p "Please input oper namespace:" ns
    clear_evicted_pod ${ns}
    ;;

"2")
    read -p "Please input oper namespace:" ns
    clear_crash_pod ${ns}
    ;;

"3")
    clear_exited_container
    ;;

"*)
    echo "input error"
    ;;
esac
```

- æÿĖċŘĚĚl'séĀŘçŁŭæĀĀçŽĎpod

```
# èŮâŔŮæL'ĀæIJL'ns
kubectl get ns | grep -v "NAME" |awk '{print $1}'

# æÿĖċŘĚĚl'séĀŘçŁŭæĀĀçŽĎpod
for ns in `kubectl get ns | grep -v "NAME" | awk '{print $1}`;do
    ↪kubectl delete pods -n ${ns} $(kubectl get pods -n ${ns} | grep
    ↪"Evicted" |awk '{print $1}');done
```

(continued from previous page)

```
# æÿĖÇŘĚâijĈâÿÿpod
for ns in `kubectl get ns | grep -v "NAME" | awk '{print $1}'`;do
  ↪kubectl delete pods -n ${ns} $(kubectl get pods -n ${ns} | grep
  ↪"CrashLoopBackOff" |awk '{print $1}');done
```

24.3 24.3 Docker æŢræ■ōè£Açğž

âIJlâôL'èçĖè£ĜçlNäÿ■æIJlæNĜâôŽdockeraŢræ■ōçZôâjŢiijNçşzçzşçZŸ50GiiNéŽŖçIĂæUúéŮ' æŎlç
éçŮéĂL æNĈèj;æŮřçĈAçZŸâĹ/dataçZôâjŢ

```
systemctl stop docker

mkdir -p /data/docker/

rsync -avz /var/lib/docker/ /data/docker/

mv /var/lib/docker /data/docker_bak

ln -s /data/docker /var/lib/

systemctl daemon-reload

systemctl start docker
```

24.4 24.4 kubescape çjŞçzIjæŎŠéŢŽ

- éŮŏéçŸæŖŖè£řijŽ

âIJlkubescapeçŽDnodeèLĈçĈzæĹŮmasterèLĈçĈzřijNæL'NâĹlâŎžâŖřâĹlâôžâŽiijNâIJlâôžâŽléĜNéic
æşççIJNæL'NâĹlâŖřâĹlçŽDâôžâŽlç;ŞçzIJäÿLèřçŽDdocker0

```
root@fd1b8101475d:/# ip a

1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
  ↪group default qlen 1

    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00

    inet 127.0.0.1/8 scope host lo

        valid_lft forever preferred_lft forever

2: tunl0@NONE: <NOARP> mtu 1480 qdisc noop state DOWN group default
  ↪qlen 1
```

(continues on next page)

(continued from previous page)

```

link/ipip 0.0.0.0 brd 0.0.0.0

105: eth0@if106: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc_
↪noqueue state UP group default

link/ether 02:42:ac:11:00:02 brd ff:ff:ff:ff:ff:ff link-netnsid_
↪0

inet 172.17.0.2/16 brd 172.17.255.255 scope global eth0

valid_lft forever preferred_lft forever

```

åĲĲpodsäy■çŽĐāōžāŽĲç;ŚçžĲçŤĲçŽĐæŸřkubepvs0

```

1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue qlen 1

link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00

inet 127.0.0.1/8 scope host lo

valid_lft forever preferred_lft forever

2: tunl0@NONE: <NOARP> mtu 1480 qdisc noop qlen 1

link/ipip 0.0.0.0 brd 0.0.0.0

4: eth0@if18: <BROADCAST,MULTICAST,UP,LOWER_UP,M-DOWN> mtu 1500_
↪qdisc noqueue

link/ether c2:27:44:13:df:5d brd ff:ff:ff:ff:ff:ff

inet 10.233.97.175/32 scope global eth0

valid_lft forever preferred_lft forever

```

• èğçâĒşæŮžæāĲijŽ

æšěçĲĲŃdockerāŖřāĲĲĲē■ç;ř

```

[root@node01 smartns-k8s-dev]# ps -ef |grep docker
root      592      1  4 Oct11 ?        02:41:27 /usr/bin/dockerd --registry-mirror=https://registry.docker-cn.com --data-root=/var/lib/docker --log-opt max-
-size=10m --log-opt max-file=3 --iptables=false --insecure-registry=harbor.devops.kubesphere.local:30280 --dns 10.17.50.3 --dns 1.2.4.8 --dns-search defau
t.svc.cluster.local --dns-search svc.cluster.local --dns-opt ndots:2 --dns-opt timeout:2 --dns-opt attempts:2
root      644      0 Oct11 ?        00:13:02 docker-containerd --config /var/run/docker/containerd/containerd.toml

```

ăĲōăŤžæŮĴăžŮ/etc/systemd/system/docker.service.d/docker-
options.confăy■ăŮžæŮĲăŖĲăŤřijŽ--iptables=false èĲŽăylăŖĲăŤřç■ĲăžŮfalseæŮŮăijŽăy■ăĒZiptables

```

[Service]
Environment="DOCKER_OPTS= --registry-mirror=https://registry.
↪docker-cn.com --data-root=/var/lib/docker --log-opt max-size=10m -
↪-log-opt max-file=3 --insecure-registry=harbor.devops.kubesphere.
↪local:30280"

```

(continues on next page)

(continued from previous page)

24.5 24.5 kubescape ažTçTlèùrçTśaijCăyŷ

âIJlkubescapeäy■āžTçTlèùrçTśingressä;ŁçTlçŽDæŸrngxijNâIJlwebçTŇéÍcéĚ■ç;őäijŽārijèĜt'äyd'äy
âŽāiyRæşlæĐRijŽingressæŌğāLúdeploymentâIJliijŽ



```
kind: Ingress
apiVersion: extensions/v1beta1
metadata:
  name: prod-app-ingress
  namespace: prod-net-route
  resourceVersion: '8631859'
  labels:
    app: prod-app-ingress
  annotations:
    desc: çTŚäžğçŌrācČāžTçTlèùrçTś
    nginx.ingress.kubernetes.io/client-body-buffer-size: 1024m
    nginx.ingress.kubernetes.io/proxy-body-size: 2048m
    nginx.ingress.kubernetes.io/proxy-read-timeout: '3600'
    nginx.ingress.kubernetes.io/proxy-send-timeout: '1800'
    nginx.ingress.kubernetes.io/service-upstream: 'true'
spec:
  tls:
    - hosts:
        - smartms.tools.anchnet.com
      secretName: smartms-ca
    - hosts:
        - smartmsd.tools.anchnet.com
      secretName: smartmsd-ca
```

(continues on next page)

(continued from previous page)

```

rules:
- host: smartms.tools.anchnet.com
  http:
    paths:
      - path: /
        backend:
          serviceName: smartms-frontend-svc
          servicePort: 80
- host: smartds.tools.anchnet.com
  http:
    paths:
      - path: /
        backend:
          serviceName: smartds-frontend-svc
          servicePort: 80

```

24.6 Jenkins 安装 Agent

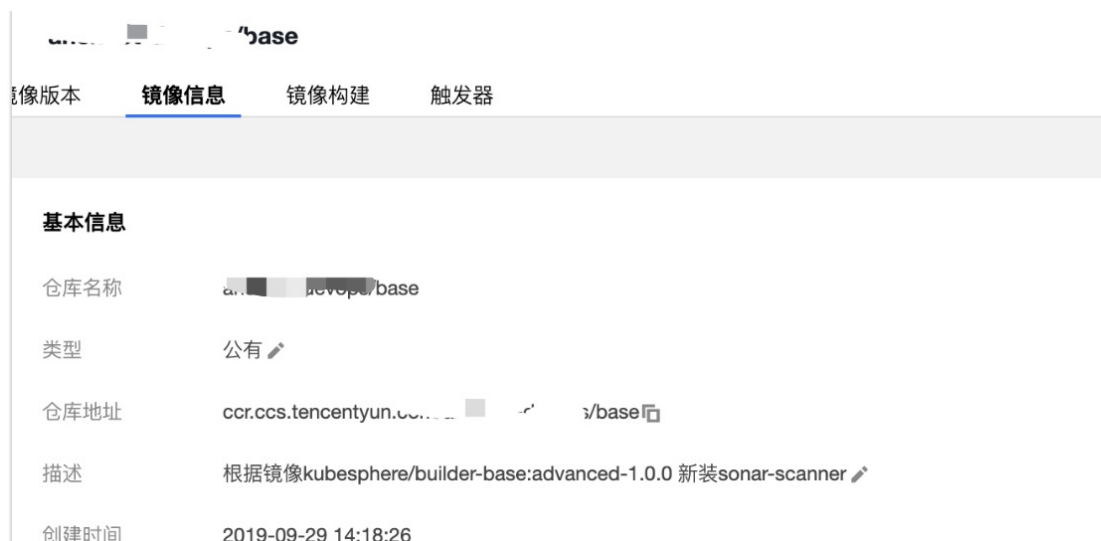
安装 Jenkins Agent 的步骤如下：

1. 安装 Jenkins 的 base 镜像，使用命令：
`docker pull jenkins/jenkins-agent`

2. 创建 Jenkins Agent 的配置文件，使用命令：
`cat > /etc/docker/daemon.json`

3. 启动 Jenkins Agent，使用命令：
`docker run -d --name jenkins-agent --restart=always --privileged --network=host jenkins/jenkins-agent`

4. 安装 Jenkins 的插件，使用命令：
`docker exec -it jenkins-agent /bin/bash`
`curl -sL https://kubernetes.io/docs/advanced-v2.0/zh-CN/devops/devops-admin-faq/#%E5%8D%87%E7%BA%A7-jenkins-agent-%E7%9A%84%E5%8C%85%E7%89%88%E6%9C%AC`



Kubernetes Pod template

名称

base

命名空间

kubesphere-devops-system

标签列表

base

用法

尽可能的使用这个节点

父级的 Pod 模板名称

容器列表

Container Template

名称

base

Docker 镜像

ccr.ccs.tencentyun.com/k8sdevops/base:v1

总是拉取镜像

☐

工作目录

/home/jenkins

运行的命令

cat

命令参数

cat

分配伪终端

☒

EnvVars

添加环境变量

设置到 Pod 节点中的环境变量列表

高级...

BUILD_NUMBER	Current build number, such as <code>AIJ153</code>
BUILD_ID	The current build ID, identical to <code>BUILD_NUMBER</code> for builds created in 1.597+, but a <code>YYYY-MM-DD_hh-mm-ss</code> timestamp for older builds
BUILD_DISPLAY_NAME	Display name of the current build, which is something like <code>AIJ#153</code> by default.
JOB_NAME	Name of the project of this build, such as <code>AIJfoo</code> or <code>AIJ-foo/bar</code> . (To strip off folder paths from a Bourne shell script, try: <code>JOB_NAME</code>) <code> BUILD_TAG Stringof"jenkins-{JOB_NAME }-\${BUILD_NUMBER}</code> . Convenient to put into a resource file, a jar file, etc for easier identification.
EX-ECUTOR_NUMBER	The unique number that identifies the current executor (among executors of the same machine) that's carrying out this build. This is the number you see in the build executor status, except that the number starts from 0, not 1.
NODE_NAME	Name of the slave if the build is on a slave, or <code>AIJmaster</code> if run on master
NODE_LABELS	Space-separated list of labels that the node is assigned.
WORKSPACE	Absolute path of the directory assigned to the build as a workspace.
JENKINS_HOME	The absolute path of the directory assigned on the master node for Jenkins to store meta.
JENKINS_URL	Full URL of Jenkins, like <code>http://server:port/jenkins/</code> (note: only available if Jenkins URL set in system configuration)
BUILD_URL	Full URL of this build, like <code>http://server:port/jenkins/job/foo/15/</code> (Jenkins URL must be set)
SVN_REVISION	Revision number that's currently checked out to the workspace, such as <code>AIJ12345</code>
SVN_URL	Subversion URL that's currently checked out to the workspace.
JOB_URL	Full URL of this job, like <code>http://server:port/jenkins/job/foo/</code> (Jenkins URL must be set)

æIJÄçZLëGłauśaEŻäZĖĖÄCăZTëGłauśäyŽăŁaçŽDălaçL'LiiJŃăRăzëçŽt' æÖëä;ŁçTl

```
mail to: 'xuel@net.com',
        charset:'UTF-8', // or GBK/GB18030
        mimeType:'text/plain', // or text/html
        subject: "Kubesphere ${env.JOB_NAME} [${env.BUILD_NUMBER}]
→] âŖŠăŸČă■čăŸŸRunning Pipeline: ${currentBuild.fullDisplayName}",
        body: ""
        -----Anchnet Devops Kubesphere Pipeline job-----
→-----

        éazçŽôăŖ■çğř : ${env.JOB_NAME}
        æďĐăžžæňæŕ : ${env.BUILD_NUMBER}
        æL'ńăŖŖăŁaæAŕ : âIJŕăİĂ:${SONAR_HOST}
        éŤIJăČŖăIJŕăİĂ : ${REGISTRY}/${QHUB_NAMESPACE}/${APP_NAME}
→:${IMAGE_TAG}
        æďĐăžžèŕçæČĖiiJŽSUCCESSFUL: Job ${env.JOB_NAME} [${env.
→BUILD_NUMBER}]
```

(continues on next page)

(continued from previous page)

```
ædǾāzžçŁũæĀ : ${env.JOB_NAME} jenkins åŖŚåÿČèŁŘèaÑæ■čāÿÿ
ædǾāzžURL : ${env.BUILD_URL} " " "
```



-----Anchnet Devops Kubernetes Pipeline job-----

项目名称: project-wG97Bo24LjEr/smartsns-offline/k8s-dev
 构建次数: 3
 扫描信息: 地址: http://...
 镜像地址: ccr.ccs.t.../smartsns:v1
 构建详情: SUCCESSFUL: Job project-wG97Bo24LjEr/smartsns-offline/k8s-dev [3]
 构建状态: project-wG97Bo24LjEr/smartsns-offline/k8s-dev jenkins 发布运行正常
 构建URL: http://jenkins.devops.kubernetes.local:30080/job/project-wG97Bo24LjEr/job/smartsns-offline/job/k8s-dev/3/

24.8 åŖĆèĀČéŞŁæŮě

- KubernetesåŖŸç;ŚæŤŻćłŃ
- Kubernetesäÿ■æŮĜçd'ŁāŃž
- äžŮKubernetesåŖŁŕCloud Native
- Kubernetes Handbook
- KubernetesäžŮåĚěéŮíåŖŁŕåŖđæŁŸ
- KubernetesæŃĜå■Ů
- awesome-kubernetes
- äžŮDockeråŖŁŕKubernetesèŁŻéŸŭ
- pythonåŖŁŕæIJ■åŖŁŕåŖđæŁŸ
- äžŚåŮŞçŤŖšäžŃèŭŕ
- CNCF Cloud Native Interactive Landscape

24.8.1 èġĖćŚ

- éĴňăŞě(dockeraőźăŹíæĹĂæIJř+k8séŹĖç;d'æĹĂæIJř)
- âĴőæIJ■âĹăăőźăŹíăŇŮăőđæĹŸ

ăĖĆăđIJă■đ'çňŤëőřăřăæĆíæIJĹ'ăzzăĴăŸăăĹĴ'ĴĴŇăŹĴ'ăđ'ŹăŮĜçňăĴĴŇăĵćèĴŌăĚşăşĴă■ŹăőcăŸĂăĴŮă

24.8.2 èŕŭæĹŚăŮîăŠŮăŢăâŸŢĴŸŔ



- âĴőăĴă



- æŢřăžŸăőĴ

CHAPTER 25

Indices and tables

- `genindex`
- `modindex`
- `search`