ALIBABA CLOUD

Alibaba Cloud

Container Service for Kubernetes

Release notes

Document Version: 20220511

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	**Danger:** Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	**Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	**Notice:** If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	**Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.Release Notes

This topic describes the release notes for Container Service for Kubernetes (ACK) and provides links to the relevant references.

Background information

- The following Kubernetes versions are supported by ACK: 1.22.3, 1.20.11, and 1.18.8.
- The following operating systems are supported by ACK: CentOS 7.9, Alibaba Cloud Linux 3.2104, Alibaba Cloud Linux 2.1903, Windows Server 2019, and Windows Server Core 1909.

February 2022

Feature	Description	Region	References
RHEL7.9 supported by worker nodes	RHEL7.9 can be selected when you select a custom OS image for worker nodes.	All regions	Create a Kubernetes cluster by using a custom image
Support for multiple security groups in a cluster that uses Terway	In a cluster that uses the Terway network plug-in, Terway uses elastic network interfaces (ENIs) to assign IP addresses to containers. You can associate at most five security groups with an ENI that is created by Terway. This helps you regulate access control on pods in a flexible manner.	All regions	Associate multiple security groups with an ENI
Custom configurations for CoreDNS	Custom configurations are supported for CoreDNS. You can customize the configurations of the CoreDNS component on the Add-ons page of the ACK console. The custom configurations are retained after you update the CoreDNS version. **Note** You can customize configurations only for CoreDNS versions later than 1.8.4.2.	All regions	Manage system components

Feature	Description	Region	References
ExternalDNS available in App Catalog	You can use ExternalDNS to configure external DNS servers for Ingresses and Services in your ACK clusters. This allows you to use public DNS servers to discover Kubernetes resources in your clusters. ExternalDNS works in a similar way as kube-dns. ExternalDNS retrieves information about Services and Ingresses from the Kubernetes API server and creates DNS records.	All regions	Use ExternalDNS to configure external DNS servers
gRPC supported by ALB Ingresses	The gRPC protocol is supported by Application Load Balancer (ALB) Ingresses. To use gRPC, add the annotation `alb.ingress.kubernetes.io/backend-protocol: "grpc"`.	All regions	None
ALB Ingresses supported by ASK Knative	Knative is an open source, serverless application framework. Knative can help you deploy serverless workloads to Kubernetes and manage these workloads. ALB is a load balancing service intended for applications that use the HTTP, HTTPS, or Quick UDP Internet Connection (QUIC) protocol. ALB is highly scalable and can distribute large amounts of network traffic at Layer 7. Serverless Kubernetes (ASK) allows you to deploy ALB Ingresses based on Knative.	All regions	None
Authorization management, node pools, and template management optimized	The ACK console is optimized to improve user experience: • You can select multiple namespaces when you manage authorizations. You can grant permissions on multiple namespaces at a time. • A filter is added to the node management module. You can use this filter to show only unschedulable nodes. This improves the O&M efficiency. • You can add variables to templates when you use the template editor to manage templates. This improves the efficiency of template development.	All regions	• Assign RBAC roles to RAM users or RAM roles • Manage node labels • Create an orchestration template

January 2022

Feature	Description	Region	References
Supplementation of preemptible instances supported by node pools	The supplementation of preemptible instances is supported by the cost optimization scaling policy of node pools. After the supplementation of preemptible instances is enabled, the system attempts to add a new preemptible instance to the scaling group 5 minutes before a preemptible instance is reclaimed from the scaling group.	All regions	Manage node pools
Custom images selectable during node pool creation	Custom images contain custom scripts and optimized parameters and can be used to deploy the operating systems of worker nodes. You can directly use custom images without the need to be added to a whitelist. You must use custom images that are based on Alibaba Cloud Linux 2.1903 or CentOS 7.9.	All regions	Manage node pools
New region	ACK managed and ACK dedicated clusters are available in the China North 2 Finance region.	China East 2 Finance	Supported regions
KMS keys that have automatic rotation enabled supported by Secret encryption	Keys that are created in KMS can be used to encrypt Secrets in ACK Pro clusters. ACK allows you to use KMS keys that have automatic rotation enabled to encrypt Secrets in your clusters. When a key is being automatically rotated, the old key version is used to encrypt the existing Secrets in your cluster.	All regions	Use KMS to encrypt Kubernetes Secrets
Resource priority-based scheduling supported by Cybernetes	Priority-based resource scheduling is provided by Alibaba Cloud to meet elasticity requirements in pod scheduling. A ResourcePolicy specifies the priorities of nodes in descending order for pod scheduling. When the system deploys or scales out pods for an application, pods are scheduled to nodes based on the priorities of the nodes that are listed in the ResourcePolicy. When the system scales in pods for an application, pods are removed from nodes based on the priorities of the nodes in ascending order.	All regions	Configure priority-based resource scheduling

December 2021

Feature	Description	Region	References
ACK One	ACK One is a distributed cloud container platform that is provided by Alibaba Cloud. ACK One allows enterprise users to manage cloud-native applications in hybrid cloud, multi-cluster, distributed computing, and disaster recovery scenarios. You can register external Kubernetes clusters that are deployed in all regions or on all types of infrastructure with ACK One. In addition, ACK One is compatible with the APIs of open source Kubernetes. This allows you to centrally manage and maintain computing resources, networks, storage, security, monitoring, logs, jobs, applications, and traffic.	All regions	ACK One overview
Kubernetes 1.22.3	Kubernetes 1.22.3 can be selected when you create clusters.	All regions	Kubernetes 1.22 release notes
Deployment sets for node pools	Deployment sets are used to manage the distribution of Elastic Compute Service (ECS) instances. ECS instances in a deployment set are distributed across multiple physical servers for high redundancy. This improves the availability of your applications. A node pool that is associated with a deployment set contains ECS nodes that are distributed across multiple physical servers. You can configure pod affinity to deploy your application pods to different ECS nodes. This way, disaster recovery is implemented and the availability of your applications is improved.	All regions	Associate a deployment set with a node pool
Workbench	Workbench is used to log on to containers. Workbench provides higher stability and compatibility compared with the previous terminal tool. To log on to a container by using Workbench, find the container on the Pods page and click Terminal in the Actions column.	All regions	Connection methodsGuidelines on instance connection

Feature	Description	Region	References
Custom configurations for the NGINX Ingress controller	The parameters of the NGINX Ingress controller can be configured on the Add-ons page of the ACK console. The custom configurations are retained after you update the NGINX Ingress controller. You can configure resource requests, resource limits, enable the host network mode, enable admission webhooks, and specify node selectors for the NGINX Ingress controller.	All regions	N/A
Prometheus monitoring dashboards	Monitoring dashboards of the cloud controller manager (CCM) and kube-controller-manager are provided for ACK Pro clusters. This helps you gain better insight into your clusters. You can view the dashboards on the Prometheus Monitoring page in the Operations module of the ACK console.	All regions	Enable ARMS Prometheus
Log center	The log of the cloud controller manager (CCM) can be collected by using the log center feature. The log center feature is available in ASK Pro clusters. This improves the observability of ASK clusters.	All regions	Collect the logs of control plane components in a managed Kubernetes cluster
OPA-based policy governance	The Policy Governance feature is provided by ACK. This feature is developed based on the Open Policy Agent (OPA) policy engine and the gatekeeper admission controller. Policy Governance provides a variety of predefined policies that apply to more Kubernetes scenarios than pod security policies (PSPs). In addition, the configuration of these policies is easy and flexible, which helps the O&M engineers of enterprises better utilize the capabilities that are provided by ACK pod security policies.	All regions	Configure and enforce ACK pod security policies
Node pool priorities	Node pool priorities can be specified in the node pool auto scaling policy. If multiple node pools meet the requirement, ACK selects the node pool with the highest priority for a scale-out activity.	All regions	Auto scaling of nodes

Feature	Description	Region	References
Open source the ALB Ingress controller on GitHub	The ALB Ingress controller is compatible with the NGINX Ingress controller, and provides improved traffic routing capabilities based on ALB instances. The ALB Ingress controller supports complex routing, automatic certificate discovery, and HTTP, HTTPS, and QUIC protocols. The ALB Ingress controller meets the requirements of cloud-native applications for ultra-high elasticity and balancing of heavy traffic loads at Layer 7.	All regions	v2.2.0

November 2021

Feature	Description	Region	References
ALB Ingresses-based routing	ALB Ingresses can be created in the ACK console. You can create Ingress rules in ALB Ingresses and customize Ingress configurations. ALB Ingresses use ALB instances to balance the traffic loads at Layer 7.	All regions	Access Services by using an ALB Ingress
Support for multiple route tables by using CCM	ACK uses the CCM to add route entries to the route table of the virtual private cloud (VPC) in which the cluster is deployed. This enables network connectivity between pods in the cluster. Multiple route tables can be configured for a VPC by using the CCM. To create multiple route tables for a VPC, configure the CCM on the Add-ons page or modify the CCM ConfigMap.	All regions	Configure multiple route tables for a VPC
Desired number of nodes	The number of nodes in a node pool can be changed by setting the desired number of nodes in the node pool. If you set the desired number of nodes to a value larger than the current number of nodes, the node pool is scaled out after you submit the change. If you set the desired number of nodes to a value smaller than the current number of nodes, the node pool is scaled in after you submit the change. The scaling activity is performed based on the scaling policy that you configure when you create the node pool.	All regions	Manage node pools

Feature	Description	Region	References
ACK quotas	The display of quotas and applications for quota increases are optimized by ACK. ACK displays the quotas of ACK managed clusters, ACK dedicated clusters, ASK clusters, ACK edge clusters, registered clusters, and other resources in the console. The ACK console also provides an entry point to Quota Center.	All regions	Limits
IPv4/IPv6 dual stack	IPv4/IPv6 dual stack can be selected when you create an ASK cluster. You can use IPv6 addresses to access services in a cluster that has IPv4/IPv6 dual stack enabled. To enable IPv4/IPv6 dual stack, the following prerequisites must be met: • The Kubernetes version of the cluster is 1.20.11-aliyun.1 or later. • The VPC in which the cluster is deployed supports IPv4/IPv6 dual stack.	All regions	Create an ASK cluster
ContainerOS	ContainerOS is an operating system that Alibaba Cloud provides for containerized development. ContainerOS is fully compatible with Kubernetes. ContainerOS is based on Alibaba Cloud Linux 3 and provides enhanced security, faster startup, and simplified system services and software packages. ContainerOS is preinstalled with components to provide out-of-the-box features in cloud-native scenarios. You can use ContainerOS in managed node pools in ACK clusters. ACK provides long-term free technical support for ContainerOS.	All regions	ContainerOS overview

October 2021

Feature	Description	Region	References

Feature	Description	Region	References
Kubernetes 1.20.11	CVE-2021-25741 is discovered by the Kubernetes community. CVE-2021-25741 is a vulnerability that can be exploited by attackers to access the host directories by using a symbolic link and creating a container that has a subPath volume mounted. The severity of CVE-2021-25741 is rated as high. Upgrade from Kubernetes 1.20 to Kubernetes 1.20.11 is supported. CVE-2021-25741 is fixed in Kubernetes 1.20.11.	All regions	• Vulnerability CVE-2021-25741 in Kubernetes • Kubernetes 1.20 release notes
ClusterRole	The lifecycles of ClusterRoles can be managed in the ACK console. This improves the efficiency of cluster management.	All regions	Customize an RBAC role
ARMS monitoring	The network topology feature of Application Real-Time Monitoring Service (ARMS) is integrated with ACK, which helps enhance the cluster topology feature. This allows you to view the network topology of Services and workloads, and the network topology of resources and Alibaba Cloud services.	All regions	View the network topologies of a cluster
Cost analysis	Application dashboards are provided by the cost analysis feature. Application dashboards provide cost trends, correlation analytic results, and cost saving suggestions and plans.	All regions	Enable cost analysis
Cloud-native AI component set	• The user guide for the cloud-native AI component set is released. • Model management is supported. • Model evaluation is supported. • Accounts other than Alibaba Cloud accounts can be used to log on to AI Dashboard and AI Developer Console. • Fluid applications that use JindoRuntime can be monitored by Prometheus Monitoring.	All regions	• Cloud-native AI component set operations and maintenance guide • Cloud-native AI component set user guide • Model management • Evaluate a model • Access the AI development console • Monitoring and observability

Feature	Description	Region	References
ALB Ingress controller	The ALB Ingress controller is compatible with the NGINX Ingress controller, and provides improved traffic routing capabilities based on ALB instances. The ALB Ingress controller supports complex routing, automatic certificate discovery, and HTTP, HTTPS, and QUIC protocols. The ALB Ingress controller meets the requirements of cloud-native applications for ultra-high elasticity and balancing of heavy traffic loads at Layer 7.	All regions	ALB Ingress overview
ASK Pro cluster	ASK Pro clusters offer higher reliability and security based on ASK clusters. In addition, ASK Pro clusters are covered by the service level agreement (SLA) that supports compensation clauses. This type of cluster is suitable for enterprise users who require high stability and security for large-scale workloads. ASK Pro clusters are now in public preview.	All regions	ASK Pro cluster overview

September 2021

Feature	Description	Region	References
ARM node pools	ARM node pools are supported by ACK clusters. You can select ECS instances of the g6r and c6r instance families when you create ARM node pools. Alibaba Cloud ARM instances offer cost savings in general-purpose computing scenarios where NGINX, Redis, and SQL are used, and provide high concurrency and high throughput in big data computing scenarios.	All regions	Manage node pools
Auto scaling for Windows node pools	Auto scaling can be enabled for Windows node pools to improve the elasticity of your applications.	All regions	Create a Windows node pool
Windows images for node pools	Windows images can be specified when you create node pools. Custom images based on Windows Server 2019 (kernel versions later than 1809) are supported.	All regions	Create a Windows node pool
Multiple security groups for a node pool	More than one security group can be configured for a node pool to enforce fine-grained access control.	All regions	Manage node pools

Feature	Description	Region	References
IPv4/IPv6 dual stack	The IPv4/IPv6 dual stack mode can be enabled when you create an ACK cluster. Then, clients can connect to the application in the ACK cluster through an IPv6 address.	All regions	Create an ACK managed cluster
CIS Kubernetes V1.20 Benchmark v1.0.0	CIS Kubernetes V1.20 Benchmark v1.0.0 is supported by the inspection feature for clusters of Kubernetes 1.20 and later.	All regions	Safety patrol inspection
Node pool scale-out policy and scale-in activity settings	The node pool scale-out policy can be specified and scale-in activities can be set to allowed or disallowed when you configure the auto scaling feature. The node pool scale-out policy decides the priorities of the node pools that are to be scaled out. You can set the node pool scale-out policy to least-waste or random.	All regions	Auto scaling of nodes
Backup center	The backup feature is in public preview. This feature allows you to back up applications deployed in both ACK clusters and self-managed clusters. The application backup feature of ACK is upgraded and renamed as backup center. This feature provides an all-in-one solution for you to back up, restore, and migrate both stateless and stateful applications deployed in ACK clusters. This solution meets the disaster recovery and migration requirements of stateful applications deployed in a hybrid cloud environment or across multiple clusters.	All regions	Backup center overview
Model evaluation for AI project acceleration	Model evaluation includes model management and evaluation. This feature is used to control the versions of models trained by using the AI component set and evaluate models based on multiple metrics, such as the accuracy and recall rate. This feature also provides metrics to help you choose the model that best suits your business.	All regions	• Model management • Evaluate a model

August 2021

Feature	Description	Region	References
ACK Scheduler V1.20-ack-4.0 released to support load-aware scheduling and Elastic Container Instance-based scheduling	• Load-aware scheduling schedules pods to nodes with lower loads based on the historical statistics of node loads. This implements load balancing and prevents application or node exceptions caused by overloaded nodes. • Elastic Container Instance-based scheduling is a scheduling policy that Alibaba Cloud provides for elastic resource scheduling. You can add annotations to specify the resources that you want to use when you deploy applications. You can specify that only ECS instances or elastic container instances are used, or enable the system to request elastic container instances when ECS resources are insufficient. Elastic Container Instance-based scheduling can meet your resource requirements in different workload scenarios.	All regions	• Use load-aware pod scheduling • Use ECI elastic scheduling
CCM 2.0.1 released to support weighted routing across multiple Services and the reuse of existing vServer groups.	• The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-vgroup-port` can be used to reuse an existing vServer group that is attached to a Server Load Balancer (SLB) instance. • When an SLB instance is shared among multiple Services, the annotation `service.beta.kubernetes.io/alicloud-loadbalancer-weight` can be used to set the weight of each Service. • The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-connection-drain` can be used to configure connection draining for an SLB instance. Only TCP and UDP are supported. • The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-connection-drain-timeout` can be used to set the timeout value when connection draining is enabled for an SLB instance. Only TCP and UDP are supported.	All regions	Cloud Controller Manager

Feature	Description	Region	References
Subscription clusters can be created. Subscription and auto-renewal are supported by the SLB instance that is attached to the Kubernetes API server of the cluster.	Subscription clusters can be created. If you select the subscription billing method for a cluster, the ECS nodes and the SLB instance attached to the Kubernetes API server of the cluster use the subscription billing method. You can buy resource plans for elastic IP addresses (EIPs), NAT gateways, and Log Service projects that are used in the cluster. This allows you to complete all payments related to an ACK cluster at a time when you create the cluster and eliminates the hassle of paying the bills on a pay-as-you-go basis.	All regions	N/A
An existing SLB instance can be reused by the NGINX Ingress controller when you create a cluster.	If you select the subscription billing method for a cluster, you can reuse an existing SLB instance for the NGINX Ingress controller. You can associate a subscription SLB instance that you previously created with the NGINX Ingress controller. This allows you to complete all payments related to an ACK cluster at a time when you create the cluster.	All regions	N/A
Intelligent O&M: Cluster diagnostics and global check released	• The global check feature is released to troubleshoot issues in ACK clusters and networks. You can perform a global check on your cluster resources, components, and configurations with a click, and then obtain suggestions on how to fix issues. No parameter configurations are required. • The cluster diagnostics feature is released to troubleshoot nodes, pods, and networks in an ACK cluster.	All regions	• Use the global check feature to troubleshoot cluster issues • Use the cluster diagnostics feature to troubleshoot cluster issues

Feature	Description	Region	References
ASK Ingress ALB controller	The ALB Ingress controller is compatible with the NGINX Ingress controller and provides improved traffic routing capabilities based on ALB instances. The ALB Ingress controller supports complex routing, automatic certificate discovery, and HTTP, HTTPS, and QUIC protocols. The ALB Ingress controller meets the requirements of cloud-native applications for ultra-high elasticity and balancing of heavy traffic loads at Layer 7. You can install the ALB Ingress controller when you create an ACK cluster or on the **Add-ons** page after the cluster is created.	All regions	ALB Ingress overview
Windows Server 2019 supported by ACK edge clusters	Windows Server 2019 can be selected when you create node pools in the cloud for an ACK edge cluster. This improves the cloud-edge coordination capability of Windows applications.	All regions	N/A
Container Network File System (CNFS) enhanced to share and automatically expand Apsara File Storage NAS (NAS) volumes	• CNFS allows you mount a NAS volume in sharepath mode to share the volume among multiple applications or pods. • CNFS allows you to configure auto expansion policies for NAS volumes. If the volume usage exceeds the threshold, the volume is automatically expanded.	All regions	• Use CNFS to share a dynamically provisioned NAS volume • Use CNFS to automatically expand NAS volumes

July 2021

Feature	Description	Region	References
Kubernetes version update	Update from Kubernetes 1.18 to Kubernetes 1.20 is supported.	All regions	• Update the Kubernetes version of an ACK cluster • Kubernetes 1.20 release notes

Feature	Description	Region	References
CoreDNS	CoreDNS is supported on the Add-ons page of the console. CoreDNS is the default plug-in used to implement DNS-based service discovery in ACK clusters and ACK edge clusters. CoreDNS provides domain name resolutions for services within the clusters.	All regions	CoreDNS
Cost analysis based on namespaces	The cost analysis feature is improved to provide resource usage trends and the cost estimation of individual CPU cores per unit time for applications and pods based on namespace.	All regions	Enable cost analysis
Enhanced security for registered clusters	The security of registered clusters is enhanced. You can install security-inspector, aliyun-acr-credential-helper, and gatekeeper in registered clusters. security-inspector is used to perform security scans. aliyun-acr-credential-helper is used to pull images without passwords. gatekeeper is used to manage OPA policies.	All regions	Overview

Feature	Description	Region	References
CNFS 1.0	The CNFS feature is released. CNFS allows you to abstract NAS file systems as Kubernetes resources by using CustomResourceDefinition (CRD) objects. You can use CNFS to create, delete, set quotas for, mount, monitor, and expand NAS file systems. CNFS provides a declarative management method, which is also used to manage containers. CNFS improves storage performance and data security.	All regions	CNFS overview

June 2021

Feature	Description	Region	References
Resource groups selectable	The resource group can be selected from a drop-down list when you create the cluster or node pool in the console. The cluster and ECS instances in the cluster are grouped into the selected resource group. Previously, the resource group was selected at the top of the console. The resource group that you select at the top of the console is used to filter resources displayed on the page, such as VPCs.	All regions	N/A

Feature	Description	Region	References
Visualized configuration for network policies	Kubernetes network policies can be used to configure policy-based network control. You can use network policies to control traffic at the IP address or port level. ACK provides a visual interface that you can use to configure network policies in a convenient manner.	All regions	Use network policies
ACK Terway Hubble	ACK Terway Hubble can be deployed in clusters by using App Catalog. ACK Terway Hubble is a network architecture, workload, and topology observability platform. You can deploy ACK Terway Hubble in an ACK managed cluster to gain observability into the network traffic and network policies.	All regions	Implement network observability by using ACK Terway Hubble
Cost analysis	Cost allocations and trends of resources, applications, and containers can be provided at the node pool level. The cost analysis feature also provides cost optimization suggestions based on the current cost and the sales strategies of node pools.	All regions	Enable cost analysis
Auto scaling	You can set this parameter to configure the interval at which the cluster is evaluated for scaling. You can specify 15 seconds, 30 seconds, and 1 minute as the scan interval.	All regions	Auto scaling of nodes

Feature	Description	Region	References
SANs modifiable for ASK clusters	Custom subject alternative names (SANs) can be modified for the API server certificate of an ASK cluster. This allows you to update the information about the API server certificate, such as the domain name, IP address, and URL, after the ASK cluster is created.	All regions	Update the SAN of the API server certificate for an existing ACK cluster
Cluster security	The inspection feature can be used to detect security risks in the workloads of a registered cluster.	All regions	Use the inspection feature to check for security risks in the workloads of a registered Kubernetes cluster
Topology-aware scheduling	The following scheduling policies are supported by topology-aware CPU scheduling: • Dynamically adjust resource usage thresholds to improve the resource utilization of workloads with different priorities. • Use the Last Level Cache (L3 cache) and Memory Bandwidth Allocation (MBA) to improve the resource isolation of tasks with different priorities.	All regions	• Dynamically adjust resource water marks to improve the resource utilization of workloads with different priorities • Resource isolation based on the L3 cache and MBA

May 2021

Feature	Description	Region	References

Feature	Description	Region	References
CIS reinforcement for worker nodes	Center for Internet Security (CIS) reinforcement is supported for worker nodes. You can enable CIS reinforcement to enhance OS security for cluster nodes. CIS is a third-party security organization that is committed to leading a global community of enterprises, public service sectors, and academia to develop security best practice solutions. CIS reinforcement supports only Alibaba Cloud Linux 2, which is the official OS image of Alibaba Cloud and the default OS image used in ACK clusters.	All regions	CIS reinforcement
New region	ACK Pro clusters are available in Nanjing Local Region.	Nanjing Local Region	N/A
New region	ACK Pro clusters are available in the China North 2 Ali Gov region on Alibaba Gov Cloud.	China North 2 Ali Gov	Supported regions
Cost analysis	The cost analysis feature is added to help IT administrators analyze resource usage and allocate costs. This feature offers suggestions on cost savings and helps improve resource utilization. This feature provides the following services: • Cost analysis of cloud resources • Cost trend analysis • Suggestions on cost savings • Real-time cost forecasting • Cost allocation based on namespaces • Optimization of application costs	All regions	Enable cost analysis
Custom SSL certificates	Custom SSL certificates can be specified for SLB instances by using annotations when you create Ingresses in ASK clusters. The SSL certificates are no longer forcibly specified by using Secrets.	All regions	N/A
Topology-aware scheduling supported by AMD CPUs	resource-controller V1.2.1-d1e280f-aliyun is released. This component works with ack-sceduler of Kubernetes 1.20.4 to support the topology-aware scheduling for AMD CPUs.	All regions	Topology-aware CPU scheduling

April 2021

Feature	Description	Region	References
Kubernetes 1.20	Kubernetes 1.20 is supported. You can select this Kubernetes version when you create a cluster.	All regions	Kubernetes 1.20 release notes

Feature	Description	Region	References
Hot migration from ACK dedicated clusters to ACK Pro clusters	Hot migration from existing ACK dedicated clusters to ACK Pro clusters is supported. You can dynamically migrate workloads from ACK dedicated clusters to ACK Pro clusters without service interruptions.	All regions	Hot migration from ACK dedicated clusters to ACK Pro clusters
NodeLocal DNSCache	ACK NodeLocal DNSCache is a local DNS cache solution developed based on the open source NodeLocal DNSCache project. This solution consists of a DNS caching agent that runs as a DaemonSet and an admission controller that runs as a Deployment to dynamically inject data to DNSConfig. The admission controller listens on pod creation requests and dynamically modifies DNSConfig. This enables pods to use local cache to accelerate DNS lookups.	All regions	ACK NodeLocal DNSCache
Registered cluster	The Kubernetes event center feature and the aliyun-acr-credential-helper component are supported in registered clusters.	All regions	Register an external Kubernetes cluster and Pull images without a password in a self-managed Kubernetes cluster
Custom configuration for control plane components	Custom control plane parameters are supported in ACK Pro clusters to meet the requirements for modifying control plane parameters in production environments. You can modify the parameters of kube-apiserver and kube-controller-manager based on your requirements.	All regions	Customize control plane parameters for a professional Kubernetes cluster
Alerting	The alerting feature is added to enable centralized alert management. ACK allows you to configure alerts to centrally manage exceptions in the cluster and provides various metrics for different scenarios. By default, the alerting feature is enabled when you create clusters. ACK allows you to deploy CRD objects in a cluster to configure and manage alert rules.	All regions	Alert management

March 2021

Feature	Description	Region	References
Data export supported by node pools	Information about nodes in a node pool can be exported to comma-separated values (CSV) files. This improves the O&M efficiency.	All regions	Manage node pools
SANs updatable for ACK managed clusters	Updates to the SANs in the API server certificates are supported for ACK standard and ACK Pro clusters.	All regions	Customize the SAN of the API server certificate when you create an ACK cluster
Temporary kubeconfig files for cluster access	Temporary kubeconfig files are supported for access to ACK clusters. The validity period of a temporary kubeconfig file used to access an ACK cluster ranges from 30 minutes to 3 days. This meets the requirements for temporary access to ACK clusters.	All regions	Connect to ACK clusters by using kubectl
containerd	The containerd runtime is supported by ACK. You can select containerd as the container runtime when you create a cluster. You can also select containerd when you create a regular node pool or a managed node pool. This allows you to deploy both containerd containers and Docker containers in a cluster. Hot migration from Docker containers to containerd containers is not supported. To migrate from Docker containers to containerd containers, you must recreate pods.	All regions	Release notes for containerd

February 2021

Feature	Description	Region	References
ACK edge Pro cluster	ACK edge Pro clusters can be created. This type of cluster provides the same reliability and stability as ACK Pro clusters. The billing methods of ACK edge Pro clusters are also the same as those of ACK Pro clusters.	All regions	Introduction to professional edge Kubernetes clusters
Log center	The log center feature is available in the ACK console. You can check the log of a cluster and the logs of control plane components in the log center.	All regions	View the logs of control plane components and View cluster logs
Prometheus monitoring dashboards	A CoreDNS dashboard is displayed on the Prometheus Monitoring page in the ACK console.	All regions	Enable ARMS Prometheus

Feature	Description	Region	References
EIPs supported by node pools	Public IPv4 addresses can be associated with regular node pools and managed node pools. When you create a regular node pool or a managed node pool, you can enable the nodes to automatically associate with EIPs. This enables the nodes to access the Internet. You can also configure a NAT gateway when you create a cluster to enable all nodes in the cluster to access the Internet by using the NAT gateway.	All regions	Manage node pools
New region	ACK Pro clusters are available in the China South 1 Finance region.	China South 1 Finance	Introduction to professional managed Kubernetes clusters

January 2021

Feature	Description	Region	References
Observability enabled for control plane components of ACK Pro clusters	The observabilities of the API server and etcd control components are enabled in ACK Pro clusters. You can observe these components in monitoring dashboards and receive alerts upon exceptions. This allows you to detect system exceptions and potential risks, and provides information to help you implement measures to ensure the stability of ACK clusters.	All regions	Enable ARMS Prometheus
Custom configuration for control plane components of ACK Pro clusters	Custom parameters are supported for kube-apiserver and kube-controller-manager in ACK Pro clusters. This meets the requirements for custom parameters of cluster control components in production environments.	All regions	Customize control plane parameters for a professional Kubernetes cluster
Log collection for control plane components	Logs of control components, such as kube-apiserver, kube-controller-manager, and kube-scheduler, can be collected. To enable log collection, select Enable for Log Collection for Control Plane Components when you create a cluster. This helps you monitor the cluster status and detect anomalies in the cluster.	All regions	View the logs of control plane components

Feature	Description	Region	References
Preemptible instances supported by node pools	Preemptible instances are supported when you set the billing method of a node pool. Preemptible instances are cost-effective. You can bid for idle resources of Alibaba Cloud, obtain the resources, and then run containers until the resources are reclaimed due to higher bids from other customers. This reduces the cost of computing resources.	All regions	Set the ratio of preemptible instances to pay-as-you-go instances
Edge node pool	Edge node pools are supported in ACK edge clusters. You can abstract a set of nodes with one or more identical attributes into an edge node pool for an ACK edge cluster. This way, you can use the edge node pool to manage and perform O&M operations on nodes from different regions in a unified manner. An edge node pool uses the basic or enhanced coordination network between the cloud and edge. The enhanced coordination network is built by using the software-defined networking (SDN) solution of ACK@Edge, and allows you to coordinate cloud and edge computing in a secure and fast network environment. This allows applications deployed in edge node pools to access the cloud through the VPC where the cluster is deployed. Compared with the basic coordination network, the enhanced coordination network provides higher network quality and improves data security.	All regions	Overview of edge node pools
Elastic node pool supported by registered clusters	Node pools are supported in registered clusters. You can use a node pool to manage a set of ECS instances with the same attributes. You can also add them to a self-managed Kubernetes cluster or a Kubernetes cluster that is deployed in the public cloud of a third-party cloud service provider. This allows you to schedule resources in the cloud, data centers, and self-managed Kubernetes clusters in a unified, flexible, and cost-effective manner.	All regions	Configure auto scaling

Feature	Description	Region	References
Application backup	The application backup feature is released. This feature meets the critical requirement for data security in Kubernetes clusters where an increasing number of applications are deployed. You can use application backups to restore applications that are accidentally disrupted for a long period of time. Different from the traditional single-server backup and disk backup, the application backup feature is used to back up applications and the relevant data, resource objects, and configurations. You can also use this feature to back up all resources in a namespace. This feature is available in ACK clusters and registered clusters. You can use this feature to back up applications, volumes, and persistent volumes (PVs) in a cluster, and also restore backups to other clusters.	All regions	Enable cluster backup
Cost reduction policy	The ratio of preemptible instances to pay-as-you-go instances can be set in a node pool. This allows you to reduce the cost. However, you must make sure that the node pool has enough pay-as-you-go instances to ensure performance stability.	All regions	Set the ratio of preemptible instances to pay-as-you-go instances

December 2020

Feature	Description	Region	References
New region	ACK is now available in the China (Guangzhou) region.	China (Guangzhou)	Limits
Hot migration from ACK standard clusters to ACK Pro clusters	Hot migration from existing ACK standard clusters to ACK Pro clusters is supported. Your services are not affected during the migration. ACK Pro clusters are developed based on ACK managed clusters. This type of cluster provides higher reliability and security in large-scale production environments for enterprise users. ACK Pro clusters are also covered by SLAs that include compensation clauses.	All regions	Hot migration from ACK standard clusters to ACK Pro clusters
SLB specification selectable for the API server	The specification of the SLB instance that is used to access the API server can be selected when you create an ACK cluster. You can select different SLB specifications based on your business requirements. This allows you to handle different traffic loads on the API server of the cluster.	All regions	Create an ACK Pro cluster

Feature	Description	Region	References
Preemptible instances supported by node pools	Preemptible instances are supported when you set the billing method of a node pool. Preemptible instances are cost-effective. You can bid for idle resources of Alibaba Cloud, obtain the resources, and then run containers until the resources are reclaimed due to higher bids from other customers. This reduces the cost of computing resources.	All regions	N/A
Kubernetes 1.18	Updates from Kubernetes 1.16 to 1.18 are supported.	All regions	Update the Kubernetes version of an ACK cluster
CronHPA	Cron Horizontal Pod Autoscaler (CronHPA) can be enabled in the ACK console for your workloads. You must install ack-kubernetes-cronhpa-controller in the cluster before you enable CronHPA.	All regions	CronHPA
CentOS 7.8	CentOS 7.8 can be used as the node OS when you create a cluster or a node pool.	All regions	Manage node pools
Reinforcement based on classified protection	Reinforcement based on classified protection is supported for the cloud-native Alibaba Cloud Linux operating system in compliance with Multi-Level Protection Scheme (MLPS) 2.0 level 3 standards. The following features are provided: • Implement identity authentication • RAM • Security auditing • Intrusion prevention • Malicious code protection To enable reinforcement based on classified protection for the node OS when you create a cluster or a node pool, you must select Alibaba Cloud Linux 2.1903 as the node OS and select **Reinforcement based on classified protection**.	All regions	
Volume snapshots supported by CSI	Volume snapshots created from disks are supported by the Container Storage Interface (CSI) component of ACK. This allows you to back up and restore workload data.	All regions	Use volume snapshots created from disks
Cluster upgrade and new components	ASK clusters can be upgraded. The metrics-server, cronhpa-controller, and alb-ingress-controller components can be installed and managed on the Add-ons page of the ACK console.	All regions	N/A

November 2020

Feature	Description	Region	References
Managed node pools	Managed node pools that are provided by ACK support auto upgrade and auto repair. This provides centralized, managed, and O&M-free lifecycle management of nodes. You do not need to be concerned about the O&M of nodes, such as component upgrading, OS upgrading, or patching to fix Common Vulnerabilities and Exposures (CVE) vulnerabilities. ACK automatically fixes node exceptions for the nodes in a managed node pool. Managed node pools are supported by ACK Pro clusters.	All regions	Overview
Kubernetes 1.18 supported by kubernetes-dashboard	Kubernetes 1.18 is supported by the kubernetes-dashboard application provided by App Catalog. This fixes the issue that the pods of Kubernetes 1.18 cannot be accessed by terminals. You can find and install the Helm chart for kubernetes-dashboard from App Catalog.	All regions	App Marketplace
Performance levels of Enhanced SSDs configurable	The performance level of an enhanced SSD can be set to PL0 or PL1 when you create a cluster. This allows you to customize the performance level of your cluster. This feature is supported by ACK Pro clusters, ACK standard clusters, ACK dedicated clusters, and ACK edge clusters.	All regions	FAQ
CCM	The CCM is updated to V1.9.3.339-g9830b58-aliyun. Hash values are supported in the configurations of LoadBalancer Services. This way, when the CCM is restarted, only the backend vServer groups of the related SLB instances are updated if the Service configuration is not changed. The configurations of the related SLB instances and listeners are not updated.	All regions	Cloud Controller Manager
Disk monitoring supported by CSI	Disk monitoring is supported by the latest version of the CSI component. This feature allows you to monitor the status of persistent volume claims (PVCs) through ARMS Prometheus when you use disks that are mounted by using the PVCs. You can also configure alerts by setting thresholds for the storage space and input/output operations per second (IOPS) of the disks.	All regions	N/A
Ingress controller and CoreDNS	Ingress controllers and CoreDNS can be installed when you create an ASK cluster. You can also install CoreDNS on the Add-ons page of the ACK console after the cluster is created.	All regions	Create an ASK cluster

Feature	Description	Region	References
Node pools supported by registered clusters	Node pools are supported in registered clusters. You can use a node pool in the ACK console to manage a set of ECS instances for a registered cluster. You can add ECS nodes from a node pool to a self-managed Kubernetes cluster or a Kubernetes cluster that is deployed in the public cloud of a third-party cloud service provider. You can also use node pools to manage the labels and taints of nodes in node pools.	All regions	Manage node pools

October 2020

Feature	Description	Region	References
Time zone	The time zone can be selected when you create a cluster. By default, the time zone of your browser is selected. This feature is supported by ACK Pro clusters, ACK standard clusters, ACK dedicated clusters, and ASK clusters.	All regions	Create an ACK managed cluster
Tagging supported for cloud disks, NAS file systems, and Log Service projects	Cloud disks, NAS file systems, and Log Service projects with tags are supported by CSI and Logtail. Cloud disks, NAS file systems, and Log Service projects that are created by ACK for a cluster are added with the cluster ID as tags. This makes it easier to allocate resource fees.	All regions	N/A

September 2020

Feature	Description	Region	References
New region	ACK is available in the China (Ulanqab) region.	All regions	Introduction to professional managed Kubernetes clusters
SMB supported by Windows containers	Server Message Block (SMB) file systems can be mounted to a Windows container. In the NAS console, you can create an SMB file system in the VPC where the cluster is deployed. You can also create a mount target for the file system. You must use the FlexVolume plug-in to mount an SMB file system.	All regions	Mount disks and SMB file systems to Windows pods

Feature	Description	Region	References
Time zone	The time zone can be selected for master nodes and worker nodes when you create an ACK dedicated cluster or an ACK managed cluster.	All regions	N/A
Kubernetes 1.18	Kubernetes 1.18.8 is supported. You can select this Kubernetes version when you create a cluster. ACK clusters of Kubernetes 1.18 or later no longer support Kubernetes Dashboard. To use Kubernetes Dashboard, we recommend that you install **kubernetes-dashboard** on the App Catalog page.	All regions	Kubernetes 1.18 release notes and [Product Changes] ACK ends support for Kubernetes Dashboard
NetworkPolicy configurable for Terway	The NetworkPolicy feature can be enabled or disabled for Terway when you create a cluster.	All regions	• Use network policies • Improve the performance of the NetworkPolicy feature for a large ACK cluster in Terway mode
Periodic inspection	Periodic inspection policies can be configured for a cluster on the Inspections page of the ACK console.	All regions	Use the inspection feature to detect security risks in the workloads of an ACK cluster
Cluster auditing	The cluster auditing feature can be enabled or disabled on the Cluster Auditing page of the ACK console.	All regions	Use cluster auditing
New components	The logtail-ds component is provided to collect container log from registered clusters, including stdout and log files of containers. The migrate-controller component is provided to migrate applications across Kubernetes clusters. This component is developed based on the open source Velero project. The ack-virtual-node component is provided to enable auto scaling for registered clusters.	All regions	• Enable Log Service for an external Kubernetes cluster • Enable cluster backup

Feature	Description	Region	References
Sandboxed-Container 2.0	Sandboxed-Container is updated to V2.0. Sandboxed-Container 2.0 has the following benefits: • Sandboxed-Container is a container runtime that is developed by Alibaba Cloud based on lightweight virtual machines. Compared with Sandboxed-Container 1.0, Sandboxed-Container 2.0 supports more lightweight and efficient deployment and simplifies the architecture and maintenance of ACK clusters. • Sandboxed-Container 2.0 reduces the resource overheads by 90% and improves the startup speed of sandboxed containers by three times. • Sandboxed-Container 2.0 increases the deployment density of sandboxed containers on a single node by 10 times. • Sandboxed-Container 2.0 supports the virtio-fs file system, which provides higher performance than the 9pfs file system.	All regions	Sandboxed-Container overview
Knative component supported by ASK clusters	Knative components are supported in ASK clusters. Knative is a cloud-native and cross-platform orchestration engine for serverless applications. You can deploy Knative in ASK clusters. This allows you to use cloud resources by calling the Knative API without the need to pay for the Knative controller.	All regions	Overview

August 2020

Feature	Description	Region	References
Gatekeeper	The gatekeeper component can be installed on the Add-ons page of the ACK console. This component facilitates the management and implementation of policies that are executed by OPA in ACK clusters.	All regions	gatekeeper
Runtime inspection	Runtime inspections can be performed on the Runtime Security page of the ACK console. This feature monitors the container runtime and triggers alerts upon the following types of security events: malicious image startups, attacks by viruses or malicious programs, intrusions into containers, container escapes, and high-risk operations on containers. To use this feature, you must first activate Security Center. If you use a Resource Access Management (RAM) user, make sure that the RAM user has the permissions to access Security Center.	All regions	Use the runtime security feature to monitor ACK clusters and configure alerts

Feature	Description	Region	References
Scheduled backup	Scheduled backups are supported for Elastic Block Storage (EBS) devices. You can create scheduled snapshots from disks. To use this feature, you must first install the cluster-storage-operator component.	All regions	N/A
IPvlan and eBPF supported by Terway	IPvlan and extended Berkeley Packet Filter (eBPF) are supported by Terway. If an elastic network interface (ENI) is shared among pods, Terway allows you to use IPvlan and eBPF for network virtualization. Terway enables pod network virtualization by using the lightweight IPvlan technology. This allows pod traffic to bypass the network stack of the host and reduces the network performance overheads. Terway uses Cilium as the BPF agent on nodes to configure BPF rules for pod ENIs. This enables Services and network policies to be configured on ENIs. This way, requests within pod networks are forwarded to ENIs through IPvlan. This reduces network complexity. **Note** This feature applies to the Alibaba Cloud Linux 2 operating system. To use this feature, you must to apply to be added to a whitelist.	All regions	Work with Terway
New regions	ACK Pro clusters are available in the China (Beijing), China (Shenzhen), Germany (Frankfurt), Indonesia (Jakarta), and China East 2 Finance regions.	China (Beijing), China (Shenzhen), Germany (Frankfurt), Indonesia (Jakarta), and China East 2 Finance	Introduction to professional managed Kubernetes clusters
ACK@Edge released for commercial use	ACK@Edge is released for commercial use. ACK@Edge is a cloud-managed solution that is provided by ACK to coordinate cloud and edge computing.	All regions	ACK@Edge overview

July 2020

Feature	Description	Region	References

Feature	Description	Region	References
ACK Pro cluster released for public preview	ACK Pro clusters are released for public preview. This type of cluster is developed based on ACK managed cluster and provides higher reliability and security in large-scale production environments for enterprise users. ACK Pro clusters are also covered by SLAs that include compensation clauses. This type of cluster is suitable for the following users: • Internet enterprises. These enterprises deploy their business in large-scale production environments and require business management with high stability, security, and observability. • Big data computing enterprises. These enterprises deploy large-scale data computing services, high-performance data processing services, and other services with high elasticity. These services require clusters with high stability, high performance, and efficient computing capabilities. • International enterprises that run their business in China. These enterprises prioritize security and services that provide SLAs with compensation clauses. • Financial enterprises. These enterprises require SLAs that include compensation clauses.	All regions	Introduction to professional managed Kubernetes clusters
New regions	ASK is available in the Japan (Tokyo) and Indonesia (Jakarta) regions.	Japan (Tokyo) and Indonesia (Jakarta)	ASK overview
CCM	The CCM is updated to V1.9.3.313-g748f81e-aliyun. The following features are provided: • Supports deletion protection for SLB instances. By default, deletion protection is enabled for newly created SLB instances. • Supports modification protection for the configurations of SLB instances. By default, modification protection is enabled for the configurations of newly created SLB instances. • Allows you to specify the resource group for an SLB instance when you create a Service. • Allows you to specify the name of an SLB instance when you create a Service. • Allows you to mount pods in Terway mode to the backend of an SLB instance.	All regions	Cloud Controller Manager

Feature	Description	Region	References
Security management	Security management is supported for your clusters. You can configure pod security policies and cluster inspections. Pod security policy is a significant method to verify the security of pod configurations before pods are deployed. This ensures that applications are running in secure pods. Cluster inspection detects the security risks of workloads in an ACK cluster and generates inspection reports for your reference. This way, you can check whether the workloads in your ACK cluster run in a secure environment.	All regions	Configure pod security policies (earlier version)
Shared VPCs supported by ACK	Shared VPCs are supported. A shared VPC can host cloud resources that are created by multiple accounts. The cloud resources include ECS instances, SLB instances, and ApsaraDB RDS instances. This provides a unified approach for you to manage cloud resources in a shared VPC. Shared VPCs are powered by the resource sharing mechanism. The Alibaba Cloud account that owns a shared VPC can share all vSwitches in the VPC with other accounts in the same organization. You can select a shared VPC when you create an ACK cluster. If you select a shared VPC for an ACK cluster, you can use only Terway as the network plug-in.	All regions	N/A
Cluster registration	Cluster registration is supported. During daily O&M, you may need to deploy multiple clusters in the cloud and data centers. In some scenarios, you may even deploy clusters in the clouds of different cloud service providers. In these cases, you can register external Kubernetes clusters in the ACK console. This allows you to manage external Kubernetes clusters in the console and reduce O&M costs.	All regions	Overview
Workload management	Redeployment and rollback of workloads are supported. ACK provides features on the workload management page in the ACK console, such as application redeployment and rollback. This makes it more convenient to manage your workloads.	All regions	Create a stateless application by using a Deployment

June 2020

Feature	Description	Region	References
Taint management	Taint management is supported for node pools. You can configure taints when you create or modify a node pool. This allows you to add taints to all nodes in the node pool. You can select **Synchronize Node Labels and Taints** to update taints for existing nodes in a node pool.	All regions	Manage taints

Feature	Description	Region	References
Application migration from virtual machines to ACK clusters	Application migration from virtual machines to ACK clusters by using Server Migration Center (SMC) is supported. SMC allows you to migrate servers to Container Registry. You can use SMC to migrate containerized applications to Container Registry at low costs.	All regions	Migrate source servers to Container Registry

May 2020

Feature	Description	Region	References
Advanced security groups	Advanced security groups are supported when you create a cluster. You can select a basic security group, an advanced security group, or an existing security group. Compared with a basic security group, an advanced security group can contain up to 65,536 private IP addresses. Advanced security groups are used for clusters where a large number of containers or instances are deployed.	All regions	Create an ACK managed cluster
Component management	The Prometheus component and Kubernetes event center can be installed from the Add-ons page of the ACK console. ACK is integrated with the most commonly used Prometheus component in the container monitoring field, and the most commonly used node-problem-detector (NPD) component in the O&M field. You can select these components when you create a cluster. You can also upgrade and maintain the components on the Add-ons page of the ACK console. The Prometheus component is provided by ARMS. NPD is a tool used for node problem detection. NPD can export events that record node exceptions, such as Docker Engine hangs, Linux kernel hangs, network access issues, and file descriptor issues. You can click the **Event Center** tab on the **Events** page to view event details.	All regions	Enable ARMS Prometheus
Kubernetes 1.16.9	Kubernetes 1.16.9 is supported. You can create a cluster of Kubernetes 1.16.9. If the Kubernetes version of your cluster is earlier than V1.16.9, go to the Clusters page and choose **More > Upgrade Cluster** in the Actions column to upgrade to Kubernetes 1.16.9. Compared with Kubernetes 1.16.6, Kubernetes 1.16.9 fixes the CVE-2020-8555 SSRF vulnerability for the kube-controller-manager component.	All regions	Vulnerability fixed: CVE-2020-8555 in kube-controller-manager

Feature	Description	Region	References
Elastic workload	Elastic workloads are supported. You can go to the **App Catalog** page and select **ack-kubernetes-elastic-workload** to install the component. You can use ACK and Virtual Kubelet in combination to proportionally schedule pay-as-you-go and preemptible instances. This allows you to schedule your workloads with elasticity.	All regions	App Marketplace
Application center	The application center feature is released in the ACK console. In earlier versions of the ACK console, after applications are deployed, the topology of the applications is not displayed in a unified view. Therefore, version management and rollback cannot be unified for continuous deployments. The application center provides a unified portal for your applications. This allows you to view the deployment of applications in a unified manner. You can also view the deployment status and changes of all ACK sub-resources that are allocated to each application. In addition, Gits and Helm charts are used to deploy applications in ACK clusters by versions. This allows you to publish or roll back different application versions deployed in ACK clusters.	All regions	Application Center overview

April 2020

Feature	Description	Region	References
AGS released for commercial use	Alibaba Cloud Genomics Service (AGS) is released for commercial use. AGS is an ACK-based big data compute service provided by Alibaba Cloud for users in the biotechnology industry. AGS provides efficient, elastic, and reliable services. AGS is faster in computing and more cost-effective than traditional methods. AGS uses the pay-as-you-go billing method and charges you based on the number of successful API calls in the backend. To submit a computing task, you only need to run a command to call the AGS API on the client. This process is counted as one API call.	All regions	AGS overview
Online expansion supported by dynamically provisioned volumes	Expansion of dynamically provisioned volumes without restarting pods is supported for Kubernetes 1.16 and later.	All regions	Expand a disk volume without service interruptions

Feature	Description	Region	References
Ingress controller	Multiple Ingress controllers can be deployed in a Kubernetes cluster. An Ingress is an important entry for Layer 7 services. If you create only one Ingress for a cluster, the routing performance may encounter a bottleneck. If an Ingress allows inbound access through the Internet and private network at the same time, security risks exist. To solve these issues, ACK provides a Helm chart for the Ingress controller when only one Ingress is used. The name of the Helm chart is ack-ingress-nginx. You can deploy multiple Ingress controllers from App Catalog. You can use YAML files to configure access to Internet-facing and internal-facing SLB instances separately.	All regions	Deploy Ingresses in a high-reliability architecture
New region	ASK is available in the India (Mumbai) region.	India (Mumbai)	Create an ASK cluster

March 2020

Feature	Description	Region	References
Component management	The following features are added for component management: • Allows you to view the YAML files of components. • Allows you to perform health checks for nodes before component upgrades. This prevents component upgrade failures that are caused by node drains or exceptions. • Allows you to manually refresh the Add-ons page.	All regions	Manage system components
Self-managed ECS instances	Nodes that run on self-managed ECS instances can be added to the backend of SLB instances by using the CCM. This way, the existing applications and containerized applications share the same SLB instances and inbound traffic. This is suitable for scenarios where existing applications are gradually replaced by containerized applications.	All regions	Cloud Controller Manager

Feature	Description	Region	References
Cluster expansion and node specification changes	Cluster expansion and node specification changes are supported by Terway. When you manually expand a cluster, you may need to create nodes in new zones. In earlier versions, to create pods in a new zone, you must first add new pod vSwitches in the zone. You can add pod vSwitches in Terway ConfigMaps. When you change the specifications of a node, the maximum number of pods that are supported by Terway on the node also changes. After this release, the K8s max-pod parameter is automatically adjusted to fit the new node specifications.	All regions	Work with Terway
Node pool management	Node pool management is supported. A node pool contains a group of nodes with the same configurations. For example, nodes in a node pool are configured with the same container runtime, OS, and security group. You can create multiple node pools for a cluster. This allows you to deploy a variety of services to different node pools in a cluster. Node pools also support auto scaling. Nodes can be automatically added when a node pool is short of required resources.	All regions	Manage node pools
Cluster inspection	Cluster inspection is optimized. Cluster inspection is the core feature provided by ACK for cluster O&M. Cluster inspection dynamically scans clusters to identify potential risks. The optimization provides the following features: • Displays information about unknown hosts. • Checks the availability of Yellow dogUpdater, Modified (YUM). • Checks the availability of systemd.	All regions	Use the global check feature to troubleshoot cluster issues
Kubernetes 1.16	Update to Kubernetes 1.16.6 is supported. You can update your clusters from Kubernetes 1.14.8 to 1.16.6. You can also create clusters that run Kubernetes 1.16.6. We recommend that you read the update notes before you update your clusters.	All regions	Update the Kubernetes version of an ACK cluster
New region	ACK managed clusters are available in the China South 1 Finance region.	China South 1 Finance	Create an ACK managed cluster

Feature	Description	Region	References
ephemeral-storage	The ephemeral-storage parameter is added for container configurations when you create an application. Ephemeral storage is a new storage resource similar to CPU and memory resources. Kubernetes uses this parameter to manage and schedule the transient storage of applications that run in Kubernetes clusters. The root directory and log directories (*/var/log*) of kubelet are stored on the primary partition of a node. In addition, emptyDir volumes, container log, image layers, and the writable layers of containers are also stored on the primary partition. Therefore, ephemeral-storage is used to manage the primary partition of a node. You can set requests and limits when you create an application. This allows you to schedule and manage the storage resources that are allocated from the primary partition to the application.	All regions	Create a stateless application by using a Deployment

February 2020

Feature	Description	Region	References
Kubernetes 1.16 and Docker 19.03.5	Kubernetes 1.16 and Docker 19.03.5 are supported to provide enhanced cloud-native capabilities. Compared with the earlier version, Kubernetes 1.16 accelerates pod creation and improves affinity, stability, and observability. You can select Docker 19.03.5 when you create a cluster. ACK accelerates container startups and the building of images that are based on Docker 19.03.5.	All regions	Kubernetes 1.16 release notes
Auto scaling	The auto scaling feature is updated to add the following options: the Aliyun Linux2 operating system, custom security groups, and preemptible instances with GPU capabilities. To use AliyunLinux 2 and custom security groups, you must first submit a ticket to enable them for your account.	All regions	
CentOS 7.7	CentOS 7.7 is supported as the node OS. You can specify the CentOS 7.7 operating system when you create worker nodes. CentOS 7.7 is automatically used when you expand clusters or enable auto scaling for clusters.	All regions	

Feature	Description	Region	References
Helm 3	Helm 3 is supported. You can install Helm 3 from App Catalog. Compared with Helm 2, Helm 3 improves the security of role assignment, provides full compatibility with Kubernetes role-based access control (RBAC) in multi-tenant scenarios, and supports hooks for more management operations.	All regions	For more information about how to update from Helm 2, see [Component Upgrades] Upgrade Helm V2 to V3.
New regions	ASK is available in the Indonesia (Jakarta) and UK (London) regions. You can create ASK clusters in these regions in the ACK console.	Indonesia (Jakarta) and UK (London)	Create an ASK cluster
ClusterIP Services in ASK clusters	ClusterIP Services are supported in ASK clusters. This provides more options when you deploy containerized applications in ASK clusters. You can create ClusterIP Services in an ASK cluster to enable access to your workloads from within the ASK cluster.	All regions	Manage Services
CCM	ECS instances and elastic container instances can be attached to the backend of SLB instances that are associated with Services by using the CCM. This enables unified scheduling for application pods across worker nodes and virtual nodes. This also improves application resilience.	All regions	Release notes for the CCM
ACK edge clusters	32-bit and 64-bit ARM nodes are supported in ACK edge clusters. This allows ACK edge clusters to support more heterogeneous infrastructures. You can add Edge Node Service (ENS) nodes or nodes from data centers to ACK edge.	All regions	Add an edge node

January 2020

Feature	Description	Region	References
Virtual nodes	ClusterIP Services can be accessed by pods that are deployed on virtual nodes. This enables Kubernetes to centrally manage virtual nodes and elastic container instances. You can deploy applications on virtual nodes without the inconvenience of resource capacity planning. This meets the requirements of scenarios such as online workload scaling, offline computing, and CI/CD, and also reduces the overall computing costs. To enable this feature, log on to the console, click App Catalog, and then find and install ack-virtual-node.	All regions	Deploy the virtual node controller and use it to create Elastic Container Instance-based pods

Feature	Description	Region	References
API server	Service account token volume projection can be enabled for the API server when you create a cluster. This enables service account authentication on pods. This feature is also required if mutual Transport Layer Security (TLS) authentication is enabled on Istio through Secret Discovery Service (SDS).	All regions	Create an ACK dedicated cluster
CSI	CSI can be selected as the volume plug-in when you create an ACK cluster. The optimized CSI plug-in provides the following features: • Object Storage Service (OSS) subdirectories can be mounted to containers. • The Memory type emptyDir volumes are supported. The Memory type volume is a RAM-based temporary file system, whose storage space is limited by memory. This type of file system provides good performance and is typically used to provide caching space in containers. • Accelerated OSSFS transmission is supported. OSSFS allows you to share data by mounting OSS buckets to local file systems in Linux. To meet the requirements of big data and AI scenarios, ACK improves read speed by adjusting concurrency, block size, and libfuse configurations. For more information, see alibaba-cloud-csi-driver.	All regions	Install CSI
Sandboxed containers	Disks and NAS file systems can be mounted to sandboxed containers to enhance cloud-native capabilities. This allows ACK to provide the same storage performance as when these storage services are used on virtual machines. ACK also supports RootFS BLKIO Limit and disk I/O throttling on pods, and optimizes its support for multi-tenancy.	All regions	Mount a NAS file system to a sandboxed container and Mount a disk to a sandboxed container

Feature	Description	Region	References
ACK clusters for confidential computing	ACK clusters for confidential computing are released for public preview. This type of cluster is developed on top of Intel Software Guard Extensions (SGX) and is particularly suitable for sensitive data protection and scenarios such as smart contracts in blockchains, user secrets processing, intellectual property protection, genomics computing in bioinformatics, and edge computing. You can create and manually expand ACK clusters for confidential computing. You can also enable auto scaling, and add different types of nodes to the clusters. For more information, see Create an ACK managed cluster for confidential computing and SGX application development guide. ACK also provides open source sgx-device-plugin to help you deploy SGX applications in ACK clusters. For more information, see Kubernetes device plugin for Intel SGX. **Note** Intel (R) SGX is a set of CPU instruction codes that are developed by Intel. Intel (R) SGX allows you to run application code and data in a special runtime environment called enclave, which is built on top of hardware silos and memory encryption technologies. Enclaves refer to Trusted Execution Environment (TEE). No application, OS Kernel, BIOS, or hardware other than the CPU can access an enclave without verification. All data in the enclave memory is encrypted. Users encrypt the code and data in an enclave with their private keys obtained from Intel. An enclave can be started only after the signature is verified through Intel Attestation Service (IAS), which is a remote certification service of Intel.	All regions	Create an ACK managed cluster for confidential computing

Feature	Description	Region	References
AGS	Gene sequencing is supported by calling AGS API operations. ACK has released a set of AGS API operations. You can call these API operations to submit gene sequencing tasks. Results are automatically uploaded to your OSS buckets. This saves you the inconvenience of cluster creation and task deployments. These API operations support different SLA levels and provide computing resources based on different requirements. This allows you to reduce costs and improve efficiency. This feature is in public preview. To use the feature, submit a ticket.	All regions	Use AGS to process WGS tasks

December 2019

Feature	Description	Region	References
Component management	Component management is supported. You can log on to the ACK console. On the **Clusters** page, find the cluster that you want to manage and choose **More > Manage System Components** in the Actions column to manage cluster components. You can manage all system components and optional components with operations such as upgrade, uninstall, and reinstall. Custom component configurations will be available soon.	All regions	Manage system components
App Catalog	The ack-node-local-dns plug-in is provided in App Catalog to speed up Domain Name Service (DNS) queries. ack-node-local-dns sends internal DNS queries to CoreDNS and directly forwards external DNS queries to external DNS resolvers. ack-node-local-dns caches all queries and provides DNS caching on each node. This significantly improves the overall DNS query rate of the cluster.	All regions	App Marketplace
New region	ACK managed clusters are available in the China East 1 Finance region. You only need to create worker nodes in an ACK managed cluster. ACK creates and manages master nodes. This type of cluster is easy to use and provides high availability at low costs. This saves you the inconvenience of master node O&M and allows you to focus on business development.	China East 1 Finance	Create an ACK managed cluster
NPU-accelerated ECS instances	Neural processing unit (NPU)-accelerated ECS instances are supported when you create ACK managed clusters or ACK dedicated clusters. The instance type is ecs.ebman1.26xlarge, which is suitable for big data analytics and AI scenarios in video and graphics industries.	All regions	Create an ACK managed cluster

Feature	Description	Region	References
Terway	The user experience of Terway is improved. The new user interface provides information about the number of pods that are supported by each ECS instance type when you create a cluster. When you expand a cluster, the user interface also provides multiple options. This allows you to select vSwitches for nodes and pods. The user interface is optimized to provide easy-to-read, accurate information.	All regions	Work with Terway

November 2019

Feature	Description	Region	References
Cluster expansion	Multiple zones and multiple data disks are supported when you expand an ACK cluster. The user interface for expanding an ACK cluster is updated to provide the same configuration options as those for creating an ACK cluster. You can select multiple zones when you expand an ACK cluster. You can also mount multiple data disks to a node and specify whether to encrypt these disks.	All regions	Increase the number of nodes in an ACK cluster
Custom node configurations	Custom scripts, tags, and Operation Orchestration Service (OOS) are supported for node configurations. You can write custom scripts to configure nodes when you create or expand an ACK cluster. To use this feature, submit a ticket to enable this feature for your account. You can use this feature to specify the node OS. Instead of building custom images, you can directly inject scripts into standard images. Auto scaling allows you to add tags to cluster nodes. This makes it easier for you to identify cluster nodes and allocate the cost of nodes. ACK integrates OOS into the node O&M. You can go to the OOS page from the ACK console and run OOS scripts to maintain nodes on the OOS page.	All regions	Increase the number of nodes in an ACK cluster
Multiple zones and log auditing supported in ASK clusters	Multiple zones and log auditing are supported in ASK clusters. After ASK is upgraded to V2.0, ASK clusters provide more cloud-native features. Cross-zone ASK clusters and log auditing are supported. You can deploy pods across zones to improve the availability of your business. You can also use log auditing to improve the security of ASK clusters. ASK clusters will be improved to provide the same features as dedicated and ACK managed clusters.	All regions	Create an ASK cluster

Feature	Description	Region	References
vGPU	vGPU resources are provided through the vgn5i instance family to meet the requirements of AI and big data industries. You can select instance types of the vgn5i instance family when you create an ACK cluster.	All regions	N/A
ENI buffer pools for Terway	ENI buffer pools are supported for Terway. Terway is a container network plug-in that is developed on top of Alibaba Cloud ENI. The update enables Terway to create a buffer pool of ENI IP addresses during node initialization. This accelerates pod creation and improves user experience.	All regions	Work with Terway
CCM	External ECS instances can be added to the backend of SLB instances by using the CCM. The CCM is a system component that associates Services with SLB instances. By default, cluster nodes that host Services are mounted to the backend of the related SLB instances. The update allows you to add ECS instances outside an ACK cluster as the backend servers to the related SLB instances. This makes it easier to perform application migration and canary releases.	All regions	Cloud Controller Manager

October 2019

Feature	Description	Region	References
AliyunLinux2	The AliyunLinux2 operating system is supported. AliyunLinux2 is the latest OS version that is developed by Alibaba Cloud on top of an advanced CentOS kernel version. AliyunLinux2.1903 is fully adapted to ACK. This OS version supports faster startups and optimized performance, and improves the efficiency and reliability of ACK clusters.	All regions	Create an ACK dedicated cluster
Ingress dashboard	The Ingress dashboard is provided. In earlier versions, you must manually configure the Ingress dashboard, which is a time-consuming and error-prone task. A check box is added to the configuration page of the Ingress controller. You need to select the check box to enable the Ingress dashboard feature. This way, the Ingress dashboard is automatically installed after the cluster is created.	All regions	Create an ACK dedicated cluster

Feature	Description	Region	References
SLB instance specifications	Multiple SLB instance specifications are supported when you create a Service. In earlier versions, when you create a LoadBalancer Service, ACK automatically creates shared-performance SLB instances. To meet your requirements in various scenarios, ACK allows you to select SLB instance specifications when you create a LoadBalancer Service. The SLB instances adopt the pay-as-you-go billing method.	All regions	Manage Services
EIPs for the API server	An EIP can be associated to or disassociated from the API server of an ACK cluster. SLB instances provide access to the API server of an ACK cluster. When you create an ACK cluster, ACK allows you to specify an Internet-facing or internal-facing SLB instance to handle traffic to the cluster. However, you may need to change the network type of the SLB instance after the cluster is created. ACK allows you to associate an EIP with or disassociate an EIP from the SLB instance after the cluster is created. This allows you to change the access mode to the API server between Internet access and internal access.	All regions	Create an ACK dedicated cluster
Auto scaling of ENS nodes for ACK edge clusters	The auto scaling of ENS nodes in ACK edge clusters is supported. To support edge computing scenarios, ACK allows you to configure auto scaling of ENS nodes in ACK edge clusters. This feature can be implemented by calling the API.	All regions	Auto scaling of nodes
New region	ASK is available in the China (Zhangjiakou) region.	China (Zhangjiakou)	Create an ASK cluster

September 2019

Feature	Description	Region	References
New region	ACK is available in the China (Chengdu) region. You can create ACK dedicated clusters in the China (Chengdu) region. To create ACK managed clusters in the China (Chengdu) region, submit a ticket.	China (Chengdu)	Create an ACK dedicated cluster

Feature	Description	Region	References
Kubernetes 1.14.6 and new features for cluster upgrades	The canary release of the upgrade from Kubernetes 1.14.6 is implemented in the following regions: China (Shanghai), China (Zhangjiakou), Singapore (Singapore), and Germany (Frankfurt). Upgrades from Kubernetes 1.14.6 will soon be available in all regions. More features are also provided to simplify the upgrade process. In the ACK console, you can click **Upgrade Cluster** on the Clusters page to upgrade your cluster. The new upgrade feature adds the following improvements to secure upgrades: • A comprehensive cluster check is performed before an upgrade. • You can manually pause or resume an upgrade. • Detailed log of upgrades is retained.	• China (Shanghai) • China (Zhangjiakou) • Singapore (Singapore) • Germany (Frankfurt)	Update the Kubernetes version of an ACK cluster
Node maintenance	Node maintenance is supported. To maintain nodes in a cluster, you must make sure that workloads are not deployed on the nodes that you want to maintain. ACK supports node maintenance. You can select one or more nodes that you want to maintain and set them to unschedulable on the Nodes page. You can also drain these nodes. • After you set a node to unschedulable, pods cannot be scheduled to the node. • If you drain a node, no new pods are scheduled to the nodes and existing pods on the node are migrated to other nodes. However, pods that are managed by DaemonSets are not migrated from the node. If you have a LoadBalancer Service, you can specify whether to remove nodes that run the pods that are associated with the Service from the backend of the related SLB instance when these nodes are set to unschedulable. This allows you to flexibly manage your workloads during node maintenance.	All regions	Set node schedulability

Feature	Description	Region	References
Custom node names	Custom node names are supported. To manage a cluster that includes a large number of nodes, you must identify nodes by name. The default node names provided by ACK are not easy to identify. ACK allows you to customize node names when you create a cluster. When you create a cluster in the ACK console, you can select **Custom Node Name** in the advanced settings of the cluster. You can define a prefix, an IP substring length, and a suffix for a custom node name. The IP substring length specifies the number of digits to be truncated from the end of a node IP address and can be used to uniquely identify a node.	All regions	Create an ACK dedicated cluster
Advanced security groups	Advanced security groups are supported when you create an ACK clusters. Compared with basic security groups, advanced security groups support more ECS instances, more ENIs, and effective management on an infinite number of private IP addresses. Advanced security groups are suitable in scenarios that require high O&M efficiency, high ECS instance specifications, and a large number of compute nodes. To meet the requirements of a large-scale cluster, you can select advanced security groups for **Security Group** in the advanced settings when you create the cluster.	All regions	Create an ACK dedicated cluster
Disk encryption and CSI	Disk encryption and the CSI component are supported. ACK allows you to encrypt data disks. You can enable disk encryption for the selected data disks when you create a cluster. This feature can automatically encrypt the data that is transmitted from an ECS instance to a data disk and automatically decrypt the data when it is read. This improves data security. In addition, Kubernetes 1.14.6 supports the standard CSI plug-in, which is generally used for volume management. You can select FlexVolume or CSI when you create a cluster.	All regions	Create an ACK dedicated cluster and CSI overview

August 2019

Feature	Description	Region	References
Kubernetes 1.14.6	Kubernetes 1.14.6 is supported. You can select Kubernetes 1.14.6 when you create a cluster in the ACK console. You cannot upgrade an existing cluster to Kubernetes 1.14.6.	All regions	Kubernetes release notes

Feature	Description	Region	References
New regions	ASK is available in the Singapore (Singapore), China (Hong Kong), and Australia (Sydney) regions. ASK allows you to create containerized applications without managing or maintaining clusters and nodes. You are billed based on the actual amount of resources that are consumed by the elastic container instances that run the applications. ASK clusters allow you to focus on the design and development of applications, instead of managing the underlying infrastructures.	Singapore (Singapore) China (Hong Kong) Australia (Sydney)	Create an ASK cluster
ASK 2.0	ASK 2.0 is released to provide more Kubernetes-native features. ASK 2.0 supports multiple namespaces, CRD objects, RBAC, PVs, and PVCs. ASK 2.0 improves the security and isolation capability of clusters. The average price of ASK clusters is reduced by 46% due to lower costs of elastic container instances. This includes a 30% reduction in CPUs and a 65% reduction in memory.	All regions	Create an ASK cluster
SCC	ACK clusters based on Super Computing Cluster (SCC) resources are supported. SCCs are powered by ECS Bare Metal (EBM) instances and use the high-speed Remote Direct Memory Access (RDMA) technology. SCCs improve network performance. SCCs are used in scenarios such as high-performance computing, AI, machine learning, scientific and engineering computing, data analytics, and audio and video processing. You can create SCC-based ACK clusters. This type of cluster combines high-performance infrastructure resources with lightweight and agile containers. SCC-based ACK clusters are applicable to high network throughput and compute-intensive scenarios.	All regions	Create an ACK dedicated cluster
Auto scaling and cross-zone scheduling	Multiple scaling groups are supported for auto scaling. Cross-zone scheduling policies are supported. The auto scaling feature is optimized. You can configure multiple scaling groups so that resources of different specifications are automatically added when the scaling threshold is reached. This feature meets the requirements of running compute-intensive applications and GPU computing tasks. When you configure auto scaling policies, you can specify different scheduling policies for multiple zones, including priority policies, cost optimization policies, and zone balancing policies. This meets the requirement for resource scheduling when the cluster is deployed across multiple zones.	All regions	Auto scaling of nodes

Feature	Description	Region	References
Custom cluster domain names	Custom cluster domain names are supported. ACK allows you to customize a cluster domain name by specifying the cluster-domain parameter. The cluster-domain parameter specifies the local domain name that is used for service discovery. If you have multiple clusters, we recommend that you customize the local domain names to simplify the management of clusters and services. ACK allows you to customize a cluster domain name when you create a cluster. This simplifies management and improves the O&M efficiency.	All regions	Create an ACK dedicated cluster
App Hub	App Hub is provided in App Catalog. App Hub provides various cloud-native and open source containerized applications. ACK integrates App Hub into App Catalog. To deploy cloud-native applications in your cluster, log on to the ACK console and click the **App Hub** tab on the **App Catalog** page to find and install the applications with one click. This saves you the inconvenience of creating clusters and deploying applications by using a CLI.	All regions	App Marketplace

Feature	Description	Region	References
CCM	The CCM is upgraded. The CCM is the core component in an ACK cluster and is responsible for managing various cloud resources, such as SLB instances and VPCs. The following features are added to the CCM: • SLB instances can be created with access control settings. You can specify an IP whitelist for an SLB instance that is created by ACK. This enhances the security of the ACK cluster. • You can specify whether to remove unschedulable nodes when you run the kubectl cordon or kubectl drain command. Cordoning and draining nodes are important features in cluster maintenance. However, the community has not reached an agreement on whether to remove a node from the backend of an SLB instance when the node is set to unschedulable for maintenance. The CCM provides an interface that allows you to specify whether to remove such nodes from the backend of the SLB instance. This ensures the flexibility of maintenance. • Pods can be mounted to the backend of an SLB instance by using Terway. Terway ENI is the latest network plug-in that is provided by ACK. The core feature of Terway ENI is to mount the ENI IP address of a node to a pod. The CCM allows you to mount pods instead of nodes to the backend of an SLB instance. This prevents traffic forwarding through nodes and improves network performance. • Node weights can be set based on the number of pods on each node for Services in Local mode. The CCM can adjust the percentage of traffic that is sent to each node based on the number of pods on each node. This balances workloads among nodes. This feature applies to only Services in Local mode.	All regions	Cloud Controller Manager

July 2019

Feature	Description	Region	References
ACK edge clusters	ACK edge clusters are released for public preview. You can add edge nodes or ENS nodes to ACK edge clusters. This type of cluster supports edge computing and manages edge nodes and ENS nodes to reduce O&M costs. This type of cluster also supports autonomous edges and networks to meet the requirements in different edge computing scenarios. You can select this type of cluster on the cluster template page.	China site	-

Feature	Description	Region	References
Multi-cluster management	The multi-cluster management feature is released for public preview. You can select **Register Kubernetes Cluster** on the cluster template page to add Kubernetes clusters from data centers and other public clouds to the ACK console. Then, you can deploy applications to these clusters in the console. You can manage hybrid cloud clusters and clusters that are deployed across multiple clouds. After you add self-managed clusters from data centers to ACK, you can manage these clusters by using the O&M feature that is provided by ACK.	China site	Register an external Kubernetes cluster
New region	ACK managed clusters are available on the Alibaba Cloud Japan site. • Saves resources. You do not need to create master nodes in an ACK managed cluster. If you use another type of cluster, you must create at least three master nodes. • Improves O&M efficiency. ACK manages master nodes. • Ensures security. ACK meets various security requirements.	Japan	Create an ACK managed cluster
Support for multiple data disks during cluster creation	Multiple data disks can be mounted to nodes when you create an ACK cluster. This saves you the inconvenience of manually adding data disks after the cluster is created. ACK formats and mounts one of the selected data disks to the docker directory. You can determine how to handle the other data disks.	All regions	Create an ACK dedicated cluster
Existing security groups selectable during cluster creation	An existing security group can be selected when you create an ACK cluster. You can specify an existing security group for the VPC of your cluster in the advanced settings. This allows you to use custom inbound and outbound security group rules to improve the security of your cluster.	All regions	Create an ACK dedicated cluster

Feature	Description	Region	References
Deletion protection	Deletion protection is released to ensure the security of your cluster. You are required to enter a Short Message Service (SMS) verification code when you delete a cluster. However, you may mistakenly delete the cluster by calling the API. To ensure the security of clusters, ACK supports deletion protection for clusters. You can enable deletion protection when you create a cluster. This way, you cannot delete the cluster in the console or by calling the API. To delete the cluster, you must first disable deletion protection. You can enable or disable deletion protection on the cluster details page.	All regions	Create an ACK dedicated cluster
Batch authorization	Multiple RAM users can be authorized at the same time. You can also grant the permissions to manage all clusters. This allows you to efficiently authorize RAM users. The authorization procedure is also optimized to improve user experience.	All regions	Authorization overview
Time zone	The time zone of an application can be synchronized to that of the node. You can select **Synchronize Timezone from Node to Container** when you create an application from an image. This ensures that the application pods and the host node use the same time zone.	All regions	Create a stateless application by using a Deployment
New region	Container Registry Enterprise Edition is available in the UK (London) region. Container Registry Enterprise Edition supports large-scale image distribution with enhanced security. This service is suitable for enterprise users that require high security and large-scale nodes.	UK (London)	What is Container Registry?
Helm 2 charts supported by Container Registry Enterprise Edition	Helm 2 charts are supported by Container Registry Enterprise Edition to make it easier for you to manage cloud-native assets. You can enable the charts component on the Overview page of your Container Registry Enterprise Edition instance. When the component is running, you can start to manage Helm chart repositories.	All regions	N/A

June 2019

Feature	Description	Region	References
New regions	ACK managed clusters are available in the Japan (Tokyo) and UK (London) regions on Alibaba Cloud public cloud.	Japan (Tokyo) UK (London)	What is Container Service for Kubernetes?

Feature	Description	Region	References
Terway	A new version of Terway is released. The exclusive ENI mode and the inclusive ENI mode are supported by this version. The default mode is the inclusive ENI mode. • The exclusive ENI mode: In this mode, the number of pods that can be deployed on a node must match the number of ENIs that can be created on the node. This mode improves network performance. • The inclusive ENI mode: In this mode, you can deploy multiple pods on a node. The pods share the same ENI.	All regions	Work with Terway
Knative	Knative is supported. Knative is a Kubernetes-based serverless framework. Knative creates a cloud-native and cross-platform orchestration standard for serverless applications. Knative implements this standard by integrating the creation of containers (or functions), workload management (auto scaling), and event models. ACK supports Knative and allows you to install and upgrade the Build, Serving, and Eventing components. You must deploy Istio before you use Knative. ACK provides instructions to deploy sample applications, and also provides best practices for tracing, monitoring, and logging applications.	All regions	Overview and Use Knative to deploy serverless applications
Pod search	Pods can be searched by node IP address or pod IP address. In the ACK console, choose **Applications > Pods** and specify a node IP address or a pod IP address to search for a pod. This saves the time to find pods that you want to manage and maintain.	All regions	N/A

May 2019

Feature	Description	Region	References
New regions	ACK managed clusters are available in the Australia (Sydney) region on Alibaba Cloud public cloud and the China East 2 Finance region on Alibaba Finance Cloud. You can create ACK managed clusters in the Australia (Sydney) region on Alibaba Cloud public cloud and the China East 2 Finance region on Alibaba Finance Cloud.	Australia (Sydney) China East 2 Finance	What is Container Service for Kubernetes?

Feature	Description	Region	References
Genomics computing clusters for genomics computing	Genomics computing clusters are released. This type of cluster uses high-performance computing (HPC) instances as worker nodes and provides a large-scale workflow engine for batch genomics computing. Genomics computing clusters are suitable for data splitting and mutation detection, and support data analytics for the following formats: BCL, FASTQ, BAM, SAM, and VCF. In the ACK console, choose **Clusters > Clusters** and click Create Kubernetes Cluster. In the Select Cluster Template dialog box, select **Genomics Computing Cluster**.	All regions	N/A
ACK clusters with FPGA-accelerated nodes	ACK clusters with FPGA-accelerated nodes are released. This type of cluster uses FPGA F3 instances as worker nodes and is used for H265 video encoding and image conversion from JPEG to HEIF. FPGA-based video encoding reduces the processing time from more than 1 week to 15 minutes. This significantly reduces the bitrate and saves bandwidth costs when transcoding videos of the same quality. In the ACK console, choose **Clusters > Clusters** and click Create Kubernetes Cluster. In the Select Cluster Template dialog box, select **Dedicated FPGA Cluster** to create an ACK dedicated cluster with FPGA-accelerated nodes.	All regions	N/A
CCM	The CCM is updated to V1.9.3.110-g4938309-aliyun. This version supports more SLB configuration options. The following features are provided: • Allows you to restrict the creation of Internet-facing SLB instances by setting parameters. • Allows you to change certificate IDs. • Allows you to specify a vSwitch when you attach an internal-facing SLB instance to a Service. • Allows you to set SLB instance configurations to redirect traffic from HTTP port 80 to HTTPS port 443.	All regions	Cloud Controller Manager

Feature	Description	Region	References
Istio	Istio is updated to V1.1.4. Istio 1.1.4 improves self-recovery capabilities, and supports automatic recovery of the control plane and automatic upgrades of earlier versions. Istio is also integrated with Time Series Database (TSDB). TSDB is a database service that supports high-speed read and write operations, compressed storage, and real-time computing. To fix the local storage issues in Prometheus, TSDB provides remote storage services with high performance and high reliability at low costs. Compared with other remote storage solutions provided by the community, TSDB is easier to use and only requires you to change the Prometheus configuration. The solution supports parallel read and write operations and is highly compatible with PromQL. TSDB is a distributed storage system with auto scaling capabilities.	All regions	N/A
Container Registry Enterprise Edition	Images can be synchronized across all regions worldwide for instances of Container Registry Enterprise Edition. This solves issues in the global delivery of applications and improves the business iteration efficiency for enterprises. Container Registry Enterprise Edition supports large-scale image distribution with enhanced security. It is suitable for enterprises that require high security and a large number of nodes.	All regions	N/A
Support for multiple zones and five master nodes during cluster creation	Multiple zones and five master nodes are supported when you create an ACK dedicated cluster. This allows you to create a cross-zone ACK dedicated cluster with five master nodes to significantly improve the availability of the cluster.	All regions	N/A

April 2019

Feature	Description	Region	References
Kubernetes 1.12.6	ACK managed clusters or ACK dedicated clusters in all regions can be updated from Kubernetes 1.11.5 to 1.12.6 in the ACK console.	All regions	N/A

Feature	Description	Region	References
Audit logs	Audit logs can be collected from ACK managed clusters. An audit log records operations on the API server and allows cluster administrators to trace the activities of different users.	All regions	Use cluster auditing
Istio	Istio is upgraded to V1.1. Istio 1.1 allows you to manage Istio applications in the ACK console. You can create and manage Istio applications and services on a graphical interface. You can create different application versions, implement canary releases, set canary release policies, and also configure fault injection policies.	All regions	N/A
GPU-accelerated pods supported by ASK	GPU-accelerated pods are supported when you create applications in an ASK cluster. When you create an application from a template, specify the pod type as GPU in the YAML file.	All regions	N/A
Container Registry Enterprise Edition	Container Registry Enterprise Edition is available in the China (Beijing) region.	China (Beijing)	What is Container Registry?
ACK clusters with FPGA-accelerated nodes	ACK clusters with FPGA-accelerated nodes are released. This type of cluster uses FPGA F3 instances as worker nodes and is used for H265 video encoding and image conversion from JPEG to HEIF. FPGA-based video encoding reduces the processing time from more than 1 week to a short period of time. This significantly reduces the bitrate and reduces bandwidth costs when transcoding videos of the same quality. In the ACK console, choose **Clusters > Clusters** and click Create Kubernetes Cluster. In the Select Cluster Template dialog box, select **Dedicated FPGA Cluster** to create an ACK dedicated cluster with FPGA-accelerated nodes.	All regions	N/A

March 2019

Feature	Description	Region	References
New regions	ACK managed clusters are available in the China (Zhangjiakou), China (Hohhot), US (Silicon Valley), and Germany (Frankfurt) regions.	China (Zhangjiakou) China (Hohhot) Germany (Frankfurt) US (Silicon Valley)	What is Container Service for Kubernetes?

Feature	Description	Region	References
Container Registry Enterprise Edition	Container Registry Enterprise Edition was officially released at the Alibaba Cloud Summit on March 21, 2019. This edition provides higher security and supports large-scale image distribution. Container Registry Enterprise Edition is in public preview in the China (Shanghai) region. To use this edition, submit a ticket.	China (Shanghai)	What is Container Registry?
Container Registry Shared Edition	Container Registry Shared Edition is available in all regions on the International site (alibabacloud.com).	All regions	What is Container Registry?
Kubernetes 1.12.6	Kubernetes 1.12.6 is supported. You can create a cluster of Kubernetes 1.12 in the console.	All regions	Create an ACK dedicated cluster
Log Service	The Log Service plug-in is supported by ACK managed clusters. You can enable Log Service when you create an ACK managed cluster or an ACK dedicated cluster. After the plug-in is installed, you can use Log Service to manage Kubernetes log.	All regions	Create an ACK managed cluster
New region	ACK managed clusters that run Windows are available. You can create this type of cluster in the ACK console or by calling the API. This way, you can create Windows containers and deploy traditional Windows applications on cloud-native platforms to achieve agility and elasticity.	All regions	Windows clusters are no longer supported.
IPVS	The IP Virtual Server (IPVS) proxy mode is supported. Compared with the traditional iptables mode, the IPVS mode significantly improves the load balancing performance in large-scale clusters. You can use this mode in all clusters and all regions.	All regions	Create an ACK dedicated cluster
Cluster templates	Multiple cluster templates are provided in the console. You can select templates of different cluster types based on your business requirements. Templates of the following cluster types are supported: ACK managed clusters, clusters with EBM instances, GPU-accelerated clusters, and Windows clusters. Cluster templates allow you to create ACK clusters based on your business requirements.	All regions	N/A
Elastic Container Instance	High-specification elastic container instances are provided for genomics computing. The maximum CPU specification is increased from 8 vCPUs to 64 vCPUs. The highest specification of an elastic container instance is 64 vCPUs and 256 GiB memory. The lowest specification of an elastic container instance is 0.25 vCPU and 0.5 GiB memory. You can select a specification based on your business requirements to achieve the highest cost efficiency.	All regions	Limits

February 2019

Feature	Description	Region	References
New region	ACK managed clusters are available in the China (Shenzhen) region. ACK managed clusters provide the following core benefits: • Saves resources. You do not need to create master nodes in an ACK managed cluster. Compared with other cluster types, this cluster type saves you the costs of three master nodes. • Improves O&M efficiency. ACK manages the master nodes. • Ensures security. ACK meets various security requirements.	China (Shenzhen)	Create an ACK managed cluster
App Catalog	Knative add-ons are provided in App Catalog. Knative is a scale-to-zero and request-driven computing runtime based on Kubernetes and Istio. Knative supports the deployment of serverless applications and functions. ACK provides Knative add-ons to help you build the Knative Serving environment in your cluster.	All regions	Overview
Cluster inspection	Cluster inspection is supported. You can use this feature to perform in-depth checks on cluster resources, components, and configurations. This can identify the causes of errors in your cluster.	Chinese mainland	Use the global check feature to troubleshoot cluster issues

January 2019

Feature	Description	Region	References
Windows containers	Windows containers are supported. This allows you to deploy and run Windows applications in containers of ACK clusters. This enables Kubernetes-based elastic scheduling and management of Window applications. You can add Windows nodes to managed and ACK dedicated clusters. Container Registry Enterprise Edition is in private preview. To use this service, submit a ticket.	All regions	Create a Windows node pool

Feature	Description	Region	References
Container Registry Enterprise Edition	Container Registry Enterprise Edition is released for internal preview. Container Registry Enterprise Edition provides container image repositories built on top of dedicated resources. This edition provides stable image building, large-scale image distribution, and image hosting with enterprise-class security. It is suitable for enterprises that require high security and a large number of nodes. Container Registry Enterprise Edition is in private preview. To use this service, submit a ticket.	All regions	What is Container Registry?
Intelligent cluster O&M	Intelligent cluster O&M is available in the China (Hangzhou) region. Intelligent O&M provides the best practices for cluster management in different scenarios. This allows you to identify the causes of errors in the cluster by performing in-depth checks on cluster resources, components, and configurations.	China (Hangzhou)	Use the global check feature to troubleshoot cluster issues
ARMS	ARMS is supported and integrated into ACK. After you install the ARMS plug-in, you can monitor the application performance in your cluster. ARMS is a monitoring service for application performance management (APM). To monitor a Java application, you need only to attach an ARMS agent to the startup script of the application. No code change is required. ARMS enables you to locate failed API operations or slow calls, reproduce API parameters, detect memory leaks, and discover system bottlenecks. This significantly improves the efficiency of service diagnostics.	All regions	Monitor application performance
Elastic Container Instance	Starting January 22, 2019, you are charged for the commercial use of Elastic Container Instance. Elastic container instances are deployed as the underlying infrastructures of ASK cluster. You are charged when you create elastic container instances in ASK clusters. ASK clusters remain free of charge.	All regions	Billing
New regions	ASK clusters are available in the China (Beijing) and China (Shenzhen) regions. ASK clusters provide improved experience with serverless containers.	China (Beijing) China (Shenzhen)	Create an ASK cluster

December 2018

Feature	Description	Region	References

Feature	Description	Region	References
New region	ACK is available in the UK (London) region on both the China site (aliyun.com) and the International site (alibabacloud.com).	UK (London)	Create an ACK dedicated cluster
New regions	ACK managed clusters are available in the China (Shanghai), Malaysia (Kuala Lumpur), and India (Mumbai) regions on both the China site (aliyun.com) and the International site (alibabacloud.com).	China (Shanghai) Malaysia (Kuala Lumpur) India (Mumbai)	Create an ACK managed cluster
Node removal	Nodes can be removed from an ACK cluster. You can also choose whether to release the related ECS instances.	All regions	Remove a node
DaemonSet	DaemonSets are supported. DaemonSet is a daemon process that ensures that each node runs one replica of a pod.	All regions	N/A
Istio	Custom Istio Ingress and Egress gateways are supported by configuring different parameters.	All regions	ASM
Istio CoreDNS	Istio CoreDNS is supported. You can use the CoreDNS plug-in to read Istio service entries and associate the IP addresses of the services to their host addresses.	All regions	ASM
Existing ECS instances selectable during cluster creation	Existing ECS instances can be added as worker nodes when you create an ACK managed cluster.	All regions	Create an ACK managed cluster

November 2018

Feature	Description	Region	References
New region	ACK managed clusters are available in the Indonesia (Jakarta) region on the International site (alibabacloud.com).	Indonesia (Jakarta)	Create an ACK managed cluster
Terway	The Terway plug-in is released. Terway enables direct communication between containers through ENIs and provides higher network performance than Flannel.	All regions	Work with Terway
Thumbnail images for worker nodes	Thumbnail images are used to display the performance metrics of worker nodes, which makes it easy for you to view the status of nodes.	All regions	N/A

Feature	Description	Region	References
Node adding	Multiple existing nodes can be added to a cluster at the same time.	All regions	N/A
Rolling renewal of cluster certificates	Rolling renewal of cluster certificates is supported to prevent certificates from expiring.	All regions	N/A

October 2018

Feature	Description	Region	References
New region	ACK is available in the China South 1 Finance region on Alibaba Finance Cloud.	China South 1 Finance	Create an ACK dedicated cluster
New regions	N/A	Regions outside China	Create an ACK managed cluster
Deployment	Version management and rollback are supported for Deployments.	All regions	N/A
Istio	Istio is deeply integrated into ACK and Istio add-ons are supported.	All regions	N/A

September 2018

Feature	Description	Region	References
Kubernetes 1.11	• Kubernetes 1.11 is supported to provide features, such as CRD upgrade, CoreDNS general availability (GA), pod priority settings, and preemptive scheduling. • Multiple Kubernetes versions are supported, such as Kubernetes 1.10 and 1.11. • Multi-container applications and stateful applications are supported in the console.	All regions	Use a StatefulSet to create a stateful application
Container Registry	Images can be pulled from the private repositories of Container Registry without a password.	All regions	

Feature	Description	Region	References
Auto scaling	Auto scaling of nodes is supported. ACK provides the auto scaling component for nodes to automatically scale in and out. Regular instances, GPU-accelerated instances, and preemptible instances can be automatically added to or removed from an ACK cluster as required. This feature is applicable to instances that are deployed across multiple zones and diverse instance types, and also supports different scaling modes.	All regions	Auto scaling of nodes
Preemptible instances are supported.	N/A	All regions	

August 2018

Feature	Description	Region	References
ACK managed clusters	ACK managed clusters are released for public preview.	All regions	Create an ACK managed cluster
Istio	Istio add-ons are supported.	All regions	N/A

July 2018

Feature	Description	Region	References
New region	N/A	Australia (Sydney)	Create an ACK dedicated cluster
Canary releases and phased releases are supported.	N/A	All regions	N/A

June 2018

Feature	Description	Region	References
New regions	N/A	Japan (Tokyo) China (Hohhot)	Create an ACK dedicated cluster
FPGA and HugePages are supported by Kubernetes 1.10.	N/A	All regions	N/A
Application monitoring and alerting	Application monitoring and alerting are supported.	All regions	N/A

Feature	Description	Region	References
Subscription supported when you create an ACK cluster	N/A	All regions	Create an ACK dedicated cluster
Ingresses and the exec and attach commands supported	N/A	All regions	Features

May 2018

Feature	Description	Region	References
New region	ACK is available in the China East 2 Finance region on Alibaba Finance Cloud. Alibaba Finance Cloud provides services in compliance with security regulations.	China East 2 Finance	Create an ACK dedicated cluster
ASK released	N/A	All regions	Create an ASK cluster
Blue-green releases, canary releases, and A/B testing supported	N/A	All regions	N/A

April 2018

Description	Region	References
ACK is available in five regions in Southeast Asia, the Middle East, and India. Kubernetes 1.9 is stably supported.	Malaysia (Kuala Lumpur) Indonesia (Jakarta) Singapore (Singapore) India (Mumbai) UAE (Dubai)	Create an ACK dedicated cluster
MySQL, RDS, RabbitMQ, and Spark are supported in Service Catalog.	All regions	This feature is phased out.
Management of applications released by using Helm is supported in App Catalog.	All regions	Manage releases by using Helm

March 2018

Feature	Description	Region	References
Kubernetes 1.9	Kubernetes 1.9.3 is supported. ACK releases Workloads API. By default, CRD is enabled. GPU scheduling is supported. You can select custom ECS images when you create a cluster. You can also reset images when you add nodes to a cluster.	All regions	N/A
Helm	App Catalog is released to allow you to deploy applications by using Helm.	All regions	Manage releases by using Helm
ServiceBroker	App Catalog is released to support ServiceBroker.	All regions	This feature is phased out.
CloudMonitor	Nodes can be monitored by using CloudMonitor.	All regions	Monitor basic resources

January 2018

Feature	Description	Region	References
ACK and Container Registry released on the International site (alibabacloud.com)	N/A	Regions outside China	What is Container Service for Kubernetes?
Kubernetes 1.8.4 is supported to provide features such as security enhancement and auto scaling.	N/A	All regions	Auto scaling of nodes
FlexVolume	The FlexVolume plug-in is released to support disks, NAS file systems, and OSS buckets.	All regions	Disk volume overview, NAS volume overview, and OSS volume overview
Network policies and bandwidth throttling	Kubernetes network policies and bandwidth throttling are supported. This improves network performance.	All regions	Use annotations to configure load balancing
EBM instances supported	N/A	All regions	N/A

October 2017

Feature	Description	Region	References

Feature	Description	Region	References
Kubernetes 1.8.1	Kubernetes 1.8.1 is supported.	All regions	What is Container Service for Kubernetes?
Blockchain solutions released for public preview	N/A	All regions	N/A

August 2017

Feature	Description	Region	References
Kubernetes 1.7.2	N/A	All regions	Create an ACK dedicated cluster

2.Release notes for Kubernetes versions supported by ACK

2.1. Overview of Kubernetes versions supported by ACK

This topic lists the Kubernetes versions that are supported by Container Service for Kubernetes (ACK).

The following Kubernetes versions are supported by ACK:

- Kubernetes 1.22 release notes
- Kubernetes 1.20 release notes
- Kubernetes 1.18 release notes
- Kubernetes 1.16 release notes
- Kubernetes 1.12 release notes

Alibaba Cloud periodically updates the list of Kubernetes versions that are supported by ACK. For more information about how ACK supports different Kubernetes versions, see ACK releases of Kubernetes.

2.2. Kubernetes 1.22 release notes

Container Service for Kubernetes (ACK) strictly abides by the terms of the Certified Kubernetes Conformance Program. This topic describes the changes that ACK has made to support Kubernetes 1.22.

Version upgrades

Components are upgraded and optimized by ACK to support Kubernetes 1.22.

Key component	Version	Description
		• A number of beta API versions are discontinued in Kubernetes 1.22. Before you upgrade to Kubernetes 1.22, take note of the following items: ◦ Existing resources that are created by using the beta API versions are not affected. After you upgrade to Kubernetes 1.22, you can use stable API versions to manage the resources.

Key component	Version	Description
Kubernetes	1.22.3-aliyun.1	◦ Before you upgrade to Kubernetes 1.22, you must recreate the controllers and applications that are managed by using the beta API versions. Otherwise, the controllers and applications cannot work as normal after you upgrade to Kubernetes 1.22. For more information, see Version details. • A number of optimizations are added in Kubernetes 1.22. For more information, see Version details. • Dockershim was deprecated in Kubernetes 1.20 and later versions, and will be removed in Kubernetes 1.24. ACK provides the following suggestions to help you handle this issue: ◦ We recommend that you use containerd as the container runtime for new nodes. ◦ containerd is not fully compatible with the Windows operating system. For Windows containers, we recommend that you use Docker Enterprise Edition (EE). ◦ Before you upgrade to Kubernetes 1.24, we recommend that you migrate workloads that run in Docker containers to containers that run other container runtimes. For more information, see Version details. • The PodSecurityPolicy (PSP) resource was deprecated in Kubernetes 1.21 and later versions, and will be removed in Kubernetes 1.25. You can use continue to use PSPs in Kubernetes 1.22. We recommend that you use the PodSecurity admission controller as an alternative. For more information, see Version details.

Key component	Version	Description
etcd	3.5.1	N/A
CoreDNS	1.8.4.1-3a376cc-aliyun	The upgrade does not affect your workloads. The following features are provided: • EndpointSlices can be monitored. • IPv6 addresses are supported by DNS resolutions.
CRI	• Docker CE 19.03.15 • Docker EE is supported for the Windows operating system.	N/A
	ContainerD 1.4.8	N/A
CSI	1.20.7-aafce42-aliyun	N/A
CNI	Flannel 0.15.1.4-e02c8f12-aliyun	The upgrade does not affect your workloads. The following features are provided: • The API version of resources such as Authorization is updated to support Kubernetes 1.22. • Services can be exposed by using the HostPort method. • The hairpin mode can be enabled.
	Terway	N/A
NVIDIA Container Runtime	3.7.0	N/A

Key component	Version	Description
Ingress Controller	1.1.0-aliyun.1	• NGINX Ingress controller 0.44.0 and earlier versions use Ingresses of the v1beta1 API version and therefore cannot run as normal in clusters of Kubernetes 1.22. To resolve this issue, upgrade the NGINX Ingress controller to 1.1.0 for a cluster of Kubernetes 1.20 before you upgrade the cluster to Kubernetes 1.22. • The upgrade may temporarily interrupt your workloads and cause compatibility issues with your workload configurations. We recommend that you evaluate the impact of the component upgrade before you upgrade to Kubernetes 1.22.

Version details

Resource changes and deprecation

- The `admissionregisration.k8s.io/v1beta1` API version of the MutatingWebhookConfiguration and ValidatingWebhookConfiguration resources is discontinued. Admission webhook configurations and mutating webhook configurations cannot be created by using this API version, which adversely affects the use of admission webhooks and mutating webhooks. You can use the `admissionregistration.k8s.io/v1` API version instead.
- The `apiextensions.k8s.io/v1beta1` API version of the CustomResourceDefinition (CRD) resource is discontinued. CRDs cannot be created by using this API version, which adversely affects the reconciliation of controllers that use CRDs. You can use the `apiextensions.k8s.io/v1` API version instead.
- The `apiregistration.k8s.io/v1beta1` API version of the APIService resource is discontinued. Extended Kubernetes APIs that are managed by using this API version cannot be used. You can use the `apiregistration.k8s.io/v1` API version instead.
- The `authentication.k8s.io/v1beta1` API version of the TokenReview resource is discontinued. TokenReviews that are created by using this API version cannot be used for authentication, which adversely affects your applications. You can use the `authentication.k8s.io/v1` API version instead.
- The `authorization.k8s.io/v1beta1` API version of the SubjectAccessReview resource is discontinued. SubjectAccessReviews that are created by using this API version cannot be used for authorization, which adversely affects your applications. You can use the `authorization.k8s.io/v1` API version instead.
- The `certificate.k8s.io/v1beta1` API version of the CertificateSigningRequest (CSR) resource is discontinued. CSRs that are created by using this API version cannot be used to apply for certificate signing and issuing. You can use the `certificate.k8s.io/v1` API version instead.

- The `coordination.k8s.io/v1beta1` API version of the Lease resource is discontinued. Leases that are created by using this API version cannot be used for leader election, which adversely affects your applications. You can use the `coordination.k8s.io/v1` API version instead.
- The `networking.k8s.io/v1beta1` and `extensions/v1beta1` API versions of the Ingress and IngressClass resources are discontinued. Ingresses that are created by using these API versions cannot be used to expose Services. You can use the `networking.k8s.io/v1` API version instead.
- The `rbac.authorization.k8s.io/v1beta1` API version of the ClusterRole, ClusterRoleBinding, Role, and RoleBinding resources is discontinued. Role-based access control (RBAC) resources that are managed by using this API version cannot be used to grant the permissions to manage applications and clusters. You can use the `rbac.authorization.k8s.io/v1` API version instead.
- The `storage.k8s.io/v1beta1` API version of the CSIDriver, CSINode, StorageClass, and VolumeAttachment resources is discontinued. If you use this API version to manage resources that are related to the Container Storage Interface (CSI) plug-in, the CSI plug-in may not run as normal and storage services in your cluster are adversely affected. You can use `storage.k8s.io/v1` instead.
- The `scheduling.k8s.io/v1beta1` API version of the PriorityClass resource is discontinued. PriorityClasses that are managed by using this API version cannot be used to configure pod priorities. You can use the `scheduling.k8s.io/v1` version instead.
- Dockershim is deprecated and will be removed in Kubernetes 1.24. For more information, see EP-2221 and cri-containerd.

  Before you upgrade to Kubernetes 1.24, we recommend that you perform the following steps to migrate workloads that run in Docker containers to containers that run other container runtimes:

  - Decide the node specifications and calculate the number of nodes that run container runtimes other than Docker based on the number of existing Docker containers.
  - Add new nodes to your cluster during off-peak hours.
  - Drain nodes that run the Docker runtime one after one. Each time a node is drained, verify that application pods on the node are successfully migrated to new nodes before you drain another node.
  - After all the nodes that run the Docker runtime are drained and no pod runs on the nodes, remove the nodes.

Feature enhancements

- By default, the ImmutableEphemeralVolumes feature is enabled in Kubernetes 1.21 and later versions. You can use this feature to set ConfigMaps and Secrets as immutable, which significantly reduces the load on the Kubernetes API server of your cluster. For more information, see Secrets and ConfigMaps.
- By default, the IPv6DualStack feature is enabled in Kubernetes 1.21 and later versions. To use IPv4/IPv6 dual stack, you must specify proper IPv4 CIDR blocks and IPv6 CIDR blocks when you create a cluster, and install a Container Network Interface (CNI) plug-in that supports IPv4/IPv6 dual stack. For more information, see IPv4/IPv6 dual stack.
- By default, the GracefulNodeShutdown feature is enabled in Kubernetes 1.21 and later versions. This feature supports only Linux nodes. After this feature is enabled, kubelet is aware of node shutdown events that are about to take place and can evict the pods on a node within a specific shutdown period. For more information, see Graceful node shutdown.
- By default, the EfficientWatchResumption feature is enabled in Kubernetes 1.21 and later versions. This feature can resume the watch cache of the Kubernetes API server in an efficient manner after the API server is restarted. This feature is suitable for large-scale clusters. For more information, see KEP-1904.

- By default, the CSIStorageCapacity feature is enabled in Kubernetes 1.22 and later versions. This feature enables kube-scheduler to schedule a pod to a node whose storage capacity is sufficient for creating the volume that is used by the pod. For more information, see Storage capacity.
- By default, the DaemonSetUpdateSurge feature is enabled in Kubernetes 1.22 and later versions. This feature allows you to use the `.spec.strategy.rollingUpdate.maxSurge` field to specify the percentage of pods that can be created above the expected number of pods during a rolling update on a DaemonSet. For more information, see Deployments.
- By default, the IndexedJob feature is enabled in Kubernetes 1.22 and later versions. This feature allows you to create an indexed Job by setting .spec.completionMode to *Indexed* in the Job configuration. This way, the annotation batch.kubernetes.io/job-completion-index and the JOB_COMPLETION_INDEX environment variable are added to each pod that is created by the Job. For more information, see Kubernetes.
- By default, the MemoryManager feature is enabled in Kubernetes 1.22 and later versions. This feature supports only Linux nodes. You can use this feature to enable non-uniform memory access (NUMA)-aware memory management. This feature is suitable for applications that require guaranteed memory resources to significantly improve application performance. ACK does not configure memory reservation for this feature. For more information, see Memory maps at runtime and Utilize the NUMA-aware memory manager.
- By default, the PodAffinityNamespaceSelector feature is enabled in Kubernetes 1.22 and later versions. This feature allows you to apply label selectors of pod affinity settings across namespaces instead of within the same namespace. This optimizes affinity-based pod scheduling. For more information, see KEP-2249.
- By default, the PodDeletionCost feature is enabled in Kubernetes 1.22 and later versions. After this feature is enabled, pods with lower resource utilization incur lower pod deletion costs. For more information, see ReplicaSet.
- By default, the PreferNominatedNode is enabled in Kubernetes 1.22 and later versions. After this feature is enabled, kube-scheduler preferably schedules pods to nominated nodes. kube-scheduler evaluates the other nodes only if all nominated nodes fail to match the pods. For more information, see KEP-1923.
- The ProbeTerminationGracePeriod feature is enabled in Kubernetes 1.22 and later versions. This feature supports only liveness probes. This feature allows you to set a probe-level or pod-level teminationGracePeriodSeconds field to shorten the time period that a pod must wait to restart after the pod fails a liveness probe. For more information, see Configure liveness, readiness, and startup probes.
- By default, the NetworkPolicyEndPort feature is enabled in Kubernetes 1.22 and later versions. This feature allows you to specify a port range in a NetworkPolicy. For more information, see Network policies.
- By default, the LogarithmicScaleDown feature is enabled in Kubernetes 1.22 and later versions. This feature provides a randomized approach to scale in pods and therefore reduces the impact of issues caused by pod topology spread constraints. For more information, see Pod topology spread constraints should be taken into account on scale down and KEP-2185.
- By default, the SuspendJob feature is enabled in Kubernetes 1.22 and later versions. This feature allows users to manage the lifecycle of Jobs in a more efficient manner. For example, you can use this feature to suspend and resume Jobs. For more information, see Introduce suspended Jobs.
- By default, the ServiceInternalTrafficPolicy feature is enabled in Kubernetes 1.22 and later versions. You can use this feature to route internal traffic to node-local endpoints that are ready or all endpoints that are ready in the cluster. For more information, see Services.
- By default, the ServiceLoadBalancerClass feature is enabled in Kubernetes 1.22 and later versions.

You can use this feature to customize load balancing. For more information, see Specify the class of load balancer implementation.

- By default, the ServiceLBNodePortControl feature is enabled in Kubernetes 1.22 and later versions. This feature allows you to disable node port allocation for a LoadBalancer Service by setting .spec.allocateLoadBalancerNodePorts to *false* in the Service configuration. This way, the Service routes traffic directly to pods. For more information, see Disable load balancer NodePort allocation.
- By default, the SizeMemoryBackedVolumes feature is enabled in Kubernetes 1.22 and later versions. This feature supports only Linux nodes. You can use this feature to specify the size of an emptyDir memory-backed volume by setting the emptyDir.sizeLimit field. This improves the observability of pod scheduling. For more information, see KEP-1967.
- By default, the Server-side Apply feature is enabled in Kubernetes 1.22 and later versions. This feature allows you to track changes to the fields of a resource configuration. You can track information about the change, such as the source, time, and operation. For more information, see Server-side apply.
- The feature of integrating the CSI plug-in with Windows containers is stabilized in Kubernetes 1.22 and later versions. This feature allows you to use CSI Proxy to perform storage operations on the host whose operating system does not support privileged containers, such as Windows Server 2019 and Windows Server version 2004. To use this feature, make sure that the CSI plug-in that you use supports this feature. For more information, see CSI Proxy.
- By default, the CSRDuration feature is enabled in Kubernetes 1.22 and later versions. After this feature is enabled, the validity period of a certificate to be signed and issued is set to the smaller value between the value of .spec.expirationSeconds in the CSR and the value of *--cluster-signing-duration* in the kube-controller-manager configuration. In ACK clusters, the default value of --cluster-signing-duration in the kube-controller-manager configuration is 10 years. For more information, see Signers.

New features

- The volume health monitoring feature is supported in Kubernetes 1.21 and later versions. This feature helps detect the health status of persistent volumes (PVs) that are provisioned by using the CSI plug-in. This prevents data from being read from or written to unhealthy PVs. By default, this feature is enabled for ACK clusters that use the CSI plug-in. To use this feature, make sure that the CSI plug-in that you use supports this feature. For more information, see Volume health monitoring.
- The memory Quality of service (QoS) feature that is developed based on cgroups v2 is supported in Kubernetes 1.22 and later versions. In situations where computing resources are insufficient, for example, resource request spikes, CPU throttling is performed to ensure the availability of CPU resources. However, memory throttling is not supported. To support memory throttling, open source Linux kernel optimizes specific interfaces in cgroups v2. By default, the memory QoS feature is enabled for ACK clusters. This feature supports only Linux nodes. To use this feature, make sure that the OS kernels of the Linux nodes that you use support this feature. For more information, see Memcg QoS feature of the cgroup v1 interface and 2570-memory-qos.
- Windows privileged containers can be created from HostProcess containers in Kubernetes 1.22 and later versions. By default, the Windows HostProcess container feature is enabled for ACK clusters. To use this feature, make sure that the OS kernels of the nodes that you use support this feature. For more information, see What's new for Windows containers on Windows Server 2022 and Create a Windows HostProcess Pod.
- The swap memory feature is supported for workloads in Kubernetes 1.22 and later versions. This feature supports only Linux nodes. For scenarios in which the swap memory feature is required, you can use the swap memory feature to improve the performance of your application. For example, a node administrator wants to improve node performance or reduce stability issues that are caused by

memory contention. The swap memory feature is disabled for ACK clusters. For more information, see Swap memory management and KEP-2400.

- Default seccomp profiles are configured for workloads in Kubernetes 1.22 and later versions. This feature supports only Linux nodes. After this feature is enabled, the RuntimeDefault seccomp profile is used by default. Specific workloads may require fewer limits on system calls than other workloads. These workloads may fail after this feature is enabled. This feature is disabled for ACK clusters. For more information, see Enable the use of RuntimeDefault as the default seccomp profile for all workloads.

Feature updates

- The PSP resource was deprecated in Kubernetes 1.21 and later versions, and will be removed in Kubernetes 1.25. By default, the pod security policy feature is enabled for ACK clusters. You can use ACK pod security policies as an alternative to the PSP resource in Kubernetes 1.22. For more information, see Pod security admission and PodSecurityPolicy deprecation: past, present, and future.
- The topologyKeys field was deprecated in Kubernetes 1.21 and later versions. Instead, the Topology Aware Hints feature is used to enable the Service topology feature. By default, the Service topology feature is disabled for ACK clusters. If the Service topology feature is enabled for a cluster of Kubernetes 1.22, you can enable the Topology Aware Hints feature to achieve the same effect as the topologyKeys field. For more information, see Topology Aware Hints.

Enhancements to Kubernetes 1.22

Observability

- More metrics about the access and requests to the Kubernetes API server are added. This improves the observability of the Kubernetes API server.
- Key metrics of control plane components can be collected for ACK Pro clusters, serverless Kubernetes (ASK) Pro clusters, and ACK edge Pro clusters. This improves the observability of control plane components.

Stability

The following enhancements are provided for all types of ACK clusters:

- Improve the protection for storage resources to reduce the load on etcd during cold starts.
- Traffic throttling can be performed on the Kubernetes API server based on the combination of the sources, types, and routes of requests. This reduces the load on etcd during cold starts.

Performance optimizations

- kubelet: During the in-place upgrade of kubelet, the system prevents pod restarts with the best effort. For more information, see kubelet's calculation of whether a container has changed can cause cluster-wide outages.
- kube-proxy: kube-proxy is compatible with Alibaba Cloud Linux 2 (kernel-4.19.91-23) and later versions. When the IP Virtual Server (IPVS) mode is enabled, conn_reuse_mode is not set to *0*. For more information, see Set conn_reuse_mode=1 on linux kernel version >= v5.9.
- ASK clusters: ASK clusters do not proactively evict Elastic Container Instance-based pods if virtual nodes are not ready. This reduces business losses.
- ACK Pro and ACK edge Pro clusters: The scheduler is optimized. Scheduling features are improved, including gang scheduling, topology-aware CPU scheduling, and topology-aware GPU scheduling. For more information, see Introduction to professional managed Kubernetes clusters.

References

- CHANGELOG-1.21.md
- CHANGELOG-1.22.md

2.3. Kubernetes 1.20 release notes

Container Service for Kubernetes (ACK) strictly conforms to the terms of the Certified Kubernetes Conformance Program. This topic lists the changes that ACK has made to support Kubernetes 1.20.

Version upgrades

All ACK components have been upgraded and optimized to support Kubernetes 1.20.

Core component	Version	Upgrade notes
Kubernetes	1.20.11	• Before you upgrade an ACK cluster to Kubernetes 1.20 or later, make sure that the required subject alternative names (SANs) are included in the self-signed server certificates of the admission webhooks in the cluster. For more information, see the sample Helm chart. • The selfLink field is deprecated. For more information, see Stop setting SelfLink in kube-apiserver. • By default, the `node-role.kubernetes.io/control-plane` label is added by ACK to the master nodes of a dedicated Kubernetes cluster. The `node-role.kubernetes.io/master` label is deprecated in Kubernetes versions later than 1.20.
Docker Runtime	19.03.5	None
Containerd Runtime	1.4.4	None
etcd	3.4.3	None
CoreDNS	1.7.0	• The deprecated upstream plug-in is no longer compatible. If the upstream plug-in is specified in the *Corefile* configurations, it will be automatically deleted in a secure way when CoreDNS is upgraded. • The names of metrics are updated. If your monitoring system is reliant on CoreDNS metrics, you must update the metric names. For more information, see Metric changes.
NVIDIA Container Runtime	3.4.1	None

Version details

Resource changes and deprecation

- The Docker runtime is deprecated. The Docker runtime is marked as deprecated in Kubernetes 1.20. However, you can continue using the Docker runtime in your clusters. The Docker runtime will not be

supported by open source Kubernetes in later versions. This change does not affect container images. You can still build Docker images. For more information, see Dockershim Deprecation FAQ.

- By default, the `node-role.kubernetes.io/control-plane` label is added by ACK to the master nodes of a dedicated Kubernetes cluster. The `node-role.kubernetes.io/master` label is deprecated in Kubernetes versions later than 1.20.
- The selfLink field is deprecated. For more information, see Stop setting SelfLink in kube-apiserver.
- The `extensions/v1beta1` and `networking.k8s.io/v1beta1` API versions are no longer used to manage Ingresses and IngressClasses, and will be deprecated in Kubernetes versions later than 1.22. Use `networking.k8s.io/v1` instead.

  > **Note** By default, the NGINX Ingress controller is installed in ACK clusters. This component enables you to use the networking.k8s.io/v1beta1 API version to manage Ingresses and IngressClasses.

- The required SANs must be included in the self-signed server certificates of the admission webhooks in ACK clusters. Before you upgrade an ACK cluster to Kubernetes 1.20 or later, make sure that the required SANs are included in the self-signed server certificates of the admission webhooks in the cluster. For more information, see the sample Helm chart.

Feature upgrades

- The issue that exec probes do not time out based on the timeout settings is fixed for kubelet. The default timeout period for exec probes is now 1 second, which may be short for some exec probes. If the timeout period is not specified for exec probes, we recommend that you specify the default timeout period.
- The API Priority and Fairness feature (APF) is a feature of Kubernetes in public preview and is enabled by default. You can use this feature to limit and prioritize requests. For more information, see API Priority and Fairness.
- By default, the EndpointSlice feature is enabled. In Kubernetes 1.19 and later, the EndpointSlice feature is automatically enabled by kube-proxy to support large-scale clusters. For more information, see EndpointSlices.
- Immutable ConfigMaps and Secrets are supported. The immutable ConfigMaps and Secrets feature is in public preview. If a ConfigMap or Secret is set to immutable, it cannot be modified. This reduces the load on kube-apiserver. For more information, see Immutable ConfigMaps.

Enhancements to Kubernetes 1.20

Control plane improvements

- Observability. Metrics are collected to monitor request operations and watch operations. This improves the observabilities of control plane components.
- Stability. Protection is provided to defend etcd against excessive requests when a cluster is started. This improves system stability.
- Performance optimizations. Indexes are added to accelerate the processing of list requests. This reduces the CPU usage of kube-apiserver.

Enhancements of Windows containers

- By default, the EndpointSlice feature is enabled.
- Device plug-ins are supported. For more information, see Device plug-ins.

Performance optimizations

In Kubernetes 1.20.11, KubeProxy is compatible with Alibaba Cloud Linux 2 whose kernel version is 4.19.91-23 or later. If you enable the IPVS mode, `conn_reuse_mode` is not set to 0. For more information, see IPVS.

References

- CHANGELOG-1.20.md
- CHANGELOG-1.19.md

2.4. Kubernetes 1.18 release notes

Container Service for Kubernetes (ACK) strictly abides by the terms of the Certified Kubernetes Conformance Program. This topic lists the changes that ACK has made to support Kubernetes 1.18.

Version upgrade

All ACK components have been upgraded and optimized to support Kubernetes 1.18.8.

Core component	Version	Upgrade notes
Kubernetes	1.18.8	Some commonly used API versions are deprecated in Kubernetes 1.18. Before you upgrade a cluster, we recommend that you upgrade the deprecated API versions that are listed in this topic.
Docker	19.03.5 (containerd 1.2.10)	None
etcd	3.4.3	None
CoreDNS	1.6.7	None

Version details

- **Resource changes and deprecation**

  The following API versions are deprecated in Kubernetes 1.18:

  - The apps/v1beta1 and apps/v1beta2 API versions of all resources are replaced by **apps/v1**.
  - The extensions/v1beta1 API version of DaemonSets, Deployments, and ReplicaSets is replaced by **apps/v1**.
  - The extensions/v1beta1 API version of NetworkPolicies is replaced by **networking.k8s.io/v1**.
  - The extensions/v1beta1 API version of PodSecurityPolicies is replaced by **policy/v1beta1**.

  The label that specifies the region of a node is changed to topology.kubernetes.io/region. The label that specifies the zone of a node is changed to topology.kubernetes.io/zone. We recommend that you update your workload configurations.

- **Feature enhancements**

  - Server-side Apply Beta 2 is introduced. You can view the relationships between the configuration items of a resource in the metadata.managedFields field of the metadata of the resource.

- The NodeLocal DNSCache feature is released to improve the DNS availability and performance of your cluster.
- The Volume Snapshot feature is in the public preview phase and supports operations such as volume backup, restoration, and scheduled backup.

Enhancements to Kubernetes 1.18.8

For Kubernetes 1.18.8, ACK enables the following feature for kubelet: Users who use raw data volumes can upgrade clusters on the fly without the need to drain nodes.

References

- CHANGELOG-1.18.md
- CHANGELOG-1.17.md

2.5. Kubernetes 1.16 release notes

Container Service for Kubernetes (ACK) strictly abides by the terms of the Certified Kubernetes Conformance Program. This topic lists the changes that ACK has made to support Kubernetes 1.16.

Version upgrade

All ACK components have been upgraded and optimized to support Kubernetes 1.16.

Core component	Version	Upgrade notes
Kubernetes	1.16.9	Vulnerability CVE-2020-8555 is fixed in Kubernetes 1.16.9. For more information, see Vulnerability fixed: CVE-2020-8555 in kube-controller-manager.
	1.16.6	In Kubernetes 1.16, the built-in CoreDNS version is V1.6.2. Compared with CoreDNS 1.3.1 in Kubernetes 1.14. The new version has the following changes: • The proxy plug-in is replaced with the forward plug-in, which offers higher performance. • By default, the ready plug-in is enabled. It is used to check the readiness of containers. The Corefile will be automatically migrated to match the new CoreDNS version.
Docker	19.03.5 (containerd 1.2.10)	None
etcd	3.4.3	None

Version details

- **Performance optimizations**

  Compared with Kubernetes 1.14, Kubernetes 1.16.6 has the following performance optimizations:

  - Optimizes PodAffinity to improve performance by about 100%.

- Optimizes serialization operations. Improves the performance of list operations on pods by 40%. Improves the performance of list operations on nodes by 30%.
- Improves the performance of processing apply requests that involve large map objects on the server side.
- Improves the heartbeat solution based on node leases. Reduces the number of lease queries per minute that are sent to the API server or etcd by 50,000 in a cluster of 8,000 nodes.
- Dramatically speeds up the pod creation process. When it comes to creating stateless pods, which does not involve mounting volumes such as ConfigMaps or Secrets to the pods:
  - Both Kubernetes 1.16.6 and 1.14 meet the SLAs defined by Special Interest Group (SIG) Scalability. 99% of pods can be started within 5 seconds given that images are already pulled.
  - Based on the statistics of the 1% of pod creation processes that require the longest time, Kubernetes 1.14 requires nearly 5 seconds to create a pod whereas Kubernetes 1.16.6 requires only 3 seconds under the same conditions.

Compared with previous versions, Docker 19.03.5 has the following optimizations:

- Adds the built-in buildkit to speed up image builds.
- Optimizes the systemd detection logic for the runC command-line tool. Containers start faster and occupy less memory.

Docker 19.03.5 has the following improvements on runtime stability:

- Fixes the issue where pods occasionally restart when health checks are performed using exec probes.
- Fixes vulnerability CVE-2018-15664, which is exposed by the **docker cp** command.
- Fixes the issue where Docker does not respond when a rich container that runs multiple processes exits.
- Fixes the handle leak issue in containerd.

- **Feature enhancements**

  Compared with Kubernetes 1.14, Kubernetes 1.16.6 has the following important changes.

  - By default, the following API versions are not supported: extensions/v1beta1, apps/v1beta1, and apps/v1beta2. apps/v1beta1 is replaced by apps/v1. The API version used by resources defined in apps/v1beta1 is replaced by apps/v1. The API version used by DaemonSets, Deployments, and ReplicaSets defined in extensions/v1beta1 is replaced by apps/v1. The API version used by NetworkPolicies defined in extensions/v1beta1 is replaced by networking.k8s.io/v1.

    > **Note** To ensure compatibility with your workloads, ACK has added support for the preceding API versions in Kubernetes 1.16.6 and will end the support in Kubernetes 1.18. We recommend that you change the API versions as soon as possible.

  - The following kubelet security control parameters are deprecated and removed: AllowPrivileged, HostNetworkSources, HostPIDSources, and HostIPCSources. Access control parameters such as PodSecurityPolicy are added for enhanced security.
  - More features have been stabilized. For example, CustomResourceDefinitions (CRDs) and admission webhooks are in general availability.

Enhancements to Kubernetes 1.16.6

ACK has enhanced Kubernetes 1.16 in the following aspects:

- Enhances stability and performance.
  - Adds retries for idempotent functions to improve the success rate of cluster creation.
  - Running containers are not restarted during kubelet upgrades.
  - Fixes kubelet startup failures caused by hugetlb.
- Improves observability
  - Optimizes the log of liveness probes that are sent from Server Load Balancer (SLB) instances to the API server.
  - Adjusts the log level of aggregationcontroller.
  - Optimizes the output of the **get cs** command in managed Kubernetes clusters.
  - Optimizes monitoring metrics on sandboxed containers based on compatibility with existing metrics APIs.

References

- CHANGELOG-1.16.md
- CHANGELOG-1.15.md

2.6. Kubernetes 1.12 release notes

The Kubernetes version used by Container Service for Kubernetes (ACK) is fully compatible with the Kubernetes open source version and shares the core code with the Kubernetes open source version. This topic lists the changes that ACK has made to support Kubernetes 1.12.

1.12.6-aliyun.1

- The log of Transport Layer Security (TLS) handshake errors between the API server and kubelet is blocked if the log is generated during health checks for Server Load Balancer (SLB) instances.

  Commit ID: 4f1d96e153b050d8374bfbb66803d7b3d9181abe
- Version verification on Docker 18.09.2 can be performed by kubeadm.

  Commit ID: 3b1ebfa1b857c44f5261a36f1420b10a08e01771
- The level of the log that can be watched for aggregationcontroller is adjusted.

  Commit ID: 01a904eed3f9486caa482c8983698075d1cea2f1
- Kubeadm
  - Retries for interactions between kubeadm and the API server are automatically performed by kubeadm when cluster resources are updated.
  - Domain Name System (DNS) servers are no longer deployed by kubeadm.
  - kube-proxy is no longer deployed by kubeadm.
  - The validity period of a certificate generated by kubeadm is changed to 10 years.

For more information, see Kubernetes 1.12.6 release notes.

References

- CHANGELOG-1.12.md
- CHANGELOG-1.11.md

3.Release notes for OS images

This topic provides the release notes for the OS images supported by Container Service for Kubernetes (ACK).

Usage notes

ACK supports the Alibaba Cloud Linux 2, Alibaba Cloud Linux 3, CentOS 7.X, and ContainerOS operating systems. Alibaba Cloud Linux 2 is maintained by Alibaba Cloud. Alibaba Cloud provides long-term technical support and patching for vulnerabilities related to Alibaba Cloud Linux 2. Alibaba Cloud Linux 2 is compatible with CentOS 7 and provides more features, improved performance, and higher stability. We recommend that you install Alibaba Cloud Linux 2 on the nodes of your ACK cluster. For more information about Alibaba Cloud Linux 2, see Use Alibaba Cloud Linux 2.

When you install Alibaba Cloud Linux 3, make sure that the following components are installed and the version of your cluster meets the following requirement.

Component or cluster	Minimum version
Cluster	1.20.4
ACK NodeLocal DNSCache	1.5.0
Flannel	v0.13.0.1-466064b-aliyun
Terway	v1.0.10.390-g5f3c461-aliyun

Release notes

December 2021

OS	OS version	Kernel version	Release date	Description
Alibaba Cloud Linux 3.2104	aliyun_3_x64_20G_alibase_20210425.vhd	5.10.23-5.al8.x86_64	2021-12-20	For more information, see Release notes for Alibaba Cloud Linux 3.

November 2021

OS	OS version	Kernel version	Release date	Description
ContainerOS 1.20.11	lifsea_3_x64_10G_k8s_1_20_11_alibase_20211103.vhd	5.10.23-5.8.lifsea8.x86_64	2021-11-06	Key component versions: • Kubelet：1.20.11 • Containerd: 1.4.8

October 2021

OS	OS version	Kernel version	Release date	Description

ContainerOS 1.20.4	lifsea_3_x64_10G_k8s_1_20_4_alibase_20210917.vhd	5.10.23-5.8.lifsea8.x86_64	2021-10-15	• The ContainerOS base image is released.
				• Key component versions:
				  ◦ kubelet: 1.20.4
				  ◦ containerd: 1.4.8
				• Kernel description:
				  ◦ The image is based on the 5.10 kernel version of the Alibaba Cloud Linux 3.2104 64-bit image. The 5.10.23-5.8.lifsea8.x86_64 kernel version is used by the base image.
				  ◦ initramfs is removed to accelerate startup.
				  ◦ Some required kernel modules are compiled as built-in modules.

OS	OS version	Kernel version	Release date	Description

> **Note** The ContainerOS image is supported by all instance families, except Elastic Compute Service (ECS) bare metal instances and heterogeneous computing instances.

August 2021

OS version	Kernel version	Release date	Description
aliyun_2_1903_x64_20G_alibase_20210726.vhd	4.19.91-24.1.al7.x86_64	2021-08-20	For more information, see Alibaba Cloud Linux 2镜像发布记录.

April 2021

OS version	Kernel version	Release date	Description
aliyun_2_1903_x64_20G_alibase_20210325.vhd	4.19.91-23.al7.x86_64	2021-04-19	For more information, see Alibaba Cloud Linux 2镜像发布记录.
centos_7_9_x64_20G_alibase_20210318.vhd	3.10.0-1160.15.2.el7.x86_64	2021-04-19	For more information, see 公共镜像发布记录.

March 2021

OS version	Kernel version	Release date	Description
centos_7_9_x64_20G_alibase_20210128.vhd	3.10.0-1160.11.1.el7.x86_64	2021-03-03	For more information, see 公共镜像发布记录.

January 2021

OS version	Kernel version	Release date	Description
aliyun_2_1903_x64_20G_alibase_20210120.vhd	kernel-4.19.91-22.2.al7	2021-01-27	For more information, see Alibaba Cloud Linux 2镜像发布记录.

December 2020

OS version	Kernel version	Release date	Description
centos_7_8_x64_20G_alibase_20200914.vhd	3.10.0-1127.19.1.el7.x86_64	2020-12-22	For more information, see 公共镜像发布记录.

October 2020

OS version	Kernel version	Release date	Description
aliyun_2_1903_x64_20G_alibase_20200904.vhd	4.19.91-21.al7.x86_64	2020-10-20	For more information, see Alibaba Cloud Linux 2镜像发布记录.

July 2020

OS version	Kernel version	Release date	Description
aliyun_2_1903_x64_20G_alibase_20200529.vhd	4.19.91-19.1.al7.x86_64	2020-07-06	For more information, see Alibaba Cloud Linux 2镜像发布记录.
centos_7_7_x64_20G_alibase_20200426.vhd	3.10.0-1062.18.1.el7.x86_64	2020-07-06	For more information, see 公共镜像发布记录.

4.Runtime release notes

4.1. Release notes for containerd

containerd is an industry-standard container runtime. containerd manages the entire lifecycle of containers on its host. containerd provides a simple and stable runtime for your containers. This topic describes the release notes for containerd.

Context

For more information about the comparison between containerd and other runtimes, see Comparison of Docker, containerd, and Sandboxed-Container.

March 2022

Version	Release date	Description	Impact
1.5.10	2022-03-22	• The following Common Vulnerabilities and Exposures (CVE) vulnerabilities are patched: ◦ CVE-2022-23648 ◦ CVE-2021-43816 ◦ CVE-2021-41190 • runC is updated to 1.0.3. The issue that a node enters the NotReady state when a process ID (PID) is leaked and the runC pipe is blocked is fixed.	No impact on workloads

August 2021

Version	Release date	Description	Impact
1.4.8	2021-08-03	• The following issue is fixed: Sandbox creation times out due to system overloading, which further causes an IP leak. • The CVE-2021-32760 vulnerability is patched.	No impact on workloads

June 2021

Version	Release date	Description	Impact
1.4.6	2021-06-03	The CVE-2021-30465 vulnerability is patched.	No impact on workloads

March 2021

Version	Release date	Description	Impact
1.4.4	2021-03-16	containerd can be selected as the container runtime when you create a Kubernetes cluster. **Note** The containerd runtime is in public preview.	No impact on workloads

Related information

- Release notes for Docker
- Release notes of Sandboxed-Container

4.2. Release notes for Docker

Docker provides a basic runtime for your containers. This topic describes the release notes for the Docker runtime.

Context

For more information about the comparison between Docker and other runtimes, see Comparison of Docker, containerd, and Sandboxed-Container.

May 2020

Version	Release date	Description	Impact
19.03.15	2021-05-06	The built-in containerd version is upgraded to 1.4.6 and the following Common Vulnerabilities and Exposures (CVE) vulnerabilities are fixed: • CVE-2021-21285 • CVE-2021-21284 • CVE-2021-30465	No impact on workloads

February 2020

Version	Release date	Description	Impact
19.03.5	2020-02-20	For more information about this upgrade, see Docker release notes.	No impact on workloads

Related information

- Release notes for containerd
- Release notes of Sandboxed-Container

4.3. Release notes of Sandboxed-Container

This topic lists the latest changes to the Sandboxed-Container runtime.

Context

For more information about Sandboxed-Container, see Sandboxed-Container overview.

April 2021

Version	Release date	Description	Impact
2.2.0	2021-04-02	The secure computing mode (Seccomp) feature is enabled for the containerd runtime. **Note** The Seccomp feature is supported by clusters of Kubernetes V1.20 or later.	No impact on workloads.

March 2021

Version	Release date	Description	Impact
2.1.2	2021-03-01	The issue where exceptions occur in privileged containers in some scenarios is fixed.	No impact on workloads.

January 2021

Version	Release date	Description	Impact
2.1.1	2021-01-07	Privileged containers are supported.	No impact on workloads.

December 2020

Version	Release date	Description	Impact
2.1.0	2020-11-26	New features are released to improve service stability and performance. New features: • A project quota is supported to limit the number of bytes that can be written to the container rootfs directory. • A disk can be mounted to a sandboxed container. • An Apsara File Storage NAS (NAS) file system can be mounted to a sandboxed container. • Custom kernel parameters are supported for sandboxed pods. • Quality of Service (QoS) policies and network traffic marking policies are supported.	No impact on workloads.

August 2020

Version	Release date	Description	Impact
2.0.0	2020-08-28	Sandboxed-Container V2.0 is released to achieve the following benefits: • Sandboxed-Container V2.0 is a container runtime developed by Alibaba Cloud on top of lightweight virtual machines. This version supports more lightweight and efficient deployment, and simplifies the architecture and maintenance of Kubernetes clusters. • Reduces the resource overheads by 90% and accelerates the startup of sandboxed containers by three times. • Increases the deployment density of standalone sandboxed containers by 10 times. • The virtio-fs file system is supported. The performance of this file system is higher than the performance of the 9pfs file system.	During the upgrade, the pods on the nodes that use the Sandboxed-Container runtime are recreated. Pay attention to pod redundancy.

July 2020

Version	Release date	Description	Impact
1.1.1	2020-07-27	The following issues that are related to the stability of Sandboxed-Container are fixed: • The security risk that is related to the container-storaged component is eliminated. • The issue where the `kubectl cp` command is blocked after you run this command is fixed. • The issue where logs cannot be printed to stdout files after containerd is restarted is fixed. • The issue where the system time of sandboxed containers may not be synchronized at regular intervals is fixed.	No impact on workloads.

March 2020

Version	Release date	Description	Impact
1.1.0	2020-03-05	New features of Sandboxed-Container V1.1.0 are released: • Alibaba Cloud disks and NAS file systems can be mounted to sandboxed containers. This provides the same performance as the volumes that are mounted to the host and avoids performance loss when storage devices are mounted over 9pfs. • RootFS block I/O throttling is supported. The stability of Sandboxed-Container V1.1.0 is significantly improved.	No impact on workloads.

September 2019

Version	Release date	Description	Impact

Version	Release date	Description	Impact
1.0.0	2019-09-05	The following features of Sandboxed-Container V1.0.0 are supported: • Strong isolation based on sandboxed and lightweight virtual machines. • Compatibility with runC in terms of application management. • High performance that is equivalent to 90% of the performance provided by applications based on runC. • The same user experience as runC in terms of logging, monitoring, and storage. • The RuntimeClass feature is released. This feature allows you to select container runtimes such as runC and runV. For more information, see RuntimeClass. • Ease of use with minimum technical skill requirements. • Higher stability compared with the open source Kata Containers runtime. For more information about Kata Containers, see Kata Containers.	No impact on workloads.

Related information

- Release notes for containerd
- Release notes for Docker

5.Release notes for components

5.1. Core components

5.1.1. Kube API Server

The Kubernetes API server is the access gateway to a Kubernetes cluster. This topic introduces the kube-apiserver component and its usage notes, and lists the latest changes to the component.

Introduction

The Kubernetes API server validates and configures data for the API objects, which include pods, Services, and ReplicationControllers. The Kubernetes API server serves REST operations and provides a frontend to the shared state of the cluster. All other components interact through this frontend.

Usage notes

The kube-apiserver component is automatically installed. You can use it without extra configurations.

Release notes

The kube-apiserver component is upgraded along with the Kubernetes version. For more information, see Overview of Kubernetes versions supported by ACK.

5.1.2. Kube Controller Manager

The Kubernetes controller manager manages the resources in a Kubernetes cluster. This topic introduces the kube-controller-manager component and its usage notes, and lists the latest changes to the component.

Introduction

The Kubernetes controller manager is a daemon that embeds the core control loops that are shipped with Kubernetes. The Kubernetes controller manager monitors the state of the cluster through the Kubernetes API server and makes sure that the cluster runs in the desired state. The Kubernetes controller manager consists of multiple controllers that manage differnet resources. Kubernetes provides the following built-in controllers: replication controller, endpoints controller, namespace controller, and service accounts controller.

Usage notes

The kube-controller-manager component is automatically installed. You can use it without extra configurations.

Release notes

The kube-controller-manager component is upgraded along with the Kubernetes version. For more information, see Overview of Kubernetes versions supported by ACK.

5.1.3. Cloud Controller Manager

This topic introduces the cloud controller manager (CCM) and provides usage notes and release notes for the component.

Introduction

The CCM allows you to integrate Kubernetes with Alibaba Cloud services, such as Classic Load Balancer (CLB) and Virtual Private Cloud (VPC). CLB is formerly known as Server Load Balancer (SLB). The CCM provides the following features:

- Manage CLB instances

  If you set `Type=LoadBalancer` for a Service, the CCM automatically creates a CLB instance for the Service, and configures listeners and backend server groups. When the endpoint of an Elastic Compute Service (ECS) instance in a vServer group for a Service is changed or the cluster nodes are changed, the CCM automatically updates the vServer groups of the CLB instance.

- Enable cross-node communication

  If Flannel is used as the network plug-in of a Kubernetes cluster, the CCM can enable network connections between containers and nodes. This allows you to implement cross-node communication. The CCM adds the pod CIDR block to the route table of the VPC where the cluster is deployed. This enables cross-node communication. This feature is ready for use after the CCM is installed.

Usage notes

- The CCM automatically creates CLB instances for Services, and configures listeners and backend server groups. For more information, see Considerations for configuring a LoadBalancer type Service.
- The CCM allows you to add annotations in Service configurations to customize load balancing. For more information, see Use annotations to configure load balancing.

Release notes

March 2022

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v2.3.0	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v2.3.0	2022-03-21	• New features: ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-hostname` can be used to specify the hostname of a Service. ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-established-timeout` can be used to specify the connection timeout period for TCP listeners of CLB instances. ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-request-timeout` can be used to specify the request timeout period for HTTP and HTTPS listeners of CLB instances. ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-health-check-method` can be used to specify the health check method for HTTP health checks of CLB instances. • Improvements: ◦ The format of vServer groups is verified when you reuse existing vServer groups. ◦ The logic of vSwitch selection is optimized to resolve the issue that the default vSwitch is not specified. ◦ The synchronization logic of vServer groups is optimized to reduce the number of API calls.	No impact on workloads

November 2021

Version	Image address	Release date	Description	Impact

v2.1.0	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v2.1.0	2021-11-22	• New features:
			  ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-xforwardedfor-proto` can be used to obtain the listener protocol of a CLB instance from the X-Forwarded-Proto header field.
			  ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-idle-timeout` can be used to specify the timeout period of idle connections.
			  ◦ The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-http2-enabled` can be used to enable HTTP2.
			• Improvements:
			The annotation `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-weight` can be set to 0 to stop distributing traffic to specific backend servers.
			• Fixed issues:
			  ◦ The issue that listeners cannot be created for a CLB instance when a large number of backend pods are added to the CLB instance.
			  ◦ The issue that the CLB instance used by a Service is not updated after the targetPort parameter of the Service is updated.
				No impact on workloads

September 2021

Version	Image address	Release date	Description	Impact
v2.0.1	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v2.0.1	2021-09-02	• New features: ◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-vgroup-port* can be used to reuse an existing vServer group that is added to a CLB instance. This annotation takes effect only when the CLB instance is reused. For more information, see Use the CCM to deploy services across clusters. ◦ When a reused CLB instance is shared among multiple Services, the annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-weight* can be used to set the weight of each Service to enable weighted round robin. This annotation takes effect only when the existing vServer group is reused. For more information, see Use the CCM to deploy services across clusters. ◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-connection-drain* can be used to configure connection draining for a CLB instance. Only TCP and UDP are supported. ◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-connection-drain-timeout* can be used to set the timeout period of connection draining for a CLB instance. Only TCP and UDP are supported. ◦ The targetPort field can be set to a String value. ◦ Finalizers can be specified for LoadBalancer Services. • Improvements: ◦ Alpine Linux is updated to V3.13 for base images.	No impact on workloads

Version	Image address	Release date	Description	Impact
			◦ The port used by Prometheus metrics is changed from 10258 to 8080. ◦ The node labels are synchronized by schedule.	

April 2021

Version	Image address	Release date	Description	Impact
v1.9.3.380-gd6d0962-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.380-gd6d0962-aliyun	2021-04-20	• The issue that the default server group cannot be updated is fixed. • Events are generated and alerts are triggered when a CLB instance is not associated with backend servers.	No impact on workloads

March 2021

Version	Image address	Release date	Description	Impact
			New features: • ECS instances other than those in the Container Service for Kubernetes (ACK) cluster can be added to a vServer group. • The label `kubernetes.reused.by.user` is automatically added to a reused CLB instance. Improvements: • The number of concurrent threads for processing Services is increased to improve processing speed. • The processing logic of virtual-kubelet nodes is optimized to ignore Service updates caused	

Version	Image address	Release date	Description	Impact
v1.9.3.378-g42eac35-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.378-g42eac35-aliyun	2021-03-08	by the status changes of virtual-kubelet nodes. • The node label `service.beta.kubernetes.io/exclude-node` is deprecated. To exclude a node from the management of the CCM, use the label `service.alibabacloud.com/exclude-node` instead. • Resource groups are verified when a CLB instance is reused. The resource group ID specified in annotations must be the ID of the resource group to which the CLB instance belongs. Otherwise, the CLB instance cannot be used to expose more than one Service. • The readability of event content is improved. • The version priority setting of annotations is optimized. If two versions of an annotation are added to the Service configurations, the later version prevails over the earlier version. Fixed issues: • The issue that route entries failed to be deleted due to incomplete node configurations. • The logic of node initialization is optimized to fix the issue of taint missing. This prevents pods from being scheduled to a node for which route entries are not created during the initialization process.	No impact on workloads

December 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3.339-g9830b58-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.339-g9830b58-aliyun	2020-12-18	• Hash values are supported in the configurations of LoadBalancer Services. This way, when the CCM is restarted, only the vServer groups of the related CLB instances are updated if the Service configuration is not changed. The configurations of the related CLB instances and listeners are not updated. • CLB API calls are optimized to reduce the chances of throttling.	No impact on workloads

September 2020

Version	Image address	Release date	Description	Impact
v1.9.3.316-g8daf1a9-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.316-g8daf1a9-aliyun*	2020-09-29	• The occasional failure to update the vServer groups of CLB instances is fixed. • The health check port is changed from 10252 to 10258.	No impact on workloads

August 2020

Version	Image address	Release date	Description	Impact
			• New features: ◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-delete-protection* can be used to set deletion protection for CLB instances. By default, deletion protection is enabled for newly created CLB instances.	

Version	Image address	Release date	Description	Impact
v1.9.3.313-g748f81e-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.313-g748f81e-aliyun*	2020-08-10	◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-modification-protection* can be used to set the configuration read-only mode for CLB instances. By default, the configuration read-only mode is enabled for newly created CLB instances. ◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-resource-group-id* can be used to specify the resource group to which a CLB instance belongs. This setting applies only when you create a CLB instance and cannot be modified after the instance is created. ◦ The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-name* can be used to specify the name of a CLB instance. ◦ The API operations of Alibaba Cloud services can be called over internal networks instead of the Internet. To call the CCM operations, Internet access is no longer required in all regions. ◦ Tags are added to a CLB instance that is created for a LoadBalancer Service. The tags are in the `ack.aliyun.com: {your-cluster-id}` format. This feature applies to only newly created clusters. ◦ The cloud provider ID can be specified in the `<cloudProvider>://<optional>/<segments>/<provider id>` format, which is compatible with open source Kubernetes.	No impact on workloads

Version	Image address	Release date	Description	Impact
			◦ When a LoadBalancer Service is created in a cluster that uses Terway, the backed pods are automatically added to the CLB instance that is associated with the Service. The IP addresses of elastic network interfaces (ENIs) that are allocated to the pods are added as the backend servers of the CLB instance. This improves network performance. For LoadBalancer Services, the targetPort field cannot be set to a string value. • Improvements: ◦ Alpine Linux is updated to V3.11.6 for base images. ◦ Listener updates are automatically synchronized to vServer groups. ◦ CLB API operations are optimized. You can call the CLB API to create CLB instances with improved speed.	

June 2020

Version	Image address	Release date	Description	Impact
			• New features: ◦ The CLB instance attached to the cluster API server cannot be reused by LoadBalancer Services. ◦ Prometheus metrics (ccm_node_latencies_duration_milliseconds, ccm_route_latencies_duration_milliseconds, and ccm_slb_latencies_duration_milliseconds) are added to monitor the synchronization latency of the CCM. ◦ Events are collected to monitor the synchronization process between a Service and the related CLB instance. • Improvements:	

Version	Image address	Release date	Description	Impact
v1.9.3.276-g372aa98-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.276-g372aa98-aliyun*	2020-06-11	◦ Weight calculation is optimized for Services in Local mode. To enable the Local mode, set externalTrafficPolicy=Local in Service configurations. This improves load balancing among pods. For more information, see How does the CCM calculate node weights in Local mode?. ◦ API calls of cloud services are optimized to improve efficiency and reduce the chances of throttling. ◦ When you delete a node with the label service.beta.kubernetes.io/exclude-node, the related route entries are no longer deleted. • Fixed issues: ◦ The issue that persistence timeout cannot be set to 0 by adding annotations during Service updates. ◦ The issue that bandwidth cannot be set to 100 by adding annotations during Service updates.	No impact on workloads

March 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3.239-g40d97e1-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.239-g40d97e1-aliyun*	2020-03-05	• New features: For *LoadBalancer* Services, the CCM allows you to specify both ECS nodes and ENIs as the backend servers of the related CLB instances. • Improvements: ◦ The API operations of Alibaba Cloud services can be called over internal networks instead of the Internet. To call the CCM operations, Internet access is no longer required in regions other than China (Beijing), China (Shanghai), and UAE (Dubai). ◦ The API operation that is used to query VPC route entries is changed to DescribeRouteEntryList. This provides higher performance when hundreds of queries are received within a short period of time.	No impact on workloads

December 2019

Version	Image address	Release date	Description	Impact
			• vSwitch IDs are supported. You can set vSwitch IDs in CloudConfig by using the following format: `:vswithid1,:vswitchid2`. • Backoff is supported when throttling is enabled. Backoff allows failed requests to rejoin the reconcile queue every 30 to 180 seconds. • The number of worker threads to be reconciled is adjusted to 2. This allows you to fully utilize the queries per second (QPS) quota on API calls to speed up the reconcile process.	

Version	Image address	Release date	Description	Impact
v1.9.3.220-g24b1885-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.220-g24b1885-aliyun*	2019-12-31	• The issue that the CCM quits unexpectedly due to concurrent Map reads and writes based on the aliyungo SDK is fixed. • When a node is removed from an ACK cluster, the related route entries are automatically deleted from the VPC route table by the CCM. • The issue that port configurations cannot be changed due to port dependencies for HTTP port forwarding is fixed. • If the backend server of a CLB instance is an ECS instance, the serverip field is no longer required when you change the backend server. This prevents errors caused by the changes of default serverip values in API requests when you add backend servers. • The route entries of a node are added to the VPC route table only if the status of the node is known. • NAT IP addresses are no longer added to node metadata by the CCM. This fixes the issue that the API server occasionally fails to connect to kubelet. • When you modify the configurations of a listener, the start listener operation is called only if the listener is in the inactive state. This prevents throttling on API requests.	No impact on workloads

November 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3.193-g6cddde4-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.193-g6cddde4-aliyun*	2019-11-19	• The label *service.beta.kubernetes.io/exclude-node* can be added to a node. After the label is added, the node is no longer managed by the CCM. • Multiple backend pods can be added to a CLB instance at a time. The network type of the pods must be Terway. • The node weight cannot be less than 1 for Services in Local mode (when *externalTrafficPolicy=Local* is set for the Services). • The issue that vServer groups are repeatedly created when concurrent requests are processed is fixed. • The issue that stale data is generated due to caching when you set node weights is fixed.	No impact on workloads

September 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3.164-g2105d2e-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3-164-g2105d2e-aliyun*	2019-09-11	• The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-cert-id* can be used to renew a certificate. • The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-forward-port* can be used to enable port forwarding from an HTTP port to an HTTPS port. • The following annotations can be used to create CLB instances with access control list (ACL) settings: *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-acl-status*, *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-acl-id*, and *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-acl-type*. • The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-remove-unscheduled-backend* can be used to remove unschedulable nodes. • When the Terway network plug-in is used, you can use the annotation *service.beta.kubernetes.io/backend-type: "eni"* to add pods that are assigned ENIs as the backend servers of a CLB instance. This improves network forwarding performance. • Services in Local mode (when *externalTrafficPolicy=Local* is set for the Services) can automatically set node weights based on the number of pods on each node.	No impact on workloads

April 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3.105-gfd4e547-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.105-gfd4e547-aliyun*	2019-04-15	• Multiple route tables can be created for a VPC. Configuration files can be used to set multiple route tables for a cluster. • The issue that updated HTTP configurations do not take effect is fixed.	No impact on workloads

March 2019

Version	Image address	Release date	Description	Impact
v1.9.3.81-gca19cd4-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.81-gca19cd4-aliyun*	2019-03-20	• Existing CLB instances that are not created by ACK can be reused by ACK managed clusters and ACK dedicated clusters. • Custom node names are supported. Node naming is no longer reliant on the nodeName field in Kubernetes. • The compatibility issue between CCM 1.8.4 and Kubernetes 1.11.5 is fixed. We recommend that you update the CCM to the latest version.	No impact on workloads

December 2018

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact

v1.9.3.59-ge3bc999-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.59-ge3bc999-aliyun*	2018-12-26	• A CLB instance can be shared by multiple Kubernetes Services.
			  ◦ If a CLB instance is created along with a Service, you cannot reuse this CLB instance when you create other Services. Otherwise, the CLB instance may be deleted. Only CLB instances that are manually created in the console or by calling the API can be used to expose multiple Services.
			  ◦ Kubernetes services that share the same CLB instance must use different frontend listening ports. Otherwise, port conflicts may occur.
			  ◦ When you reuse a CLB instance, you must use the listener name and vServer group name as identifiers. Do not modify the names of listeners or VServer groups.
			  ◦ You can modify the CLB instance name.
			  ◦ You cannot share CLB instances across clusters.
			• VPC route tables are managed in sequence instead of in parallel. This prevents throttling.
				No impact on workloads

August 2018

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3.10-gfb99107-aliyun	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.10-gfb99107-aliyun*	2018-08-15	• The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-master-zoneid* can be used to specify the primary zone for an automatically created CLB instance. • The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-slave-zoneid* can be used to specify secondary zones for an automatically created CLB instance. **Note** This parameter does not take effect in regions that do not support CLB instances that are deployed across the primary zone and secondary zones. • The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-force-override-listeners* can be used to overwrite the existing listeners when you reuse an existing CLB instance. • The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-bandwidth* can be used to specify the bandwidth when you create a pay-by-bandwidth CLB instance. The bandwidth is shared among listeners of the CLB instance.	No impact on workloads

June 2018

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.9.3	*registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3*	2018-06-25	• The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-backend-label* can be used to add worker nodes with specific labels as the backend servers of a CLB instance. • The annotation *service.beta.kubernetes.io/alibaba-cloud-loadbalancer-spec* can be used to specify the CLB instance type, such as shared-resource or high-performance. • The configuration of `externalTrafic: Local` can be used to set the Local mode for Services. If this mode is enabled, only nodes that host the pods are added as the backend servers of the related CLB instance. • If a node is added to or removed from a cluster, the node is automatically added to or removed from the backend servers of the related CLB instances. • When the labels of a node are changed, the node is automatically added to or removed from the backend servers of the related CLB instances. • Sticky sessions are supported. • Listeners are no longer managed by the system when you create a Service by using an existing CLB instance. You must manually add listeners.	No impact on workloads

5.1.4. edge-controller-manager

edge-controller-manager is the control component of an edge Kubernetes cluster. This topic lists the latest changes to edge-controller-manager.

Introduction

edge-controller-manager is the core control component of an edge Kubernetes cluster. It provides lifecycle management, IP Address Management (IPAM), and network enhancement for edge nodes.

> **Notice** edge-controller-manager is automatically installed and is not displayed in the ACK console. It is displayed in the ACK console only when it is upgradable. You can upgrade it in the ACK console.

Instruction

edge-controller-manager is automatically installed. You can use it without extra configurations.

Release notes

January 2021

Version	Image address	Release date	Description	Impact
v1.16.9-aliyunedge.1	registry.cn-hangzhou.aliyuncs.com/acs/edge-controller-manager:v1.16.9-aliyunedge.1	2021-01-14	The IPAM and network enhancement features are provided.	N/A

5.2. Application management

5.2.1. appcenter

The appcenter component enables centralized management of application deployment and application lifecycle across multiple clusters. You can manage the appcenter component in the Application Center module of the Container Service for Kubernetes (ACK) console. This topic describes the appcenter component and provides usage notes and release notes for the component.

Introduction

Application Center provides a unified portal for your applications. This helps you understand how your applications are deployed in a unified view. Application Center allows you to view your applications in a global manner. You can also view the deployment status and changes of all Kubernetes subresources allocated to each application. In addition, Git and Helm charts are used to deploy applications in ACK clusters based on application versions. This allows you to publish or roll back applications that are deployed in ACK clusters.

Usage notes

For more information about how to use appcenter, see Application Center overview.

Release notes

March 2022

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v2.2.5.1-7fcb8c2-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v2.2.5.1-7fcb8c2-aliyun	2022-03-31	The automatic container image update feature is added.	No impact on workloads
v2.2.5.0-35d527f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v2.2.5.0-35d527f-aliyun	2022-03-15	CVE vulnerabilities are patched.	No impact on workloads

January 2022

Version	Image address	Release date	Description	Impact
v2.1.3.3-7f4fce2-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v2.1.3.3-7f4fce2-aliyun	2022-01-11	The issue that template cannot be associated with RAM users when RAM users using a trigger to enable automatic template application update is fixed.	No impact on workloads

December 2021

Version	Image address	Release date	Description	Impact
v2.1.3.2-308e5f3-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v2.1.3.2-308e5f3-aliyun	2021-12-01	The issue that the component does not run as normal due to network policies is fixed.	No impact on workloads

October 2021

Version	Image address	Release date	Description	Impact
v2.1.3.1-64ae609-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v2.1.3.1-64ae609-aliyun	2021-10-14	The stability of the component is enhanced.	No impact on workloads

June 2020

Version	Image address	Release date	Description	Impact
v1.0.1.1-a97c8f0-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v1.0.1.1-a97c8f0-aliyun	2020-06-22	New features: • An application can be deployed to multiple clusters at a time. • Application triggers are supported.	No impact on workloads

5.2.2. progressive-delivery-tool

This topic introduces progressive-delivery-tool and lists the latest changes to progressive-delivery-tool.

Introduction

You can use progressive-delivery-tool to implement canary releases of your applications in a progressive manner. You can implement a canary release to forward a custom percentage of traffic to the new application version. This way, you can reduce service interruptions caused by errors of the new application version and minimize potential risks. You can also perform rollbacks if required. This allows you to publish the new application version in a progressive manner.

Release notes

November 2020

Version	Image address	Release date	Description	Impact
v1.0.3.7-ea68c41-aliyun	registry..cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v1.0.3.7-ea68c41-aliyun	2020-11-23	Canary releases can be manually implemented.	N/A

August 2020

Version	Image address	Release date	Description	Impact
v1.0.3.6-79c468b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/appcenter-installer:v1.0.3.6-79c468b-aliyun	2020-08-26	Canary releases can be implemented by using Ingresses.	progressive-delivery-tool is released for the first time.

5.2.3. OpenKruise

OpenKruise is a set of standard extensions for Kubernetes. It can be used with native Kubernetes to manage application pods, sidecar containers, and image distribution in an efficient manner. This topic introduces OpenKruise and describes the usage notes and release notes for OpenKruise.

Introduction to OpenKruise

OpenKruise is an open source automation engine that Alibaba Cloud provides for cloud-native applications. It is used as a deployment base to migrate the business of Alibaba Group to the cloud. OpenKruise has joined the Cloud Native Computing Foundation (CNCF) Sandbox project. OpenKruise contains a variety of custom workloads. You can use the workloads to deploy and manage stateless applications, stateful applications, sidecar containers, and daemon applications. OpenKruise also supports advanced strategies such as in-place updates, canary releases, stream updates, and priority configuration.

Usage notes

For more information about the usage notes of OpenKruise, see Use OpenKruise to deploy cloud-native applications.

Release notes

January 2022

Version	Image address	Release date	Description	Impact
v1.0.0	registry-vpc.{{.Region}}.aliyuncs.com/acs/kruise-manager:v1.0.0	2022-01-25	New features: • ConfigMaps and Secrets can be distributed across namespaces. • The launch sequence of containers can be specified. • In-place updates of environment variables can be performed by calling the Downward API.	N/A

October 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.10.0	registry-vpc.{{.Region}}.aliyuncs.com/acs/kruise-manager:v0.10.0	2021-10-15	New features: • Elastic topology management is supported for applications. This feature is implemented by using the WorkloadSpread resource. • Security protection is supported for application pods. This feature is implemented by using the PodUnavailableBudget resource. • Pod topology spread constraints are supported by the CloneSet workload. • Rate-limited scaling is supported by the Advanced StatefulSet workload. Image pre-download is supported by Advanced StatefulSet for in-place updates.	N/A

May 2021

Version	Image address	Release date	Description	Impact
v0.9.0	registry.cn-hangzhou.aliyuncs.com/acs/kruise-manager:v0.9.0	2021-05-20	New features: • Container restart and recreation are supported. • Cascading deletion protection is supported. • Pod deletion based on priorities, image pre-download, and efficient rollback are supported by the CloneSet workload. • Hot upgrades of mesh containers are supported by the SidecarSet workload.	N/A

5.3. Logs and monitoring

5.3.1. alicloud-monitor-controller

The alicloud-monitor-controller component is a system component provided by Container Service for Kubernetes (ACK) for integration with CloudMonitor. This topic introduces alicloud-monitor-controller and provides usage notes and release notes for alicloud-monitor-controller.

Introduction

The alicloud-monitor-controller component is a system component provided by ACK for integration with CloudMonitor. When an application is created, modified, or deleted, alicloud-monitor-controller automatically updates the application metadata to CloudMonitor. This allows you to monitor the containers that are provisioned for the application. In addition, alicloud-monitor-controller provides alert templates. You can enable this feature in the console.

Usage notes

For more information about how to use alicloud-monitor-controller, see Alert management.

Release notes

November 2021

Version	Image address	Release date	Description	Impact
v1.5.13-6990db0e-aliyun	registry.{REGION}.aliyuncs.com/acs/alicloud-monitor-controller:v1.5.13-6990db0e-aliyun	2021-11-04	Default alert rules for basic cluster resources are supported by the ACK alert center.	No impact on workloads

August 2021

Version	Image address	Release date	Description	Impact
v1.5.12-f7aedb4a-aliyun	registry.{REGION}.aliyuncs.com/acs/alicloud-monitor-controller:v1.5.12-f7aedb4a-aliyun	2021-08-23	The alert configurations are optimized.	No impact on workloads
v1.5.10-a75e4ee2-aliyun	registry.{REGION}.aliyuncs.com/acs/alicloud-monitor-controller:v1.5.10-a75e4ee2-aliyun	2021-08-19	The alert configurations are optimized.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
v1.5.7-e1d5de8a-aliyun	registry.{REGION}.aliyuncs.com/acs/alicloud-monitor-controller:v1.5.7-e1d5de8a-aliyun	2021-04-23	Event alerting is supported. Alerts can be configured based on events by using CustomResourceDefinition (CRD) objects.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact
v1.4.0-49ff2362-aliyun	registry.{REGION}.aliyuncs.com/acs/alicloud-monitor-controller:v1.4.0-49ff2362-aliyun	2020-07-24	• Metric collection is supported in a cluster where both Windows node pools and Linux node pools are deployed. • The following scaling periods are supported: 15 seconds, 20 seconds, 30 seconds, and 60 seconds. • The issue that Horizontal Pod Autoscaler (HPA) is triggered during rolling updates of applications is fixed.	No impact on workloads

Troubleshooting

No application group is found in the CloudMonitor console

Perform the following checks:

- Check whether the pod where alicloud-monitor-controller is deployed runs as normal in the kube-system namespace.
- Check whether alicloud-monitor-controller is upgraded to the latest version. We recommend that you upgrade the component to the latest version.
- Check the log file of alicloud-monitor-controller for connection timeout errors and SDK errors.

If no error is found after you perform the preceding checks, use the following ticket template to submit a ticket.

Ticket template

1. Is alicloud-monitor-controller upgraded to the latest version?

   Yes

2. Are SDK errors or connection timeout errors found in the log file of alicloud-monitor-controller? If these errors are found, delete the pod.

   No such errors are found. The issue persists after the pod is recreated.

3. Provide the complete log of alicloud-monitor-controller in the ticket.

   Compress the log file and then upload it.

No data is found in the application group in the CloudMonitor console

Perform the following checks:

Check whether the name of the instance in the application group is the same as that of the pod where alicloud-monitor-controller is deployed.

- If the instance name is different from the pod name, perform the steps provided in the preceding section.
- If the instance name is the same as the pod name, check whether the pod where the metrics-server component is deployed runs as normal. Check whether log data is collected as normal. If the `Succesful write 164190 bytes metrics to monitor server` message is found in the log, this indicates that log data is collected as normal.

If no error is found after you perform the preceding checks, use the following ticket template to submit a ticket.

Ticket template

1. Is the name of the instance in the application group the same as that of the pod where alicloud-monitor-controller is deployed?

   Yes

2. Can the log of metrics-server in the kube-system namespace be printed as normal?

   Yes

3. Provide the cluster ID, application name, and pod name in the ticket.

5.3.2. ack-node-problem-detector

The ack-node-problem-detector component is used in Container Service for Kubernetes (ACK). This component is optimized and enhanced based on the open source Node Problem Detector (NPD) that is provided by the Kubernetes community. The ack-node-problem-detector component is used to monitor cluster nodes and integrate third-party monitoring plug-ins. This component detects node anomalies in an ACK cluster and supports the event center feature. You can use ack-node-problem-detector to integrate custom monitoring plug-ins. This allows you to enhance the monitoring of nodes and detect more node anomalies.

Introduction

- The DaemonSet of ack-node-problem-detector detects node anomalies. For more information

about open source node-problem-detector, see node-problem-detector.

- If you specify a sink parameter when the event center feature is enabled, ack-node-problem-detector-eventer is configured for the ack-node-problem-detector component. The ack-node-problem-detector-eventer component is used to monitor events of the cluster and report the events to the event center. For more information about kube-eventer, see kube-eventer.

Usage notes

For more information about the installation, scenarios, and features of ack-node-problem-detector, see Event monitoring.

Release notes

February 2022

Version	Image address	Release date	Description
v1.2.9	• ack-node-problem-detector: registry.aliyuncs.com/acs/node-problem-detector:v0.8.10-e0ff7d2 • kube-eventer: registry-vpc.__ACK_REGION_ID__.aliyuncs.com/acs/kube-eventer-amd64:v1.2.6-f0efecf-aliyun • kube-event-init: registry.{.Values.controller.regionId}.aliyuncs.com/acs/kube-eventer-init:v1.6-a92aba6-aliyun	2022-02-22	• Kernel inspection is supported. • Security is enhanced.

January 2022

Version	Image address	Release date	Description
v1.2.8	• ack-node-problem-detector: registry.aliyuncs.com/acs/node-problem-detector:v0.8.10-e0ff7d2 • kube-eventer: registry-vpc.__ACK_REGION_ID__.aliyuncs.com/acs/kube-eventer-amd64:v1.2.5-cc7ec54-aliyun • kube-event-init: registry.{.Values.controller.regionId}.aliyuncs.com/acs/kube-eventer-init:v1.6-a92aba6-aliyun	2022-01-20	• Compatibility with different modes of containerd is supported. • The Quality of service (QoS) limits of the resources of ack-node-problem-detector are optimized to improve stability.

November 2021

Version	Image address	Release date	Description
v1.2.7	• ack-node-problem-detector: registry.aliyuncs.com/acs/node-problem-detector:v0.8.10-e0ff7d2 • kube-eventer: registry-vpc.__ACK_REGION_ID__.aliyuncs.com/acs/kube-eventer-amd64:v1.2.5-cc7ec54-aliyun • kube-event-init: registry.{.Values.controller.regionId}.aliyuncs.com/acs/kube-eventer-init:v1.6-a92aba6-aliyun	2021-11-25	• This version is compatible with Alibaba Cloud Linux 3 and CentOS 8. • ARM architecture environments are supported.

April 2021

Version	Image address	Release date	Description
v1.2.5	• ack-node-problem-detector: registry.aliyuncs.com/acs/node-problem-detector:v0.6.3-28-160499f • kube-eventer: registry-vpc.__ACK_REGION_ID__.aliyuncs.com/acs/kube-eventer-amd64:v1.2.4-0f5aaee-aliyun • kube-event-init: registry.{.Values.controller.regionId}.aliyuncs.com/acs/kube-eventer-init:1.5-5e0e7c1-aliyun	2021-04-25	• The following issue is fixed: kube-event-init in the kube-system namespace returns the "414 Request Too Large" error when the event center feature is enabled. • The eventer list-watch mechanism is optimized. This prevents etcd from receiving excessive requests. For more information, see eventer list-watch. • The following issue is fixed: kube-eventer fails to parse the timestamps of some system events. For more information, see fix FailedScheduling event write to sls with wrong timestamp.

July 2020

Version	Image address	Release date	Description

Version	Image address	Release date	Description
v0.6.3-28-160499f	registry.aliyuncs.com/acs/node-problem-detector:v0.6.3-28-160499f	2020-07-27	• The following information can be added to out of memory (OOM) Killer events: the name of the relative pod, the namespace to which the pod belongs, and the user IDs (UIDs) of the killed processes. • The efficiency of the check_fd plug-in is improved. • Node events are optimized to report that the process ID (PID) usage of cluster nodes exceeds the threshold. • Plug-ins that detect network connections are upgraded. • Alert plug-ins are added to send alerts when the inode usage in the system disks of cluster nodes exceeds the threshold.

5.3.3. ack-node-repairer

When a Node Problem Detector (NPD) detects node exceptions, node events or node conditions are generated and reported to the Container Service for Kubernetes (ACK) cluster. ACK Node Repairer automatically listens to the events or conditions on each node and fixes the issues based on the related configurations. This topic describes how to install and configure ACK Node Repairer.

Prerequisites

- One of the following types of ACK cluster is created:
  - Create an ACK managed cluster.
  - Create an ACK dedicated cluster.
- ACK Node Problem Detector is installed. For more information, see Event monitoring.
- Connect to ACK clusters by using kubectl.

Context

ACK Node Repairer is programmed with a predefined list of commonly occurring node exceptions and the actions to fix these exceptions. When a node exception occurs, ACK Node Repairer automatically triggers the corresponding action on the node to fix the exception. After the node exception is fixed, NPD automatically changes the state of the node exception. This creates a closed-loop system for detecting and repairing node exceptions. The operations and maintenance (O&M) engineers can also define the node exceptions that need to be fixed and the actions to fix the exceptions.

NPD is a tool for diagnosing Kubernetes nodes. NPD detects node exceptions and generates node events when the following exceptions are detected: Docker engine hangs, Linux kernel hangs, outbound traffic anomalies, and file descriptor anomalies.

Install ack-node-repairer

Before you use ACK Node Repairer, you must first install ack-node-repairer.

1.
2.
3. Find and click ack-node-repairer. On the **App Catalog - ack-node-repairer** page, click the **Parameters** tab and specify the AccessKey pair of the current account.

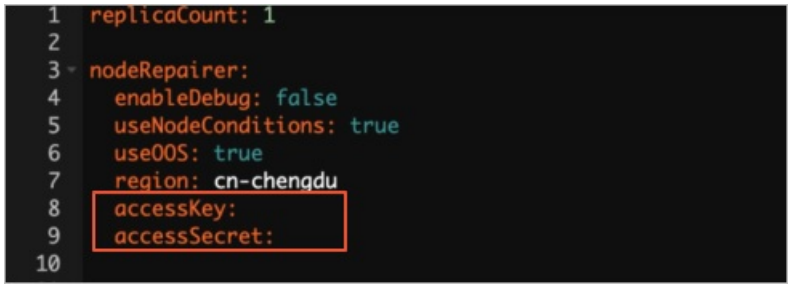

Set the AccessKey pair based on the following description:

- accessKey: nodeRepairer.accessKey
- accessSecret: nodeRepairer.accessSecret

4. On the right side of the page, click **Create**.

Configure ack-node-repairer

After ACK Node Repairer is installed, all auto repairing operations that are supported by the current version are automatically performed. You can configure and enable or disable auto repairing for a type of node exception. The following example shows how to configure ack-node-repairer to automatically fix the issue of Network Time Protocol (NTP) service failure.

1. View the YAML file of default-node-repairer.

   After ACK Node Repairer is installed, a noderepairers.nodes.alibabacloud.com type resource object named default-node-repairer is automatically created in the kube-system namespace. This resource object defines the node exceptions that are monitored by ACK Node Repairer and the related actions to fix the node exceptions. Run the following command to view the content of the YAML file:

   ```
   kubectl -n kube-system get noderepairers.nodes.alibabacloud.com default-node-repairer -o yaml
   ```

2. Modify the configurations of default-node-repairer.

   In the spec.rules field, add the detector parameter to detect the `NTPProblem` condition and add the healers parameter to fix NTP issues. The following code block is an example:

   ```
   spec:
     rules:
     # Specify the detector parameter to detect NTP issues and the nodeOperation parameter to fix NTP issues.
     - detector:
         conditionType: NTPProblem
         type: conditionType
         paused: false
       healers:
       - nodeOperation: restart-ntpd
         type: nodejob
   ```

> **Note** To configure auto repairing for each node exception, you must associate the node condition with the action that is performed to fix the node exception. The rules.detector.conditionType parameter specifies the node condition. If you set rules.detector.paused to *true*, auto repairing is disabled for this type of node condition.

After you perform the preceding steps, when NTP issues occur on a node in the cluster, ACK Node Repairer automatically runs the `systemctl restart chronyd.service` command on the node through OOS to restart NTP on the node.

Records of auto repairing events and results

A noderepairers.nodes.alibabacloud.com type resource object is automatically created in the kube-system namespace to record each auto repairing event. To view the content of this resource object, run the following command. You can also view the auto repairing result by checking the Status field in the output.

```
kubectl -n kube-system get nodejobs.nodes.alibabacloud.com {nodejob_cr_name} -o yaml
```

5.3.4. logtail-ds

The logtail-ds component is an agent provided by Log Service to collect Kubernetes logs. This topic introduces logtail-ds and describes the usage notes and release notes of the component.

Introduction

The logtail-ds component enables high-performance log collection with low resource overhead. By default, the component is installed in the kube-system namespace. The system automatically performs the following operations when you install logtail-ds:

1. Creates a CustomResourceDefinition (CRD) named aliyunlogconfigs that is used to register CRDs with the Kubernetes system.
2. Creates a Deployment named alibaba-log-controller that is used to manage CRDs.
3. Installs Logtail as a DaemonSet that is used to collect log data.

You can use logtail-ds to dynamically filter the containers whose logs you want to collect. logtail-ds can collect multiple types of log data, such as standard output, files, and syslog. The component also supports multiple log parsing methods and configuration methods. For more information about the log collection features, see Use Logtail to collect data.

Usage notes

For more information about the usage notes of logtail-ds, see Collect log data from containers by using Log Service.

Release notes

For more information about the release notes of logtail-ds, see logtail-ds release notes.

5.3.5. metrics-server

The metrics-server component is used to collect resource usage data. You can use metrics-server to collect and monitor resource metrics of your cluster. This topic introduces metrics-server and describes the usage notes and release notes for metrics-server.

Introduction

The metrics-server component is a resource monitoring tool built based on the open source Metrics Server by Container Service for Kubernetes (ACK). The component collects resource usage metrics for all pods in your cluster and enables Horizontal Pod Autoscaler (HPA). You can call the Metrics API to retrieve monitoring metrics.

For more information about the open source Metrics Server, see Kubernetes Metrics Server.

Usage notes

The metrics-server component collects resource usage metrics of all pods in the cluster. For more information, see Monitor basic resources.

Release notes

February 2022

Version	Image address	Release date	Description	Impact
v0.3.8.6-307cf45-aliyun	registry.{REGION}.aliyuncs.com/acs/metrics-server:v0.3.8.6-307cf45-aliyun	February 2022	Compatibility with Kubernetes earlier than 1.22 is added.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact
v0.3.8.5-307cf45-aliyun	registry.{REGION}.aliyuncs.com/acs/metrics-server:v0.3.8.5-307cf45-aliyun	May 2021	• Compatibility with Kubernetes earlier than 1.20 is added. • Container Service Monitoring is enabled in CloudMonitor.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.2.2-b4bf266-aliyun	registry.{REGION}.aliyuncs.com/acs/metrics-server:v0.2.2-b4bf266-aliyun	April 2021	• Collecting metrics from Windows and Linux nodes is supported. • The following scaling periods are supported: 15 seconds, 20 seconds, 30 seconds, and 60 seconds. • The issue that HPA adds excess pods during an application rolling update is fixed.	We recommend that you upgrade the component to the latest version.

Related information

- FAQ about metrics-server

5.3.6. ack-arms-prometheus

The ack-arms-prometheus component is a resource monitoring tool used by Application Real-Time Monitoring Service (ARMS) Prometheus to monitor Container Service for Kubernetes (ACK) clusters. This topic introduces ack-arms-prometheus and describes the usage notes and release notes for ack-arms-prometheus.

Introduction

The ack-arms-prometheus component is a resource monitoring tool provided by ACK for integration with ARMS Prometheus. You must install this component if you want to enable Prometheus monitoring for an ACK cluster. The component uploads data collected from your cluster to ARMS Prometheus. ARMS Prometheus then generates data links based on collected data. You can view dashboards and data links in ARMS Prometheus.

Usage notes

For more information about ack-arms-prometheus, see Enable ARMS Prometheus.

Release notes

February 2022

Version	Release date	Description	Impact

Version	Release date	Description	Impact
1.1.3	February 2022	This image version is compatible with ACK clusters that run Kubernetes 1.22.	No impact on workloads

January 2022

Version	Release date	Description	Impact
1.1.2	January 2022	kube-state-metrics v2.3.0-755434c-aliyun is updated.	No impact on workloads

December 2021

Version	Release date	Description	Impact
1.1.0	December 2021	PodMonitor can be used to collect metrics from pods.	No impact on workloads

September 2021

Version	Release date	Description	Impact
0.1.11	September 2021	• This image version is compatible with ACK clusters that run Kubernetes 1.22. • The version of kube-state-metrics is updated to v2.3.0-755434c-aliyun.	No impact on workloads

August 2021

Version	Release date	Description	Impact
0.1.9	August 2021	GPU metrics can be collected by using GPU Exporter.	No impact on workloads

July 2021

Version	Release date	Description	Impact
0.1.8	July 2021	Compatibility is improved.	No impact on workloads

October 2020

Version	Release date	Description	Impact
0.1.5	October 2020	• Clusters of Kubernetes 1.18 are supported. • Image pulling through internal endpoints is supported.	No impact on workloads We recommend that you update the component to the latest version.

July 2020

Version	Release date	Description	Impact

Version	Release date	Description	Impact
0.1.4	July 2020	• Out-of-the-box Kubernetes monitoring, including pod monitoring, node monitoring, and resource monitoring, is supported. This feature is used to monitor the Kubernetes container runtime where applications are deployed. • Web-based component monitoring, including nine common components such as MySQL, Redis, Kafka, ZooKeeper, and NGINX, is supported. This feature is suitable for scenarios where applications depend on middleware. • Fully-managed Prometheus monitoring system, including the prometheus.yaml collection rule, Grafana dashboard, and alerting system, is supported. This feature meets the needs of migrating data from self-managed Prometheus systems to Alibaba Cloud.	No impact on workloads We recommend that you update the component to the latest version.

April 2020

Version	Release date	Description	Impact

Version	Release date	Description	Impact
0.1.3	April 2020	Control of resources used by the agent is supported.	No impact on workloads We recommend that you update the component to the latest version.

August 2019

Version	Release date	Description	Impact
0.1.2	August 2019	ack-arms-prometheus is released.	No impact on workloads We recommend that you update the component to the latest version.

Related information

- FAQ about ack-arms-prometheus

5.3.7. ags-metrics-collector

This topic describes the latest changes to ags-metrics-collector.

Introduction

ags-metrics-collector is a monitoring component for Alibaba Cloud Genomics Service (AGS) customers. You can use it to monitor the resources that are used by each node in AGS workflows.

Instruction

ags-metrics-collector is automatically installed. You can use it without extra configurations. For more information about how to use ags-metrics-collector, see AGS overview.

Release notes

Version	Image address	Release date	Description	Impact
v21.3.9.0-e5e121c-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/ags-metrics-collector:v21.3.9.0-e5e121c-aliyun	2021-03-09	Alpine Linux is upgraded to V3.11.8 for basic images that are used by ags-metrics-collector.	N/A

5.3.8. logtail-windows

This topic introduces logtail-windows and describes the usage notes and release notes for logtail-windows.

Introduction

Container Service for Kubernetes (ACK) allows you to use the logtail-windows component to collect log data from Windows containers and send the data to Log Service.

Usage notes

logtail-windows must be manually installed. For more information, see Use Logtail to collect application logs from Windows nodes.

- logtail-windows 1.0.x: logtail-windows 1.0.x is compatible with all Windows worker nodes. logtail-windows 1.0.20 and later are incompatible with Windows Server version 1903. To install logtail-windows 1.0.20 or later on a worker node, you must use Windows Server version 1909 or 2004 for the node, and configure the Semi-Annual Channel (SAC). For more information about how to manually install logtail-windows 1.0.x, see Install Logtail V1.0.x on Windows nodes.
- logtail-windows 1.1.x: logtail-windows 1.1.x is dependent on the privileged proxy processes on Windows worker nodes. The image size of logtail-windows 1.1.x is 90% smaller than that of logtail-windows 1.0.x. The startup of the logtail-windows 1.1.x image is faster than that of the logtail-windows 1.0.x image. logtail-windows 1.1.x can be installed only on Windows worker nodes that are added from node pools and are added with the `windows.alibabacloud.com/deployment-topology=2.0` label when these nodes are added to the cluster. If the `windows.alibabacloud.com/deployment-topology=2.0` label is added to a Windows worker node after the node is added to the cluster, logtail-windows 1.1.x cannot be installed on the node. For more information about how to manually install logtail-windows 1.1.x, see Install Logtail V1.1.x on Windows nodes.

Release notes

February 2022

Version	Image address	Release date	Description	Impact
v1.1.30-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.1.30-aliyun.1	2022-02-24	• ilogtail is updated to support the addon token feature. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads

August 2021

Version	Image address	Release date	Description	Impact
v1.0.20	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.0.20	2021-08-21	• The base images are updated. The original version number is overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads
v1.1.18-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.1.18-aliyun.1	2021-08-20	• The base images are updated. The original version number is overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads

July 2021

Version	Image address	Release date	Description	Impact
v1.1.18-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.1.18-aliyun.1	2021-07-05	• The base images are updated. The original version number is overwritten. • The size of the logtail-windows image is reduced from 2.7 GB to 220 MB. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1999), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1052).	No impact on workloads

June 2021

Version	Image address	Release date	Description	Impact
v1.1.18-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.1.18-aliyun.1	2021-06-17	• *Rancher Wins* is used as the proxy of privileged operations. • The size of the logtail-windows image is reduced from 2.7 GB to 220 MB. • The following Windows systems are supported: Windows Server version 1809 (10.0.17763.1935), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.985).	The update is applicable only to Windows worker nodes that are installed with *Rancher Wins*.

April 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.20	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.0.20	2021-04-09	Windows Server version 1809 (10.0.17763.1817), Windows Server version 1909 (10.0.18363.1440), and Windows Server version 2004 (10.0.19041.867) are supported. Windows Server version 1903 is no longer supported.	• The update is not applicable to worker nodes that run Windows Server version 1903. • No impact on worker nodes that run Windows Server version 1809, Windows Server version 1909, or Windows Server version 2004. • Windows Server version 1903 is no longer supported in later logtail-windows versions.

January 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.19	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.0.19	2021-01-27	• Logstores can be created over the Internet. • Windows Server version 1809 (10.0.17763.1697), Windows Server version 1903 (10.0.18362.1256), Windows Server version 1909 (10.0.18363.1316), and Windows Server version 2004 (10.0.19041.746) are supported.	No impact on workloads
v1.0.18	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.0.18	2021-01-25	• The availability issue in the China South 1 Finance region is fixed. • Logtail is updated to 1.0.0.18. • Windows Server version 1809 (10.0.17763.1697), Windows Server version 1903 (10.0.18362.1256), Windows Server version 1909 (10.0.18363.1316), and Windows Server version 2004 (10.0.19041.746) are supported.	No impact on workloads
v1.0.10	registry-vpc.${region}.aliyuncs.com/acs/logtail-windows:v1.0.10	2021-01-10	• Logtail is updated to 1.0.0.10. • Windows Server version 1809 (10.0.17763.1637), Windows Server version 1903 (10.0.18362.1256), Windows Server version 1909 (10.0.18363.1256), and Windows Server version 2004 (10.0.19041.685) are supported.	No impact on workloads

5.3.9. ack-cost-exporter

This topic lists the latest changes to ack-cost-exporter.

Introduction

The cost analysis feature is implemented by using the following methods:

- ack-cost-exporter submits the billing data of a cluster for cost analysis. The billing data includes the costs, real-time prices, billing methods, discounts, coupons, and vouchers of cloud resources.
- ack-arms-prometheus submits the resource usage data of a cluster.

- ack-arms-prometheus performs cost analysis based on the billing data and the resource usage data.

ack-cost-exporter consists of alibaba-cloud-price-exporter and alibaba-cloud-billing-exporter:

- alibaba-cloud-price-exporter
  - Periodically obtains the instance types and instance prices in a Kubernetes cluster, and converts these details into Prometheus metrics.
  - Supports public cloud, hybrid cloud, and multi-cloud scenarios.
  - Supports multiple billing methods, such as subscription, pay-as-you-go, and preemptible instances.
- alibaba-cloud-billing-exporter

  Regularly obtains the bills of a cluster and converts them into Prometheus metrics.

Instruction

For more information about how to use ack-cost-exporter, see Enable cost analysis.

Release notes

May 2021

Version	Image address	Release date	Description	Impact
v0.1.0	• registry.cn-zhangjiakou.aliyuncs.com/acs/alibaba-cloud-price-exporter:v0.1.0-64dae5a-aliyun • registry.cn-zhangjiakou.aliyuncs.com/acs/alibaba-cloud-billing-exporter:v0.1.0-4194980-aliyun	2021-05-11	• Cost data of a pod can be submitted. • Cost forecasts can be performed based on different sales strategies of computing resources.	No impact on workloads.

5.3.10. ack-arms-cmonitor

You can install the ack-arms-cmonitor component in your Container Service for Kubernetes (ACK) cluster to use the Kubernetes Monitoring feature provided by Application Real-Time Monitoring Service (ARMS). This topic introduces ack-arms-cmonitor and describes the usage notes and release notes for ack-arms-cmonitor.

Introduction

Kubernetes Monitoring is an all-in-one observability service that is developed for Kubernetes clusters. Kubernetes Monitoring provides IT developers and O&M engineers with a comprehensive observability solution that monitors Kubernetes clusters in various aspects, such as metrics, traces, logs, and events.

Usage notes

For more information about how to use ack-arms-cmonitor, see What is Kubernetes Monitoring?.

Release notes

December 2021

Version	Release date	Description	Impact
v1.0.0	2021-12-02	Custom AccessKey IDs and AccessKey secrets are supported.	This update has no negative impact on workloads. We recommend that you update the component to the latest version.

November 2021

Version	Release date	Description	Impact
v0.4.0	2021-11-30	• The permissions to read Secrets globally are revoked. Only ARMS tokens can be read by ack-arms-cmonitor. • The hostIPC configuration is removed.	This update has no negative impact on workloads. We recommend that you update the component to the latest version.
v0.3.0	2021-11-10	• The ARMS Addon Token feature is supported. • The ClusterId and region_id environment variables are optimized.	This update has no negative impact on workloads. We recommend that you update the component to the latest version.

September 2021

Version	Release date	Description	Impact

Version	Release date	Description	Impact
v0.2.0	2021-09-16	• Domain Name System (DNS), Cassandra, and the S3 protocol are supported. • The user_id and region_id environment variables are added.	This update has no negative impact on workloads. We recommend that you update the component to the latest version.

June 2021

Version	Release date	Description	Impact
v0.1.0	2021-06-15	The first version is released.	This update has no negative impact on workloads. We recommend that you update the component to the latest version.

5.4. Storage

5.4.1. csi-plugin

You can use csi-plugin to mount and unmount volumes. Container Service for Kubernetes (ACK) allows you to mount the following types of volumes by using csi-plugin: disks, Apsara File Storage NAS (NAS) file systems, and Object Storage Service (OSS) buckets. This topic describes the release notes for csi-plugin.

Introduction

The csi-plugin component provided by Alibaba Cloud follows the standard specification of Container Storage Interface (CSI) and enables you to mount storage media such as Elastic Block Storage (EBS) devices, NAS file systems, and OSS buckets. By default, the latest version of csi-plugin is installed when you create ACK clusters of Kubernetes 1.16 or later. You can directly use csi-plugin to mount storage services provided by Alibaba Cloud after your cluster is created. The csi-plugin component supports volume lifecycle management. You can use csi-plugin to create, mount, unmount, delete, and expand volumes.

Usage notes

For more information, see CSI overview.

Release notes

November 2021

Version	Image address	Release date	Description	Impact
v1.20.7-aafce42-aliyun	registry.cn-{{regionID}}.aliyuncs.com/acs/csi-plugin:v1.20.7-aafce42-aliyun	2021-11-26	• The issue that ossfs cannot synchronize data in real time is fixed. • The issue that NAS volumes fail to be mounted is fixed.	No impact on workloads

October 2021

Version	Image address	Release date	Description	Impact
v1.20.6-2be29b1-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.20.6-2be29b1-aliyun	2021-10-29	Images that contain the Container OS are supported.	No impact on workloads

September 2021

Version	Image address	Release date	Description	Impact
v1.20.5-ff6490f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.20.5-ff6490f-aliyun	2021-09-24	• The sharepath feature of NAS is supported. • When the Quota feature is enabled in NAS, automatic expansion and storage monitoring are supported. • ARM64 and AMD64 architectures are supported.	No impact on workloads

August 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.18.8.51-c504ef45-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.51-c504ef45-aliyun	2021-08-19	• The time parameter for the recycle bin is added to Container Network File System (CNFS). • The `apiVersion` of CNFS is changed from `v1alpha1` to `v1beta1`. • The issue that ossfs cannot synchronize data in real time is fixed. • By default, the Detachdisk option is disabled.	No impact on workloads

July 2021

Version	Image address	Release date	Description	Impact
v1.18.8.47-06405694-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.47-06405694-aliyun	2021-07-06	• The issue that NAS file systems cannot be expanded by using CNFS is fixed. • OSS buckets can be mounted to nodes that are deployed by using the Alibaba Cloud Linux 3 image.	No impact on workloads

June 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.18.8.47-30ba5d25-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.47-30ba5d25-aliyun	2021-06-25	• Persistent volumes (PVs) can be created by using CNFS. • The `snapshot.storage.k8s.io/v1` and `snapshot.storage.k8s.io/v1beta1` API versions are supported by VolumeSnapshots, VolumeSnapshotClasses, and VolumeSnapshotContents. Upgrade the API version to `snapshot.storage.k8s.io/v1` at the earliest opportunity. • csi-snapshotter and snapshot-controller are updated to V4.0.0. • The occasional issue that the token expires when you perform the List Snapshot operation is fixed.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact
v1.18.8.47-906bd535-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.47-906bd535-aliyun	2021-05-20	• Disk partitions can be mounted. • Disk partitions can be expanded.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
v1.18.8.46-afb19e46-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.46-afb19e46-aliyun	2021-04-08	• NAS monitoring is supported. • Disks are supported in the China (Qingdao) region.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
v1.18.8.45-1c5d2cd1-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.45-1c5d2cd1-aliyun	2021-01-13	• By default, volume monitoring is enabled. • The VolumeSnapshot List feature is supported. • Quota groups are supported by NAS volumes. • Custom disk types are supported.	No impact on workloads

November 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.16.9.43-f36bb540-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.16.9.43-f36bb540-aliyun	2020-11-02	• The deployment template is updated to merge drivers into one container. • The issue that subdirectories fail to be created in Extreme NAS file systems is fixed. • Kubernetes 1.18 is supported. • Labels can be added to NAS volumes when you create NAS volumes.	No impact on workloads

August 2020

Version	Image address	Release date	Description	Impact
v1.14.8.41-9efe2ede-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.8.41-9efe2ede-aliyun	2020-08-05	• The issue that snapshots cannot be created from disks is fixed. • The issue that dynamic provisioning of NAS volumes fails due to residual data is fixed. • The check logic of BDF nodes when csi-plugin is started is fixed. • The use of UUID to obtain device paths is no longer supported.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact
v1.14.8.40-146fd1d8-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.8.40-146fd1d8-aliyun	2020-07-13	• EBS snapshots are supported. You can use EBS snapshots to restore data to a beta version. • Extreme NAS volumes can be created and deleted. • The Config SysConfig parameter of EBS volumes is supported when you configure PVs. • The issue that block volumes are loaded twice in BDF mode is fixed. • EBS and NAS volumes are allowed to access APIs by using internal domain names. • The Cloud Paralleled File System (CPFS) driver is updated and the dependency on the kernel is removed.	No impact on workloads

April 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.14.8.38-fe611ad1-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.8.38-fe611ad1-aliyun	2020-04-20	• EBS volumes can be unmounted before you delete the volumes. • The disk creation policy is updated. Standard SSDs are created in preference to ultra disks. Ultra disks are created only when no standard SSD is available. • UUID is supported as a high-priority search option to search for devices that use EBS volumes. • The authentication management in managed Kubernetes clusters is updated. • Security Token Service (STS) is supported to connect to OSS buckets. • DuplicateMount Point errors in EBS are fixed. • The BDF protocol is supported to bind EBS volumes after the volumes are connected.	No impact on workloads

March 2020

Version	Image address	Release date	Description	Impact
v1.14.8.36-93f2b131-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.8.36-93f2b131-aliyun	2020-03-14	• OSSFS is updated to support Key Management Service (KMS) and Bring Your Own Key (BYOK). • Kubernetes 1.16 is supported. • EBS volumes can be connected by using the controller server. • EBS metric values can be obtained.	No impact on workloads

February 2020

Version	Image address	Release date	Description	Impact
v1.14.8.37-aae7f495-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.8.37-aae7f495-aliyun	2020-02-18	• Kubernetes clusters that use CSI and have no Internet access are supported. • The issues related to mount path checks in EBS are fixed.	No impact on workloads

December 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.14.8.32-c77e277b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.8.32-c77e277b-aliyun	2019-12-20	• The EBS PV name can be used as the disk ID. This feature is also supported by FlexVolume. • Mount options can be configured for EBS volumes in MKFS Stage. • Mount options can be configured to have a higher priority than the volume attributes of NAS volumes. • Mount options of OSS volumes can be validated in OSS connectors. • Subpaths of OSS buckets can be mounted as volumes. • Volume topology can be used to dynamically configure Logical Volume Manager (LVM).	No impact on workloads

Version	Image address	Release date	Description	Impact
v1.14.5.60-5318afe-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.14.5.60-5318afe-aliyun	2019-12-20	• The mounted subpath and file system type can be configured for a NAS volume. • Labels can be added to disk volumes when you configure disk volumes. • Multiple regions can be configured for disk volumes that are used in multi-region clusters in the StorageClass definition. • KMS-based authentication can be configured when you configure disk volumes. • Disk volumes can be expanded. This feature is in internal preview. • The topology awareness feature is supported when you configure disk volumes.	No impact on workloads

5.4.2. csi-provisioner

csi-provisioner allows you to automatically create volumes. This topic describes the introduction, usage notes, and release notes for csi-provisioner.

Introduction

Alibaba Cloud provides the csi-provisioner component to help you automate volume creation. The component allows you to create volumes from disks and Apsara File Storage NAS (NAS) file systems. The Kubernetes version of the cluster must be 1.14 or later.

Usage notes

For more information, see CSI overview.

Release notes

November 2021

Version	Image address	Release date	Description	Impact
v1.20.7-aafce42-aliyun	registry.cn-{{regionID}}.aliyuncs.com/acs/csi-plugin:v1.20.7-aafce42-aliyun	2021-11-26	• The issue that ossfs cannot synchronize data in real time is fixed. • The issue that NAS volumes fail to be mounted is fixed.	No impact on workloads

October 2021

Version	Image address	Release date	Description	Impact
v1.20.6-2be29b1-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.20.6-2be29b1-aliyun	2021-10-29	Images based on ContainerOS are supported.	No impact on workloads

September 2021

Version	Image address	Release date	Description	Impact
v1.20.5-ff6490f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.20.5-ff6490f-aliyun	2021-09-24	• The sharepath feature of NAS is supported. • The issue that the component is frequently restarted due to the unstable etcd is fixed. • ARM64 and AMD64 architectures are supported.	No impact on workloads

August 2021

Version	Image address	Release date	Description	Impact
v1.18.8.51-c504ef45-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.51-c504ef45-aliyun	2021-08-19	• The time parameter for the recycle bin is added to Container Network File System (CNFS). • The `apiVersion` of CNFS is changed from `v1alpha1` to `v1beta1`. • The issue that ossfs cannot synchronize data in real time is fixed. • By default, the forcibly detach disk option is disabled.	No impact on workloads

July 2021

Version	Image address	Release date	Description	Impact
v1.18.8.48-cd524404-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.48-cd524404-aliyun	2021-07-06	• The issue that NAS file systems cannot be expanded by using CNFS is fixed. • Object Storage Service (OSS) buckets can be mounted to nodes that are deployed by using the Alibaba Cloud Linux 3 image.	No impact on workloads

June 2021

Version	Image address	Release date	Description	Impact
v1.18.8.47-30ba5d25-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.47-30ba5d25-aliyun	2021-06-25	• The `volumeCapacity` field is deleted from NAS volume configurations. The `allowVolumeExpansion` field is used to specify whether to enable the quota feature. • The `selflink` field is deleted from NAS volume configurations.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact
v1.18.8.47-906bd535-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.18.8.47-906bd535-aliyun	2021-05-20	• Disk partitions can be mounted. • Disk partitions can be expanded.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
v1.6.0-e360c7e43-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.6.0-e360c7e43-aliyun	2021-04-08	• Kubernetes 1.20 is supported. The `metadata.selflink` field is deleted. • The tag of the cluster ID is automatically added to disks. • NAS volumes can be expanded within the quota limit.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
v1.6.0-b6f763a43-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.6.0-b6f763a43-aliyun	2021-01-13	• Database File System (DBFS) volumes are supported. • By default, volume monitoring is enabled. • Local volumes of the QuotaPath type are supported. • The VolumeSnapshot List feature is supported. • Quota groups are supported by NAS volumes. • Custom disk types are supported.	No impact on workloads

November 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.6.0-b6f763a43-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.6.0-b6f763a43-aliyun	2020-11-02	• The deployment template is updated to merge drivers into one container. • The issue that subdirectories fail to be created in Extreme NAS file systems is fixed. • Kubernetes 1.18 is supported. • Labels can be added to NAS volumes when you create NAS volumes.	No impact on workloads

August 2020

Version	Image address	Release date	Description	Impact
v1.4.0-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.4.0-aliyun	2020-08-05	• The issue that snapshots cannot be created from disks is fixed. • The issue that dynamic provisioning of NAS volumes fails due to residual data is fixed. • The check logic of BDF nodes when csi-provisioner is started is fixed. • The use of UUID to obtain device paths is no longer supported.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact
v1.4.0-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.4.0-aliyun	2020-07-13	• Elastic Block Storage (EBS) snapshots are supported. You can use EBS snapshots to restore data to a beta version. • Extreme NAS volumes can be created and deleted. • The Config SysConfig parameter of EBS volumes is supported when you configure PVs. • The issue that block volumes are loaded twice in BDF mode is fixed. • EBS and NAS volumes are allowed to access APIs by using internal domain names. • The Cloud Paralleled File System (CPFS) driver is upgraded and the dependency on the kernel is removed.	No impact on workloads

April 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.4.0-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.4.0-aliyun	2020-04-20	• EBS volumes can be unmounted before you delete the volumes. • The disk creation policy is updated. Standard SSDs are created in preference to ultra disks. Ultra disks are created only when no standard SSD is available. • UUID is supported as a high-priority search option to search for devices that use EBS volumes. • The authentication management in managed Kubernetes clusters is updated. • Security Token Service (STS) is supported to connect to OSS buckets. • DuplicateMount Point errors in EBS are fixed. • The BDF protocol is supported to bind EBS volumes after the volumes are connected.	No impact on workloads

February 2020

Version	Image address	Release date	Description	Impact
v1.4.0-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.4.0-aliyun	2020-02-18	• Kubernetes clusters that use CSI and have no Internet access are supported. • The issues related to mount path checks in EBS are fixed.	No impact on workloads

December 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.2.2-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-provisioner:v1.2.2-aliyun	2019-12-20	• The EBS PV name can be used as the disk ID. This feature is also supported by FlexVolume. • Mount options can be configured for EBS volumes in MKFS Stage. • Mount options can be configured to have a higher priority than the volume attributes of NAS volumes. • Mount options of OSS volumes can be validated in OSS connectors. • Subpaths of OSS buckets can be mounted as volumes. • Volume topology can be used to dynamically configure Logical Volume Manager (LVM).	No impact on workloads

5.4.3. storage-operator

The storage-operator component is used to manage the lifecycle of storage components. This topic describes the features, usage notes, and release notes for storage-operator.

Introduction

Storage Operator is used to manage the lifecycle of storage components. Storage Operator runs as a Deployment, which deploys and upgrades storage components based on the default configurations inherited from the image and the custom configurations provided by ConfigMaps. This helps reduce the complexity of container development and maintenance.

- Default configurations: Storage Operator provides the default configurations of storage components. The default configurations vary based on the version of Storage Operator.
- Custom configurations: ConfigMaps can be used to define custom configurations of storage components, such as version information, and whether to install the component.

Storage Operator preferably uses custom configurations. The default configurations are used only when the custom configurations are not specified.

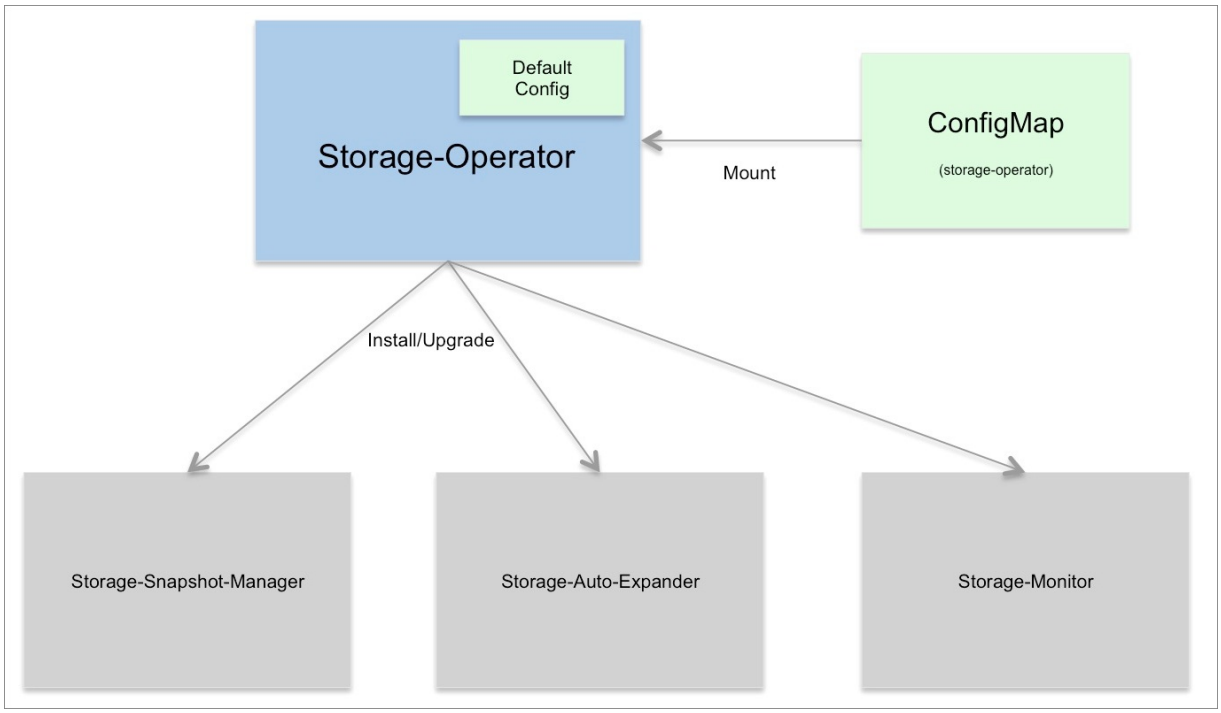

> **Note**
>
> - Each image of Storage Operator contains the default configuration file.
> - When Storage Operator runs as a Deployment, the Deployment reads configurations from a ConfigMap file that is mounted on the Deployment and contains configurations of storage components.
> - Storage Operator determines whether to deploy and upgrade a storage component by combining the default and custom configurations.

Usage notes

For more information about how to deploy storage-operator from App Catalog, see Use Storage Operator to deploy and upgrade storage components.

Release notes

April 2022

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.22.0.75-5bc07f7-aliyun	registry.cn-{{region}}.aliyuncs.com/acs/storage-operator:v1.22.0.75-5bc07f7-aliyun	2022-04-14	Advanced features for disk snapshots are supported.	N/A

March 2022

Version	Image address	Release date	Description	Impact
v1.18.8.68-9078250-aliyun	registry.cn-{{regionID}}.aliyuncs.com/acs/storage-operator:v1.18.8.68-9078250-aliyun	2022-03-16	Vulnerabilities in the CentOS 7 base image are patched.	N/A

December 2021

Version	Image address	Release date	Description	Impact
v1.18.8.67-c1aef60-aliyun	registry.cn-{{regionID}}.aliyuncs.com/acs/storage-operator:v1.18.8.67-c1aef60-aliyun	2021-12-22	The issue that automatic expansion of Apsara File Storage NAS (NAS) file systems fails because the quota increase takes effect slowly is fixed.	N/A

September 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.18.8.60-a5ba617-aliyun	registry-vpc.${region}.aliyuncs.com/acs/storage-operator:v1.18.8.60-a5ba617-aliyun	2021-09-24	• When NAS file systems are managed by Container Network File System (CNFS), the Quota feature is enabled by default and automatic expansion is supported. • The issue that pods are scheduled to unavailable nodes is fixed. • Pods can be scheduled to Linux nodes and nodes other than the virtual-kubelet node. • The issue that the status of CNFS cannot be changed is fixed.	N/A

August 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.18.8.55-e398ce5-aliyun	registry-vpc.${region}.aliyuncs.com/acs/storage-operator:v1.18.8.55-e398ce5-aliyun	2021-08-16	• NAS file systems of the Capacity type can be created by using CNFS. • The archiveOnDelete parameter in the StorageClass that is automatically created by CNFS can be set to *false*. This way, the subdirectory mounted by a persistent volume (PV) is automatically deleted after you delete the PV. • The issue that storage-monitor occupies an excessive amount of CPU resources is fixed.	N/A

June 2021

Version	Image address	Release date	Description	Impact
v1.18.8.37-c63030b-aliyun	registry-vpc.${region}.aliyuncs.com/acs/storage-operator:v1.18.8.37-c63030b-aliyun	2021-06-25	• Automatic expansion is supported. • CNFS is supported.	N/A

March 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.18.8.28-18cca7b-aliyun	registry-vpc.${region}.aliyuncs.com/acs/storage-operator:v1.18.8.28-18cca7b-aliyun	2021-03-25	New features: • Batch snapshots are supported. • Scheduled snapshots are supported. • Cluster monitoring is supported.	N/A

5.4.4. alicloud-disk-controller

Disks can be automatically created by using alicloud-disk-controller. This topic describes the release notes for the alicloud-disk-controller component.

Introduction

To create dynamically provisioned persistent volumes (PVs) in a Kubernetes cluster that uses FlexVolume, you must install alicloud-disk-controller. This is because FlexVolume cannot be used to create dynamically provisioned PVs.

Usage notes

For more information about how to install alicloud-disk-controller, see Install and upgrade FlexVolume.

Release notes

April 2021

Version	Image address	Release date	Description	Impact
v1.16.9.55	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.16.9.55-c43698d4-aliyun	2021-04-05	Kubernetes 1.20 is supported. The selflink field is deleted in metadata.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.14.8.51	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.14.8.51-842f0a81-aliyun	2020-07-09	• Region-specific domain names are supported. • Alpine Linux is upgraded to V3.11.6 for base images.	No impact on workloads

April 2020

Version	Image address	Release date	Description	Impact
v1.14.8.44	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.14.8.44-c23b62c5-aliyun	2020-04-14	• Permissions on managed Kubernetes clusters are reduced. • The disk creation rules in a StorageClass of the available type are modified. The system attempts to create a standard SSD first. If the attempt fails, the system attempts to create an enhanced SSD.	No impact on workloads

March 2020

Version	Image address	Release date	Description	Impact
v1.14.8.37	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.14.8.37-bd3fd891-aliyun	2020-03-06	Region-specific domain names are supported.	No impact on workloads

August 2019

Version	Image address	Release date	Description	Impact
v1.12.6.21	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.12.6.21-54d91d6-aliyun	2019-08-06	• The template is updated and StorageClasses that specify enhanced SSDs are supported. • Leader Selector is disabled. • The following issue is fixed: When a persistent volume claim (PVC) is in the Lost state, duplicate persistent volumes (PVs) are created. • Bring Your Own Key (BYOK) encryption for disks is supported.	No impact on workloads

February 2019

Version	Image address	Release date	Description	Impact
v1.11.2.5	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.11.2.5-1bea041-aliyun	2019-02-18	The topology feature is added to StorageClasses and the WaitForFirstConsumer mode is supported.	No impact on workloads

September 2018

Version	Image address	Release date	Description	Impact
v1.10.4	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.10.4-1847e0f-aliyun	2018-09-13	multi-zone round robin is supported in StorageClasses.	No impact on workloads

July 2018

Version	Image address	Release date	Description	Impact
v1.10.4	registry.cn-hangzhou.aliyuncs.com/acs/alicloud-disk-controller:v1.10.4-821d9c2-aliyun	2018-07-24	Disk encryption is supported.	No impact on workloads

5.4.5. FlexVolume

Volumes can be mounted and unmounted by using FlexVolume. By default, Container Service for Kubernetes (ACK) allows you to mount the following types of volume: disk volumes, Apsara File Storage NAS (NAS) volumes, and Object Storage Service (OSS) volumes. This topic lists the latest changes to FlexVolume.

Introduction

FlexVolume is a traditional mechanism developed by the Kubernetes community to extend the storage of Kubernetes clusters. ACK supports FlexVolume. FlexVolume consists of the following parts:

- FlexVolume: Mounts and unmounts volumes. By default, ACK allows you to mount the following types of volume: disk volumes, NAS volumes, and OSS volumes.
- Disk-Controller: automatically creates disk volumes.
- Nas-Controller: automatically creates NAS volumes.

For more information, see FlexVolume overview.

Usage notes

For more information about how to use FlexVolume, see CSI overview.

Release notes

May 2021

Version	Image address	Release date	Description	Impact
v1.18.8.166	registry.cn-hangzhou.aliyuncs.com/acs/flexvolume:v1.18.8.166-b39c52a-aliyun	2021-05-27	• By default, disk monitoring is enabled. • Server Message Block (SMB) file systems can be mounted to Windows containers.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact
v1.14.8.109	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.14.8.109-649dc5a-aliyun	2020-07-24	Disks that are created from snapshots can be used as persistent volumes (PVs).	No impact on workloads
v1.14.8.104	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.14.8.104-24bde93-aliyun	2020-07-09	• Alpine Linux is upgraded to V3.11.6 for base images. • The internal endpoints of the China (Hong Kong) and China South 1 Finance regions are supported. • The logic for creating subdirectories in Extreme NAS file systems is optimized.	No impact on workloads

April 2020

Version	Image address	Release date	Description	Impact
v1.14.8.88	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.14.8.88-00a6585-aliyun	2020-04-14	• The issue that a disk cannot be expanded when the target size is the same as the current size is fixed. • The issue that the `/dev/vd* not exist` error is returned when you mount a disk is fixed. • Permissions on managed Kubernetes clusters are limited.	No impact on workloads

March 2020

Version	Image address	Release date	Description	Impact
v1.14.8.66	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.14.8.66-1d0843f-aliyun	2020-03-06	• The deployment method of the Cloud Paralleled File System (CPFS) driver is updated. • Subdirectories of OSS buckets can be mounted to containers. • Private API endpoints are supported.	No impact on workloads

October 2019

Version	Image address	Release date	Description	Impact
v1.14.6.15	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.14.6.15-8d3b7e7-aliyun	2019-10-21	• Disks can be automatically expanded. • Monitoring of disk volumes is supported. • The CPFS driver is updated.	No impact on workloads

August 2019

Version	Image address	Release date	Description	Impact
v1.12.6.52	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.12.6.52-f6604e5-aliyun	2019-08-20	The NAS Mount Config configuration is supported. By default, the optimal kernel settings are used.	No impact on workloads

April 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.12.6.16	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.12.6.16-1f4c6cb-aliyun	2019-04-16	• The Sysctl configuration is added for NAS file systems. • Waitforattach operations can be performed on disks.	No impact on workloads
v1.12.6.11	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.12.6.11-ab46951-aliyun	2019-04-09	• The CPFS driver is upgraded. • Disks are locked when they are detached. This way, disks are not concurrently mounted and drive letters are not removed. • The limits on the `kubectl -f` command are updated when NAS file systems are unmounted. • By default, NFSv3 is supported by NAS file systems. • AccessKey pairs can be specified by using Secrets when you mount OSS buckets.	No impact on workloads

December 2018

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.11.2.32	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.11.2.32-af2d48c-aliyun	2018-12-10	• The issue that the OSSFS driver is restarted along with kubelet is fixed. • The waiting period for mounting disks is increased.	No impact on workloads

October 2018

Version	Image address	Release date	Description	Impact
v1.11.2.5	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.11.2.5-85c062f-aliyun	2018-10-18	The issue that disks are concurrently mounted is fixed.	No impact on workloads
v1.11.2	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.11.2-9139592-aliyun	2018-10-17	• The issue that NAS file systems are repeatedly mounted in high concurrency mode is fixed. • The issue that the process of mounting disks is time-consuming is fixed. The issue that the Device parameter is empty is fixed.	No impact on workloads

September 2018

Version	Image address	Release date	Description	Impact
v1.10.4	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.10.4-bdab325-aliyun	2018-09-27	The issue that disks fail to be unmounted is fixed.	No impact on workloads

August 2018

Version	Image address	Release date	Description	Impact
v1.10.4	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.10.4-dfe877b-aliyun	2018-08-22	The NFS Options configuration is supported.	No impact on workloads

May 2018

Version	Image address	Release date	Description	Impact
v1.9.7	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.9.7-42e8198-aliyun	2018-05-31	The issue that disks are remounted when kubelet is restarted is fixed.	No impact on workloads

February 2018

Version	Image address	Release date	Description	Impact
v1.8.4	registry.cn-hangzhou.aliyuncs.com/acs/Flexvolume:v1.8.4-bcfda92-aliyun	2018-02-05	• By default, OSS volumes, NAS volumes, and disk volumes are supported. • The following authentication methods are supported: OSS buckets support AccessKey pairs and disks support Resource Access Management (RAM) roles. • When an ACK cluster is created, FlexVolume is automatically installed.	No impact on workloads

5.4.6. csi-local-plugin

csi-local-plugin is a Container Storage Interface (CSI) plug-in and is used to provision on-premises storage in edge Kubernetes clusters. This topic lists the latest changes to csi-local-plugin.

Introduction

csi-local-plugin is developed on top of the CSI drivers of Kubernetes. csi-local-plugin allows you to manage on-premises storage resources by using Logical Volume Manager (LVM). You can create, delete, mount, and unmount on-premises volumes by using persistent volumes (PVs) and persistent volume claims (PVCs).

Instruction

csi-local-plugin is developed on top of the standard CSI driver. It allows you to create PVCs by calling the Kubernetes API and specifying them in pod configurations.

Release notes

March 2021

Version	Image address	Release date	Description	Impact
v1.16.9-a9145c7-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/csi-plugin:v1.16.9-a9145c7-aliyun	2021-03-24	csi-local-plugin is released for the first time.	N/A

5.4.7. node-resource-manager

The node-resource-manager component is used to manage the resources of nodes in edge Kubernetes clusters. This topic introduces node-resource-manager, and describes the usage notes and release notes for node-resource-manager.

Introduction

node-resource-manager provides automatic management of the computing and storage resources of nodes. It allows you to manage storage resources by using Logical Volume Manager (LVM). For example, you can use LVM to automatically create volume groups (VGs) and persistent volumes (PVs) based on the configurations specified by the spec parameter.

Usage notes

For more information about how to use node-resource-manager, see Use LVM to manage local storage.

Release notes

October 2021

Version	Image address	Release date	Description	Impact
v1.18.8.0-5b1bdc2-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/node-resource-manager:v1.18.8.0-5b1bdc2-aliyun	2021-10-10	The unit testing feature is supported. Basic features are optimized.	No impact on workloads

March 2021

Version	Image address	Release date	Description	Impact
v1.0-ac86658-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/node-resource-manager:v1.0-ac86658-aliyun	2021-03-24	node-resource-manager is released for the first time.	No impact on workloads

5.4.8. flexvolume-windows

The flexvolume-windows component is a container storage interface used in Container Service for Kubernetes (ACK) clusters. This topic introduces flexvolume-windows and provides usage notes and release notes for the component.

Introduction

The flexvolume-windows component is a container storage interface that allows you to mount block storage, such as Alibaba Cloud disks, and Apsara File Storage NAS (NAS) to Windows containers in ACK clusters.

Usage notes

The Kubernetes version of the cluster must be v1.16.9 or later. You must select the FlexVolume plug-in as the volume plug-in when you create the cluster.

> **Note** By default, the FlexVolume plug-in is unavailable when you create clusters. You must to apply to be added to a whitelist.

You must manually install flexvolume-windows as a DaemonSet in the cluster. For more information, see Step 2: Install FlexVolume on a Windows node of Mount disks and SMB file systems to Windows pods.

Release notes

August 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.0	registry-vpc.${region}.aliyuncs.com/acs/flexvolume-windows:v1.0.0	2021-08-23	The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads

5.5. Networking

5.5.1. Terway

Terway is an open source Container Network Interface (CNI) plug-in developed by Alibaba Cloud. Terway works with Virtual Private Cloud (VPC) and allows you to use standard Kubernetes network policies to regulate how containers communicate with each other. You can use Terway to set up network connectivity within a Kubernetes cluster. The topic introduces Terway and provides the usage notes and release notes for the component.

Introduction

Terway is a CNI plug-in developed by Container Service for Kubernetes (ACK). The plug-in builds networks based on elastic network interfaces (ENIs) of Alibaba Cloud to make full use of cloud resources. Terway supports the use of extended Berkeley Packet Filter (eBPF) to accelerate network traffic and reduce latency. Terway supports standard Kubernetes network policies that define how containers communicate with each other and provides compatibility with Calico network policies.

In a cluster that has Terway installed, each pod has a separate network stack and is assigned a separate IP address. Pods on the same Elastic Compute Service (ECS) instance communicate with each other by forwarding packets inside the ECS instance. Pods on different ECS instances communicate with each other through ENIs in the VPC in which the ECS instances are deployed. This improves communication efficiency because no tunneling technologies, such as Virtual Extensible Local Area Network (VXLAN), are required to encapsulate packets.

Usage notes

For more information about how to use Terway, see Work with Terway.

Release notes

March 2022

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.2.2	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.2.2	2022-03-31	The ARP inspection issue in VLAN mode is fixed.	No impact on workloads
v1.2.1	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.2.1	2022-03-15	• The issue that calico-felix changes to the legacy iptables mode when calico-felix runs on Alibaba Cloud Linux 3 is fixed. The issue that calico-felix changes to the legacy iptables mode when calico-felix runs on Alibaba Cloud Linux 3 is fixed. • The issue that IP addresses cannot be revoked when metadata errors occur is fixed.	No impact on workloads

January 2022

Version	Image address	Release date	Description	Impact
v1.2.0	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.2.0	2022-01-11	• The efficiency of file locking when CNI is being executed is improved. • A Service can be accessed by using its backend pods in IPVLAN mode. • In IPVLAN mode, if you access an external IP address or a Server Load Balancer (SLB) instance from within a cluster, the traffic is routed to the backend Service. This feature is automatically enabled for newly created clusters. To enable load balancing within an existing cluster in Terway IPVLAN mode, see How do I enable load balancing within a cluster in Terway IPVLAN mode?. • The issue that false positive alerts are generated in Terway VPC mode is fixed. • The performance of calico-felix is improved.	No impact on workloads

December 2021

Version	Image address	Release date	Description	Impact
v1.1.1	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.1.1	2021-12-20	• The reclaiming of ENIs is improved. The issue that IPVLAN cannot be set up due to netns leaks when containerd is used is fixed. • An ENI can be associated with multiple security groups. For more information about how to attach the policy to the RAM role, see Associate multiple security groups with an ENI.	No impact on workloads

November 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.1.0	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.1.0	2021-11-22	• IPv4/IPv6 dual stack is supported. • Felix is updated to V3.20.2 • CNI is updated from V0.3.0 to V0.3.1. • The deployment template of Terway is compatible with Kubernetes 1.22. Terway of this version and later versions are compatible with Kubernetes 1.18 and later versions. • The issue that ENI configuration errors occasionally occur in IPVLAN mode is fixed. For more information, see #261.	• If Typha is deployed in your cluster to improve the performance of network policies, update Terway to this version. For more information about how to attach the policy to the RAM role, see Improve the performance of the NetworkPolicy feature for a large ACK cluster in Terway mode. • If Typha is not deployed in your cluster, your workloads are not affected.

September 2021

Version	Image address	Release date	Description	Impact
v1.0.10.443-gaa1bfcc-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.443-gaa1bfcc-aliyun	2021-09-14	Network latency is reduced. This update applies only to Terway in exclusive ENI mode. Terway in other modes does not require this update.	No impact on workloads

August 2021

Version	Image address	Release date	Description	Impact
v1.0.10.398-g63d2e57-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.398-g63d2e57-aliyun	2021-08-04	• The issue that network policies occasionally fail to take effect when Terway is used in inclusive ENI mode is fixed. • The Cilium plug-in is updated to V1.10.	No impact on workloads

July 2021

Version	Image address	Release date	Description	Impact
v1.0.10.390-g5f3c461-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.390-g5f3c461-aliyun	2021-07-02	• The issue that pod annotations cannot be modified when elastic IP addresses (EIPs) are used in Terway mode is fixed. • The issue that Terway cannot be started when a dedicated ENI is assigned to each pod in Terway mode is fixed. • The network namespace path is automatically adjusted in containerd containers. • Alibaba Cloud Linux 3 is supported.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact
v1.0.10.368-g2890967-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.368-g2890967-aliyun	2021-05-24	• The issue that the `Resource Invalid` alert occurs when EIPs are used is fixed. • Communication between pods and the node is supported in IPVLAN mode. Data can be transmitted within the node and does not need to be transmitted over the VPC in which the node is deployed. • Proactive checks and corrections of the `ip forwarding` setting are supported.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
v1.0.10.333-gfd2b7b8-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.333-gfd2b7b8-aliyun	2021-04-26	• The conflicts that occur when stateful applications use EIPs are fixed. • Hubble can be enabled in IPVLAN mode. For more information about Hubble, see What is Hubble.	No impact on workloads

March 2021

Version	Image address	Release date	Description	Impact
v1.0.10.323-g778c128-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.323-g778c128-aliyun	2021-03-22	• The number of API calls per minute is reduced to prevent API abuse. • Event alerts are optimized when security group configurations are inspected.	No impact on workloads

February 2021

Version	Image address	Release date	Description	Impact
v1.0.10.317-g0652857-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.317-g0652857-aliyun	2021-02-22	• The issue of identity leakage in Cilium is fixed. • The issue that occasionally occurs in the IP management of StatefulSet pods is fixed.	No impact on workloads
v1.0.10.309-g5314eee-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.309-g5314eee-aliyun	2021-02-05	• Errors in the configurations of security group rules are no longer automatically fixed. The system generates alerts only when errors are detected. You can follow the suggestions to fix the errors. • The issue that the IP addresses of pods may be reclaimed and reassigned to the pods when the nodes are overloaded is fixed.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
v1.0.10.301-g0115576-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.301-g0115576-aliyun	2021-01-21	New features: • Custom maximum transmission units (MTUs) are supported. • Routing based on the host network stack in exclusive ENI mode is supported. Fixed issues: • Packet loss due to traffic throttling of Terway. • Residual IP policies. • Incorrect counting of the number of IP addresses provided by ENIs.	No impact on workloads

December 2020

Version	Image address	Release date	Description	Impact
v1.0.10.280-gdc2cb6c-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.280-gdc2cb6c-aliyun	2020-12-25	• Proactive checks of Terway are supported. The following items are checked: ◦ ENIs and secondary IP addresses of ENIs. ◦ The consistency of security groups. ◦ The pod network configurations. ◦ The host network configurations. • Routing based on the host network stack in IPVLAN mode is supported.	No impact on workloads
v1.0.10.263-gdbe50a9-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.263-gdbe50a9-aliyun	2020-12-03	The Terway error that occurs in exclusive ENI mode is fixed.	No impact on workloads

November 2020

Version	Image address	Release date	Description	Impact
v1.0.10.261-g8342155-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.261-g8342155-aliyun	2020-11-27	• Synchronous calls to the ECS API are supported. • The issue that error messages are returned when the CNI plug-in is used is fixed.	No impact on workloads
v1.0.10.250-gb7bb10a-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.250-gb7bb10a-aliyun	2020-11-23	• The network issue of pods when the ENI driver fails to be loaded is fixed. • The status issue of the IP addresses that are allocated by ENIs when the ENI API is throttled is fixed.	No impact on workloads

October 2020

Version	Image address	Release date	Description	Impact
v1.0.10.247-g4cb77d0-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.247-g4cb77d0-aliyun	2020-10-26	ECS instances that are deployed on dedicated hosts are supported.	No impact on workloads

September 2020

Version	Image address	Release date	Description	Impact
v1.0.10.237-g6a0f948-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.237-g6a0f948-aliyun	2020-09-16	The time required to bind ENIs to pods is reduced.	No impact on workloads

August 2020

Version	Image address	Release date	Description	Impact
v1.0.10.221-g8d6386a-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.221-g8d6386a-aliyun	2020-08-11	IPVLAN and the eBPF are supported for network virtualization. To use this feature, Submit a ticket to apply to be added to the whitelist.	No impact on workloads

Version	Image address	Release date	Description	Impact
v1.0.10.213-g27145cc-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.213-g27145cc-aliyun	2020-08-04	The network issue of pods due to occasional ENI failures is fixed.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact
v1.0.10.208-gf3144bf-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.208-gf3144bf-aliyun	2020-07-20	• The issue that policy-based routes for nodes in inclusive ENI mode are exposed is fixed. • API calls over internal networks are supported. • The issue that pod IP addresses cannot be released when the number of IP addresses provided by the vSwitch reaches the upper limit is fixed. • The error report page that returns CNI errors is optimized.	No impact on workloads
v1.0.10.211-gef088a4-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.211-gef088a4-aliyun	2020-07-24	Cluster ID tags can be added to ENIs.	No impact on workloads

April 2020

Version	Image address	Release date	Description	Impact
v1.0.10.156-g8660a0f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.156-g8660a0f-aliyun	2020-04-22	• The cache efficiency is improved when ENIs are used. • The pre-installed Felix is updated to V3.5.8. • Reclaiming network resources from pods that are in the Completed or Failed state is supported.	No impact on workloads

February 2020

Version	Image address	Release date	Description	Impact
v1.0.10.139-g14a4f84-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.139-g14a4f84-aliyun	2020-02-12	The issue that pod creation requests occasionally time out is fixed.	No impact on workloads

January 2020

Version	Image address	Release date	Description	Impact
v1.0.10.133-g001396b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.133-g001396b-aliyun	2020-01-10	• The NetworkPolicy feature can be disabled. • IPVLAN is supported for network virtualization in inclusive ENI mode.	No impact on workloads

December 2019

Version	Image address	Release date	Description	Impact
v1.0.10.122-gd0be015-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.122-gd0be015-aliyun	2019-12-24	The efficiency of IP address allocation is optimized in inclusive ENI mode.	No impact on workloads

October 2019

Version	Image address	Release date	Description	Impact
v1.0.10.100-g92a3fa5-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.10.100-g92a3fa5-aliyun	2019-10-11	The issue that the host node is not ready when a large number of jobs concurrently request resources is fixed.	No impact on workloads

August 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.9.20-g35ae000-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.9.20-g35ae000-aliyun	2019-08-23	Kubernetes 1.14.6 is supported.	No impact on workloads

April 2019

Version	Image address	Release date	Description	Impact
v1.0.9.15-g3957085-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.9.15-g3957085-aliyun	2019-04-11	The issue that the Terway update occasionally fails is fixed.	No impact on workloads

March 2019

Version	Image address	Release date	Description	Impact
v1.0.9.14-ga0346bb-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/terway:v1.0.9.14-ga0346bb-aliyun	2019-03-28	• The issue that Terway fails to obtain the ENI information when throttling is applied to the meta server is fixed. • The issue that the **failed to move veth to host netns: file exists** error is returned when you create a container is fixed. • Periodic scanning is supported to check the status of ENIs. ENIs that are abnormally released are periodically reclaimed. • Health checks are supported. TCP port check is performed instead of HTTP path check.	No impact on workloads

Related information

- Work with Terway

5.5.2. Flannel

Flannel is a Container Network Interface (CNI) plug-in that you can use to create a virtual network for containers based on Virtual Private Cloud (VPC). You can use Flannel to enable internal communication in a Kubernetes cluster. This topic describes the release notes for Flannel.

Introduction

The Flannel network plug-in provided by Container Service for Kubernetes (ACK) allocates the pod CIDR block, which is independent of the VPC CIDR block. In an ACK cluster, the Flannel network plug-in works with the VPC in which the cluster is deployed. Packets are directly forwarded based on the VPC route table. This improves communication efficiency because no tunneling technologies such as Virtual Extensible Local Area Network (VXLAN) are required to encapsulate the packets. For more information, see Overview.

Usage note

Flannel is an open source CNI plug-in, which is simple and stable. You can use Flannel with VPC of Alibaba Cloud. This allows your clusters and containers to run in high-performance and stable networks. For more information, see Flannel.

Release notes

January 2022

Version	Image address	Release date	Description	Impact
v0.15.1.5-11d1c700-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/flannel:v0.15.1.5-11d1c700-aliyun	2022-01-25	Instances that use the ARM64 architecture are supported.	No impact on workloads

December 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.15.1.4-e02c8f12-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/flannel:v0.15.1.4-e02c8f12-aliyun	2021-12-01	• The API version of resources such as Authorization is updated to support Kubernetes 1.22. • Services can be exposed by using the HostPort method. • By default, the hairpin mode is enabled. By default, the hairpin mode is disabled for Flannel of earlier versions. For more information about how to manually enable the hairpin mode, see How do I enable a pod to access a Service that is used to expose the pod?.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.13.0.1-466064b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/flannel:v0.13.0.1-466064b-aliyun	2021-05-24	• Iptables is supported to enable compatibility with CentOS 8 and Alibaba Cloud Linux 3 and later. • Proactive checks and corrections of IP forwarding configurations are supported.	No impact on workloads

August 2019

Version	Image address	Release date	Description	Impact
v0.11.0.2-g6e46593e-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/flannel:v0.11.0.2-g6e46593e-aliyun	2019-08-02	The issue that source IP addresses are not retained after Flannel is upgraded is fixed.	No impact on workloads

5.5.3. Nginx Ingress Controller

This topic introduces the NGINX Ingress controller, and describes the usage notes and release notes for the NGINX Ingress controller.

Introduction

Introduction to Ingresses

How the NGINX Ingress controller works

Usage notes

For more information about the NGINX Ingress controller, see Use Ingresses to implement canary releases and blue-green releases and Use an Ingress controller to mirror network traffic.

Release notes

March 2022

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.1.2-aliyun.2	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v1.1.2-aliyun.2	2022-03-21	• The version of the NGINX component is rolled back to V1.19.9, which is the same as the version of open source NGINX. This NGINX version is more stable. • The following issue is fixed: The NGINX Ingress controller crashes if the `cors-allow-origin` configuration is invalid. • The following issue is fixed: The Ingresses that use the same webhook URL conflict with each other when the system checks the webhook URLs of the Ingresses that belong to different IngressClasses. • The following issue is fixed: InitContainer modifies the kernel parameters of nodes if hostNetwork is set to true. • The CVE-2022-0778 and CVE-2022-23308 vulnerabilities are patched.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

January 2022

Version	Image address	Release date	Description	Impact
v1.1.0-aliyun.2	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v1.1.0-aliyun.2	2022-01-12	• The Application High Availability Service (AHAS) Sentinel plug-in is updated and the Java module is replaced by the C++ module. This greatly improves performance. • Protocol Buffers (Protobuf) is used to communicate with the Kubernetes API server of a cluster. This improves communication efficiency.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

December 2021

Version	Image address	Release date	Description	Impact
v1.1.0-aliyun.1	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v1.1.0-aliyun.1	2021-12-17	• **NGINX Ingress controller V1.X.X supports only Container Service for Kubernetes (ACK) clusters that run Kubernetes V1.20.0 and later. For ACK clusters that run earlier Kubernetes versions, you must use NGINX Ingress controller V0.X.X.** • `networking v1` Ingresses are used to support ACK clusters that run Kubernetes 1.22 and later. • You can specify multiple origins in the `cors-allow-origin` field. Requested resources are fetched based on the specified origins. • `Session affinity` can be enabled to define the behavior of canaries. You can also reset to the default behavior of canaries. • Canaries can be configured even when no host is specified. • Admission webhooks are accelerated. • Stability issues are fixed. For more information, see Ingress-NGINX changelog.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

October 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.44.0.9-7b9e93e7e-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.44.0.9-7b9e93e7e-aliyun	2021-10-28	• The allow-snippet-annotations annotation is added to reduce the impact of vulnerability CVE-2021-25742. For more information, see Vulnerability fixed: CVE-2021-25742. • SSL builtin cache is disabled to prevent potential memory leaks. • The following vulnerabilities are fixed: CVE-2021-22945, CVE-2021-22946, CVE-2021-3711, and CVE-2021-3712. For more information, see CVE-2021-22945, CVE-2021-22946, CVE-2021-3711, and CVE-2021-3712. • The AHAS sentinel SDK is updated to V1.9.7.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

September 2021

Version	Image address	Release date	Description	Impact

v0.44.0.5-e66e17ee3-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.44.0.5-e66e17ee3-aliyun	2021-09-06	• The AHAS sentinel plug-in is updated.
			  ◦ The performance and stability are improved.
			  ◦ Traffic throttling for clusters is supported.
			• Vulnerability CVE-2021-36159 is fixed. For more information, see CVE-2021-36159.
			• By default, the kernel parameter kernel.core_uses_pid is disabled. This prevents coredump files from occupying excessive disk space.
				The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

June 2021

Version	Image address	Release date	Description	Impact
v0.44.0.3-8e83e7dc6-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.44.0.3-8e83e7dc6-aliyun	2021-06-01	Vulnerability CVE-2021-23017 is fixed. For more information, see Updating NGINX for a DNS Resolver Vulnerability (CVE-2021-23017).	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

April 2021

Version	Image address	Release date	Description	Impact
v0.44.0.2-abf1c6fe4-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.44.0.2-abf1c6fe4-aliyun	2021-04-01	Compatibility with the the_real_ip field in the `log_format` parameter of NGINX Ingress controller V0.30 and earlier is added.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

March 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.44.0.1-5e842447b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.44.0.1-5e842447b-aliyun	2021-03-08	• By default, validating admission webhooks are enabled. For more information, see How the NGINX Ingress controller works. • Validity check is performed on the value of the `service-weight` annotation. • The performance of persistent connections and short-lived connections is increased by 20% to 50%. • Online Certificate Status Protocol (OCSP) stapling is supported. • LuaJIT is updated to V2.1.0. • NGINX is updated to V1.19.6. • Alpine Linux is updated to V3.13 for base images. • CVE vulnerabilities related to OpenSSL are fixed. • By default, Transport Layer Security (TLS) 1.3 is enabled. **Note** By default, only TLS 1.2 and TLS 1.3 are supported by HTTPS. For more information about how to enable HTTPS to support TLS 1.0 and TLS 1.1, see Which SSL or TLS protocol versions are supported by Ingresses?. • The Kubernetes version must be 1.16 or later. • The NGINX Ingress controller is updated based on open source Ingress-NGINX 0.44.0. For more information, see Ingress-NGINX changelog.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

April 2020

Version	Image address	Release date	Description	Impact

v0.30.0.1-5f89cb606-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.30.0.1-5f89cb606-aliyun	2020-04-02	• FastCGI Backend is supported.
			• By default, the Dynamic SSL Cert Update mode is enabled.
			• Traffic mirroring is supported.
			• NGINX is updated to V1.17.8 and OpenResty is updated to V1.15.8. The operating system of base images is updated to Alpine Linux.
			• Ingress validating admission webhooks are supported.
			• The following vulnerabilities are fixed: CVE-2018-16843, CVE-2018-16844, CVE-2019-9511, CVE-2019-9513, and CVE-2019-9516.
			• Major updates:
			  ◦ The lua-resty-waf, session-cookie-hash, and force-namespace-isolation configurations are deprecated.
			  ◦ The data type of x-forwarded-prefix is changed from BOOLEAN to STRING.
			  ◦ The the_real_ip field in the log-format parameter will be deprecated in the next version and replaced with the remote_addr field.
			• The NGINX Ingress controller is updated based on Ingress-NGINX 0.30.0. For more information about the updates, see Ingress-NGINX changelog.
				The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

October 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.22.0.5-552e0db-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.22.0.5-552e0db-aliyun	2019-10-24	Wildcard domain names, whitelists, and rewrite rules are supported if you enable dynamic update for NGINX upstream servers.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

July 2019

Version	Image address	Release date	Description	Impact
v0.22.0.4-5a14d4b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.22.0.4-5a14d4b-aliyun	2019-07-18	Canary release rules are optimized and the Perl regular expressions are supported.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

April 2019

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.22.0.3-da10b7f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.22.0.3-da10b7f-aliyun	2019-04-25	• The NGINX Ingress controller is updated based on Ingress-NGINX 0.22.0. For more information about the updates, see Ingress-NGINX. • Blue-green releases and canary releases are supported if you enable dynamic update for NGINX upstream servers. • By default, dynamic update is enabled for NGINX upstream servers. • Major updates: Capture groups are used for rewrite-target annotations. For more information, see rewrite-target. For more information about how to smoothly update the NGINX Ingress controller, visit GitHub.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

January 2019

Version	Image address	Release date	Description	Impact
v0.20.0.2-cc39f1b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.20.0.2-cc39f1b-aliyun	2019-01-17	• The default number of NGINX worker processes is limited. This avoids the issue that an excessive number of NGINX processes occupy host resources. • The port numbers of Services that route traffic to the old application version and the new application version can be different during blue-green releases and canary releases. • The NGINX configuration verification failure is fixed when no pod is active on the backend servers of the new application version during canary releases. • The issue that Ingress address endpoints are not updated due to failed connections to the Kubernetes API server is fixed.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

November 2018

Version	Image address	Release date	Description	Impact
v0.20.0.1-4597ce2-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-ingress-controller:v0.20.0.1-4597ce2-aliyun	2018-11-29	• The NGINX Ingress controller is updated based on Ingress-NGINX 0.20.0. For more information about the updates, see Ingress-NGINX. • NGINX is updated to V1.15.6 and HTTP/2-related vulnerabilities are fixed. • Regular expressions are supported by the path parameter. • The default-http-backend Service is removed and custom default backend Services are supported. • Blacklists based on IP addresses, user agents, and referer headers are supported. • The default permissions are optimized and the privileged permissions are removed. • Apache JServ Protocol (AJP) is supported.	The update may temporarily interrupt your service. We recommend that you update the NGINX Ingress controller during off-peak hours.

5.5.4. kube-flannel-ds

kube-flannel-ds is a network component used to configure overlay networks for containers in edge Kubernetes clusters. This topic introduces kube-flannel-ds and lists the latest changes to kube-flannel-ds.

Introduction

The kube-flannel-ds component is automatically installed when you create an edge Kubernetes cluster. The component is deployed as a DaemonSet and used to configure Virtual Extensible LAN (VXLAN) devices, forwarding database (FDB) tables, and routing information for cluster nodes. This allows you to enable communication between containers on the same node or different nodes.

Usage notes

The kube-flannel-ds component is automatically configured in the Container Network Interface (CNI) plug-in. You can directly use kube-flannel-ds when you call the Kubernetes API to create applications.

Release notes

May 2021

Version	Image address	Release date	Description	Impact
v0.11.0.6-1283a29-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/flannel:v0.11.0.6-1283a29-aliyun	May 2021	A startup parameter named enable-cloud-edge-isolate is added. Overlay routing between edge nodes and cloud nodes is supported.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
v0.11.0.5-437a359-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/flannel:v0.11.0.5-437a359-aliyun	January 2021	None.	No impact on workloads

5.5.5. ACK NodeLocal DNSCache

This topic introduces the ACK NodeLocal DNSCache component and describes the usage notes and release notes for the component.

Introduction

ACK NodeLocal DNSCache is a local DNS caching solution that is developed based on the open source NodeLocal DNSCache project. You can deploy ACK NodeLocal DNSCache by installing the ack-node-local-dns Helm chart. This solution consists of a DNS caching agent that runs as a DaemonSet and an admission controller that runs as a Deployment to dynamically inject DNSConfig. The admission controller listens on pod creation requests and dynamically modifies DNSConfig. This enables pods to use local DNS cache to accelerate DNS lookups.

Usage notes

You can use ACK NodeLocal DNSCache to enable pods to access the local DNS cache on nodes. For more information, see Configure NodeLocal DNSCache.

Release notes

February 2022

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
1.5.0	• DNS caching agent: registry.cn-hangzhou.aliyuncs.com/acs/k8s-dns-node-cache:v1.21.4.3-a38fc90-aliyun • Admission controller: registry.cn-hangzhou.aliyuncs.com/acs/node-local-dns-admission-controller:v1.1.1-aliyun	2022-02-22	• The ARM64 architecture is supported. • The `serve_stale` disaster recovery mode is supported. • Iptables is supported to enable compatibility with CentOS 8, and Alibaba Cloud Linux 3 and later versions.	N/A

November 2021

Version	Image address	Release date	Description	Impact
1.4.0	• DNS caching agent: registry.cn-hangzhou.aliyuncs.com/acs/k8s-dns-node-cache:v1.15.13-6-7e6778ac • Admission controller: registry.cn-hangzhou.aliyuncs.com/acs/node-local-dns-admission-controller:v1.1.0-b1bd0e8-aliyun	2021-11-24	The API version of the MutatingWebhookConfiguration resource is updated to support Kubernetes 1.22.	N/A

September 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
1.3.5	• DNS caching agent: registry.cn-hangzhou.aliyuncs.com/acs/k8s-dns-node-cache:v1.15.13-6-7e6778ac • Admission controller: registry.cn-hangzhou.aliyuncs.com/acs/node-local-dns-admission-controller:v1.0.3-8fe673f-aliyun	2021-09-23	The CPU usage of the admission controller is limited to 1 vCPU by default. This provides enhanced security.	N/A
1.3.4	• DNS caching agent: registry.cn-hangzhou.aliyuncs.com/acs/k8s-dns-node-cache:v1.15.13-6-7e6778ac • Admission controller: registry.cn-hangzhou.aliyuncs.com/acs/node-local-dns-admission-controller:v1.0.2-8b46b2f-aliyun	2021-09-16	The memory usage is limited by default. The image is pulled over the internal network by default.	N/A

April 2021

Version	Image address	Release date	Description	Impact
1.3.3	• DNS caching agent: registry.cn-hangzhou.aliyuncs.com/acs/k8s-dns-node-cache:v1.15.13-6-7e6778ac • Admission controller: registry.cn-hangzhou.aliyuncs.com/acs/node-local-dns-admission-controller:v1.0.2-8b46b2f-aliyun	2021-04-21	The first version is released.	N/A

5.5.6. CoreDNS

CoreDNS is the default plug-in used to implement Domain Name System (DNS)-based service discovery in Container Service for Kubernetes (ACK) clusters and ACK edge clusters. The topic introduces CoreDNS and provides the usage notes and release notes for CoreDNS.

Introduction

CoreDNS is a plug-in used to implement DNS-based service discovery in ACK clusters and ACK edge clusters. CoreDNS follows the specifications of DNS-based service discovery in Kubernetes. For more information, see Kubernetes DNS-Based Service Discovery. CoreDNS provides DNS resolution capabilities within Kubernetes clusters. ACK maintains a mapping between CoreDNS versions and Kubernetes versions. When you create or upgrade a Kubernetes cluster, the CoreDNS version that is installed or upgraded depends on the Kubernetes version of the cluster. For more information about the mapping, see CoreDNS version in Kubernetes.

Usage note

For more information about the features and usage notes of CoreDNS, see Configure DNS resolution.

> **Notice**
>
> Before you upgrade CoreDNS, take note of the following items:
>
> - Read and understand the CoreDNS upgrade considerations. For more information, see Configure ACK to automatically update CoreDNS.
> - Back up the CoreDNS ConfigMap in the kube-system namespace.

Release notes

February 2022

Version	Applicable cluster	Image address	Release date	Description	Impact

Version	Applicable cluster	Image address	Release date	Description	Impact
v1.8.4.3-644f4735-aliyun	ACK managed clusters, ACK dedicated clusters, and ASK clusters that run Kubernetes 1.20.4 or later	registry.{{.Region}}.aliyuncs.com/acs/coredns:v1.8.4.3-644f4735-aliyun	2022-02-22	This image version is in canary release. • Caching for DNS resolution results of the ServError type is disabled. • Pod anti-affinity settings based on hostnames are changed from preferred to required. The system is forced to schedule pods based on the anti-affinity settings.	The system is forced to schedule pods based on the anti-affinity settings. If the number of CoreDNS pods is greater than that of nodes, some CoreDNS pods change to the Pending state. You must scale out cluster nodes or scale in CoreDNS pods before you upgrade to this image version.

January 2022

Version	Applicable cluster	Image address	Release date	Description	Impact
v1.8.4.2-7d597cff-aliyun	ACK managed clusters, ACK dedicated clusters, and serverless Kubernetes (ASK) clusters that run Kubernetes 1.20.4 or later	registry.{{.Region}}.aliyuncs.com/acs/coredns:v1.8.4.2-3a376cc-aliyun	2022-01-10	• Custom parameters are supported. • By default, log parsing is enabled.	No impact on workloads

October 2021

Version	Applicable cluster	Image address	Release date	Description	Impact
v1.8.4.1-3a376cc-aliyun	ACK managed clusters, ACK dedicated clusters, and ASK clusters that run Kubernetes 1.20.4 or later	registry.{{.Region}}.aliyuncs.com/acs/coredns:v1.8.4.1-3a376cc-aliyun	2021-10-26	• EndpointSlice objects can be monitored. • IPv6 addresses are supported by DNS resolutions.	No impact on workloads

July 2021

Version	Applicable cluster	Image address	Release date	Description	Impact
v1.7.0.0-f59c03d-aliyun	ACK managed clusters, ACK dedicated clusters, and ASK clusters	registry.{{.Region}}.aliyuncs.com/acs/coredns:v1.7.0.0-f59c03d-aliyun	2021-07-08	• The default duration of graceful shutdown for CoreDNS and the memory limit of containers in the CoreDNS Deployment are modified. • The names of metrics are updated. If your monitoring system is reliant on CoreDNS metrics, you must update the metric names. For more information, see Metric changes. • The issue that only the first upstream DNS server specified in the forward plug-in is used is fixed. • The deprecated upstream plug-in is no longer compatible. If the upstream plug-in is specified in the Corefile configurations, the upstream plug-in will be automatically deleted in a secure way when CoreDNS is upgraded.	If you modify the DNS configuration file */etc/resolv.conf* on your Elastic Compute Service (ECS) instance, CoreDNS pods are upgraded or recreated based on the modified */etc/resolv.conf* file. You must make sure that the DNS servers specified in this file work as normal before you upgrade CoreDNS.

April 2021

Version	Applicable cluster	Image address	Release date	Description	Impact
1.6.7.edge	ACK edge clusters	registry.{{.Region}}.aliyuncs.com/acs/coredns:1.6.7.edge	2021-04-23	This version is optimized based on open source CoreDNS 1.6.7. For more information, see CoreDNS-1.6.7 Release.	No impact on workloads

March 2021

Version	Applicable cluster	Image address	Release date	Description	Impact
1.7.0	ACK managed clusters, ACK dedicated clusters, and ASK clusters	registry.{{.Region}}.aliyuncs.com/acs/coredns:1.7.0	2021-03-18	• The deprecated upstream plug-in is no longer compatible. If the upstream plug-in is specified in the Corefile configurations, the upstream plug-in will be automatically deleted in a secure way when CoreDNS is upgraded. • The names of metrics are updated. If your monitoring system is reliant on CoreDNS metrics, you must update the metric names. For more information, see Metric changes. • The issue that only the first upstream DNS server specified in the forward plug-in is used is fixed.	If you modify the DNS configuration file */etc/resolv.conf* on your ECS instance, CoreDNS pods are upgraded or recreated based on the modified */etc/resolv.conf* file. You must make sure that the DNS servers specified in this file work as normal before you upgrade CoreDNS.

November 2018

Version	Applicable cluster	Image address	Release date	Description	Impact

Version	Applicable cluster	Image address	Release date	Description	Impact
1.6.7	ACK managed clusters, ACK dedicated clusters, and ASK clusters	registry.{{.Region}}.aliyuncs.com/acs/coredns:1.6.7	2018-11-28	This version is optimized based on open source CoreDNS 1.6.7. For more information, see CoreDNS-1.6.7 Release.	No impact on workloads

5.5.7. kube-flannel-ds-windows

This topic introduces kube-flannel-ds-windows and describes the usage notes and release notes for kube-flannel-ds-windows.

Introduction

kube-flannel-ds-windows is a container network plug-in used by managed Kubernetes clusters. kube-flannel-ds-windows is used to set up a l2bridge network to connect Windows containers.

Usage notes

If you choose Flannel as the network plug-in when you create a managed Kubernetes cluster, kube-flannel-ds-windows is automatically deployed as a DaemonSet in the cluster. kube-flannel-ds-windows is dependent on the privileged proxy processes on Windows nodes. kube-flannel-ds-windows can be deployed only on Windows nodes that are added with the `windows.alibabacloud.com/deployment-topology=2.0` label. By default, the `windows.alibabacloud.com/deployment-topology=2.0` label is added to a Windows node when the node is created.

If your cluster does not contain Windows nodes or the Windows nodes in your cluster are not added with the `windows.alibabacloud.com/deployment-topology=2.0` label, no replicated pods are provisioned for kube-flannel-ds-windows.

If a Windows node is added with the `windows.alibabacloud.com/deployment-topology=2.0` label but no replicated pods are provisioned for kube-flannel-ds-windows, you must redeploy kube-flannel-cfg-windows and kube-flannel-ds-windows.

1.
2.
3.
4. In the left-side navigation pane, choose **Operations > Add-ons**.
5. On the **Add-ons** page, find kube-flannel-ds, click , and then click **View in YAML**.
6. In the **View in YAML** panel, copy the configurations of kube-flannel-cfg-windows and kube-flannel-ds-windows.
7. In the left-side navigation pane of the cluster details page, choose **Workloads > Deployments**.
8. On the **Deployments** page, click **Create from YAML** in the upper-right corner.
9. Set **Sample Template** to **Custom**. Paste the configurations of kube-flannel-cfg-windows and

kube-flannel-ds-windows to the **Template** field. Click **Create**.

Release notes

August 2021

Version	Image address	Release date	Description	Impact
v0.15.1.1-6e1a57c8-aliyun	registry-vpc.${region}.aliyuncs.com/acs/flannel-windows:v0.15.1.1-6e1a57c8-aliyun	2022-01-21	• Flannel is updated. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads.

August 2021

Version	Image address	Release date	Description	Impact
v0.13.1-aliyun.3	registry-vpc.${region}.aliyuncs.com/acs/flannel-windows:v0.13.1-aliyun.3	2021-08-22	• The base images are updated. The original version number is overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads.

July 2021

Version	Image address	Release date	Description	Impact
v0.13.1-aliyun.3	registry-vpc.${region}.aliyuncs.com/acs/flannel-windows:v0.13.1-aliyun.3	2021-07-05	• The base images are updated. The original version number is overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1999), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1052).	No impact on workloads.

June 2021

Version	Image address	Release date	Description	Impact
v0.13.1-aliyun.3	registry-vpc.${region}.aliyuncs.com/acs/flannel-windows:v0.13.1-aliyun.3	2021-06-17	• The base images are updated. • The following Windows systems are supported: Windows Server version 1809 (10.0.17763.1935), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.985).	No impact on workloads.

May 2021

Version	Image address	Release date	Description	Impact
v0.13.1-aliyun.2	registry-vpc.${region}.aliyuncs.com/acs/flannel-windows:v0.13.1-aliyun.2	2021-05-13	• Container Network Interface (CNI) configurations can be automatically shifted when containerd is used as the container runtime. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1879), Windows Server version 1909 (10.0.18363.1500), and Windows Server version 2004 (10.0.19041.928).	No impact on workloads.

April 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.13.1-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/flannel-windows:v0.13.1-aliyun.1	2021-04-22	• Rancher Wins is automatically installed on a Windows node with the `windows.alibabacloud.com/deployment-topology=2.0` label. You can deploy Flannel on a Windows node by using Rancher Wins. • kube-flannel-ds-windows can be used to construct l2bridge and overlay networks. By default, Alibaba Cloud uses l2bridge networks. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1879), Windows Server version 1909 (10.0.18363.1500), and Windows Server version 2004 (10.0.19041.928).	No impact on workloads.

5.6. Security

5.6.1. aliyun-acr-credential-helper

You can use the aliyun-acr-credential-helper component to pull private images without a password from instances of Container Registry Enterprise Edition and Personal Edition. This topic describes the features, usage notes, and release notes for aliyun-acr-credential-helper.

Introduction

aliyun-acr-credential-helper retrieves the required information from the acr-configuration ConfigMap that is created in the kube-system namespace and then pulls private images. The following features are supported:

- You can use aliyun-acr-credential-helper to pull private images from instances of Container Registry Enterprise Edition and Personal Edition.
- You can use aliyun-acr-credential-helper to pull private images from your Container Registry instances. You can also pull private images from other accounts after authorization or by using an AccessKey ID and the corresponding AccessKey secret.
- You can use aliyun-acr-credential-helper to pull private images from Container Registry instances that are deployed in different regions.

Usage notes

For more information about how to use aliyun-acr-credential-helper, see Use the aliyun-acr-credential-helper component to pull images without a password or Pull images without a password in a self-managed Kubernetes cluster.

Release notes

March 2022

Version	Image address	Release date	Description	Impact
v22.03.25.1-efe240e-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v22.03.25.1-efe240e-aliyun	2022-03-25	• The version conflict that occurs when multiple controllers update a Secret at the same time is fixed. • Auto scaling is supported by default node pools.	This update has no negative impact on workloads.

January 2022

Version	Image address	Release date	Description	Impact
v22.01.04.0-f637776-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v22.01.04.0-f637776-aliyun	2022-01-04	• The Advanced RISC Machine (ARM) architecture is supported. • aliyun-acr-credential-helper is no longer automatically installed when you create a cluster. You must manually install aliyun-acr-credential-helper after you create a cluster. • The performance of aliyun-acr-credential-helper is improved.	This update has no negative impact on workloads.

November 2021

Version	Image address	Release date	Description	Impact
v21.11.15.0-19d8bc1-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v21.11.15.0-5d5867b-aliyun	2021-11-15	The API version of ClusterRoles is updated.	This update has no negative impact on workloads.

September 2021

Version	Image address	Release date	Description	Impact
v21.09.22.0-450db22-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v21.09.22.0-450db22-aliyun	2021-09-22	A cluster event is generated when an error occurs on the component.	This update has no negative impact on workloads.

January 2021

Version	Image address	Release date	Description	Impact
v21.01.26.0-9ac7d9b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v21.01.26.0-9ac7d9b-aliyun	2021-01-26	The authentication methods are optimized.	This update has no negative impact on workloads.

August 2020

Version	Image address	Release date	Description	Impact
v20.08.20.0-c2da10b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v20.08.20.0-c2da10b-aliyun	2020-08-24	The issue that private images fail to be pulled because the token is expired is fixed.	This update has no negative impact on workloads.

July 2020

Version	Image address	Release date	Description	Impact
v20.07.13.0-2866ccd-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v20.07.13.0-2866ccd-aliyun	2020-07-13	New features: • API operations can be called over internal networks. • Images can be pulled by using a Secret that contains an AccessKey ID and an AccessKey secret. Optimization: The required number of API calls to pull an image is reduced.	This update has no negative impact on workloads.

March 2020

Version	Image address	Release date	Description	Impact
v20.03.16.0-36d5d7e-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-credential-helper:v20.03.16.0-36d5d7e-aliyun	2020-03-16	New features: Private images can be pulled from a different account.	This update has no negative impact on workloads.

5.6.2. kritis-validation-hook

kritis-validation-hook is a key component that is used to verify the signatures of container images. This topic describes the features, usage notes, and release notes for kritis-validation-hook.

Introduction

kritis-validation-hook is a key component that is used to verify the signatures of container images. You can use the signature verification feature to ensure that only images signed by trusted authorities are deployed. This reduces the risk of malicious code execution. For more information about kritis-validation-hook, see Introduction to kritis-validation-hook.

Usage notes

For more information about how to use kritis-validation-hook, see Use kritis-validation-hook to automatically verify the signatures of container images.

Release notes

December 2021

Version	Image address	Release date	Description	Impact
v0.6.0.5-gce1cc2d-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.6.0.5-gce1cc2d-aliyun	2021-12-17	Kubernetes 1.22 is supported. v0.6.0.5-gce1cc2d-aliyun and later versions support only Kubernetes 1.16 and later.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

November 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.5.0.6-g525daee-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.5.0.6-g525daee-aliyun	2021-11-15	• A new image signature format is supported by Container Registry. • The ARM64 architecture is supported.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

June 2021

Version	Image address	Release date	Description	Impact
v0.4.0.1-gb2862c4-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.4.0.1-gb2862c4-aliyun	2021-06-10	New feature: kritis-validation-hook can be installed in registered Kubernetes clusters.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

March 2021

Version	Image address	Release date	Description	Impact
v0.3.1.4-ga89b624-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.3.1.4-ga89b624-aliyun	2021-03-24	New features: The signatures of images stored in repositories whose names contain forward slashes (/) can be verified.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

November 2020

Version	Image address	Release date	Description	Impact
v0.2.7.2-g5fa671a-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.2.7.2-g5fa671a-aliyun	2020-11-24	The signature verification whitelist feature is supported. kritis-validation-hook does not verify the signatures of images that are included in a signature verification whitelist.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.
v0.2.6.4-g94b0940-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.2.6.4-g94b0940-aliyun	2020-11-16	New features: Signature verification is supported for Container Service for Kubernetes (ACK) images whose versions are immutable. For more information, see Configure a repository to be immutable.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

August 2020

Version	Image address	Release date	Description	Impact
v0.2.5.26-g75d5297-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.2.5.26-g75d5297-aliyun	2020-08-12	• If a container image fails to pass signature verification, a cluster event is generated in the kube-system namespace. The cause of the event is **FailedKritisAdmission**. • The dry run mode is supported. By default, this mode is disabled. If the dry run mode is enabled, container images that fail to pass signature verification can be deployed. If an image that fails to pass signature verification is deployed, a cluster event is generated in the kube-system namespace. The cause of the event is **DryRunKritisAdmission**.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

Version	Image address	Release date	Description	Impact

June 2020

Version	Image address	Release date	Description	Impact
v0.2.4.1-ge5c1265-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.2.4.1-ge5c1265-aliyun	2020-06-22	The signatures of Container Registry images stored in regions other than the current region can be verified.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

April 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.2.3.1-00e70883-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.2.3.1-00e70883-aliyun	2020-04-07	Performance is improved and log content is optimized.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

March 2020

Version	Image address	Release date	Description	Impact
v0.2.2.3-fe8a6319-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kritis-server:v0.2.2.3-fe8a6319-aliyun	2020-03-18	kritis-validation-hook is integrated with Container Registry. You can verify the signatures of images that are signed by Key Management Service (KMS). This ensures that only trusted images are deployed in ACK clusters.	If exceptions occur when the system upgrades kritis-validation-hook, cluster resources may fail to be updated. We recommend that you upgrade the component during off-peak hours.

5.6.3. security-inspector

The security-inspector component is a key component for performing security inspections. This topic describes the features, usage notes, and release notes for security-inspector.

Introduction

You can use security-inspector to scan workload configurations from various perspectives. This helps you better understand the security risks of your workloads. The following figure shows the architecture of security-inspector.

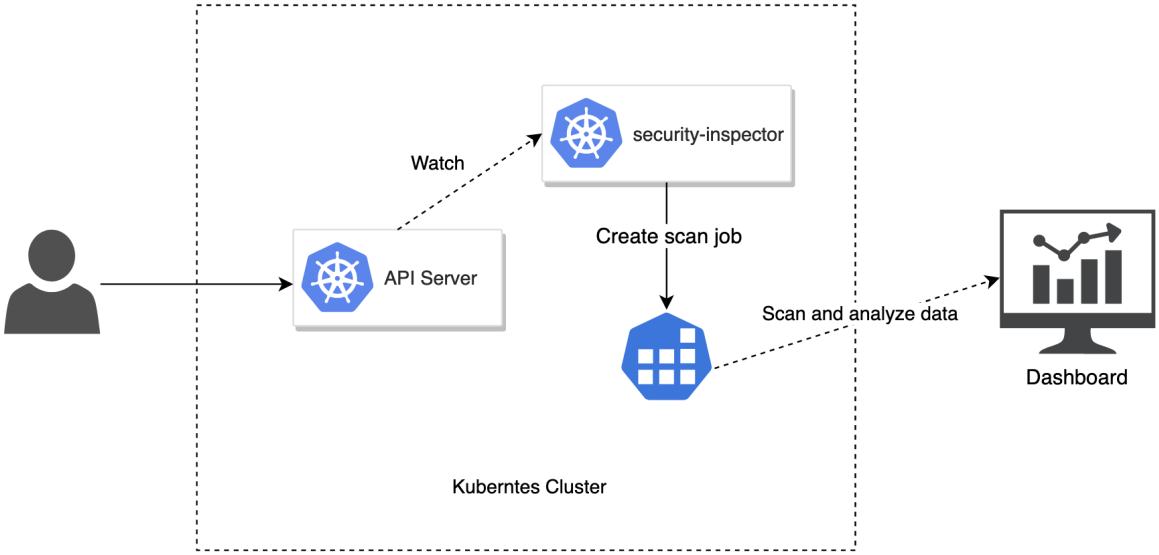

Usage note

security-inspector provides the following inspection features:

- security-inspector uses Polaris to perform security inspections. This allows you to detect security risks of workload configurations in your cluster in real time.

  > **Note** Polaris is an open source project that is used to identify security risks of workload configurations in a Kubernetes cluster. For more information, see Polaris.

- security-inspector can scan workload configurations from various perspectives and provide reports that contain the following information: health checks, images, networks, resources, and security. This allows you to better understand the security risks of your applications in real time and provides security suggestions to reinforce your system. For more information, see Use the inspection feature to detect security risks in the workloads of an ACK cluster.

Release notes

February 2022

Version	Image address	Release date	Description	Impact
v0.8.0.0-gb0edd1d-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.8.0.0-gb0edd1d-aliyun	2022-02-15	• The severity level of the `privilegeEscalationAllowed` inspection item is set to medium. • Support for clusters of Kubernetes 1.16 is optimized and the issue caused by #84880 is fixed.	No impact on workloads

December 2021

Version	Image address	Release date	Description	Impact
v0.7.0.5-g8cc37b6-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.7.0.5-g8cc37b6-aliyun	2021-12-03	• Kubernetes 1.22 is supported. security-inspector 0.7.0.5 and later versions support only clusters of Kubernetes 1.16 and later. • The ARM64 architecture is supported.	No impact on workloads

September 2021

Version	Image address	Release date	Description	Impact
v0.6.0.4-gc12ad66-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.6.0.4-gc12ad66-aliyun	2021-09-20	• Center for Internet Security (CIS) Kubernetes V1.20 Benchmark v1.0.0 is supported. For more information, see Safety patrol inspection. • Case sensitivity is removed for the capabilitiesAdded inspection item. For more information, see Use the inspection feature to detect security risks in the workloads of an ACK cluster.	No impact on workloads

June 2021

Version	Image address	Release date	Description	Impact
v0.5.0.2-g5e33765-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.5.0.2-g5e33765-aliyun	2021-06-24	The issue that inspection reports are not displayed as normal when one Log Service project is shared among multiple clusters is fixed.	No impact on workloads

March 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.4.0.0-g541eb31-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.4.0.0-g541eb31-aliyun	2021-03-15	• The CIS Kubernetes benchmark is supported. For more information, see Safety patrol inspection. • The following Kubernetes events are added. You can find the events in event center of your cluster when a scan is triggered. ◦ SecurityInspectorConfigAuditStart: Configuration inspection is started. ◦ SecurityInspectorConfigAuditFinished: Configuration inspection is completed. ◦ SecurityInspectorConfigAuditHighRiskFound: High-risk configurations are found after configuration inspection is completed. ◦ SecurityInspectorBenchmarkStart: The benchmark check is started. ◦ SecurityInspectorBenchmarkFinished: The benchmark check is completed. ◦ SecurityInspectorBenchmarkFailedCheckFound: Failed inspection items are found after the benchmark check is completed.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
v0.3.0.2-gcb49252-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.3.0.2-gcb49252-aliyun	2021-01-05	Permissions of anonymous users can be scanned to detect risky role-based access control (RBAC) permissions that are granted to the users.	No impact on workloads

December 2020

Version	Image address	Release date	Description	Impact
v0.2.0.22-gd1fbaff-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.2.0.22-gd1fbaff-aliyun	2020-12-16	• The CustomResourceDefinition (CRD) resource can be used to store the latest inspection results. • Specified inspection items can be enabled or disabled based on your needs. • The workload whitelist feature is supported.	No impact on workloads

July 2020

Version	Image address	Release date	Description	Impact
v0.1.0.3-g69f71f6-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/security-inspector:v0.1.0.3-g69f71f6-aliyun	2020-07-06	Inspection tasks can be manually triggered to inspect the workloads in your cluster and generate inspection reports.	No impact on workloads

5.6.4. gatekeeper

The gatekeeper component facilitates the management and enforcement of policies executed by Open Policy Agent (OPA) in Kubernetes clusters. This allows you to manage the labels of namespaces. This topic describes the features and usage notes of gatekeeper. It also lists the latest changes to gatekeeper.

Introduction

OPA is an open source policy engine that is commonly used to implement policies in stacks in a standardized and context-aware manner. You can use the gatekeeper component to manage and implement OPA policies, and manage labels of namespaces in ACK clusters. For more information about OPA, see Open Policy Agent. The following figure shows the architecture of gatekeeper.

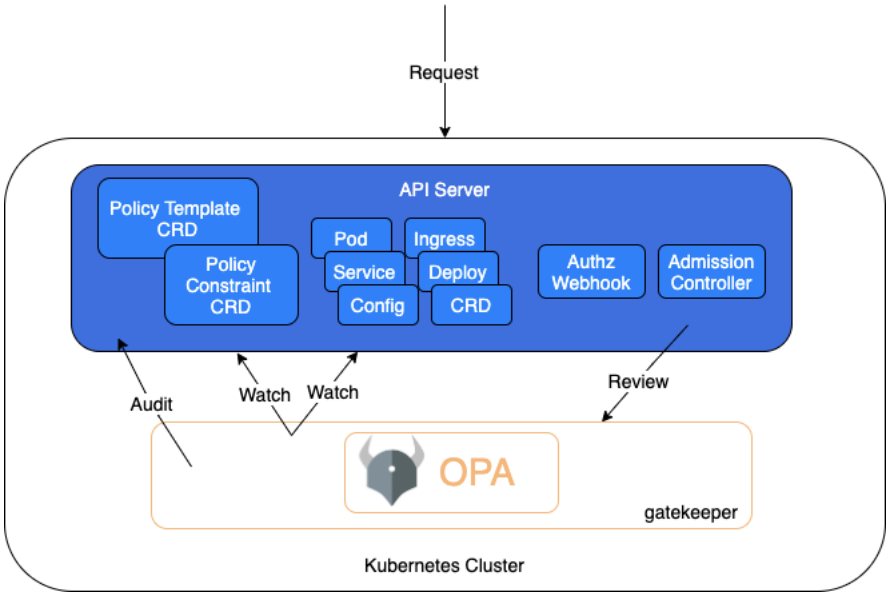

Usage notes

You can use gatekeeper to constrain pod deployments in specified namespaces based on labels. In this example, a Constraint is defined to declare that all pods created in a specified namespace must be labeled with gatekeeper-test-label. For more information about how to use gatekeeper, see Use gatekeeper.

1. Run the following commands in sequence to create a test-gatekeeper namespace and add the name=test-gatekeeper label to the namespace:

   ```
   kubectl create ns test-gatekeeper
   kubectl label ns test-gatekeeper name=test-gatekeeper
   ```

2. Run the following command to create a Constraint template that can be used to define Constraints on pod labels:

```
kubectl apply -f - <<EOF
apiVersion: templates.gatekeeper.sh/v1beta1
kind: ConstraintTemplate
metadata:
  name: k8srequiredlabels
spec:
  crd:
    spec:
      names:
        kind: K8sRequiredLabels
      validation:
        openAPIV3Schema:
          properties:
            labels:
              type: array
              items:
                type: string
  targets:
    - target: admission.k8s.gatekeeper.sh
      rego: |
        package k8srequiredlabels
        violation[{"msg": msg, "details": {"missing_labels": missing}}] {
          provided := {label | input.review.object.metadata.labels[label]}
          required := {label | label := input.parameters.labels[_]}
          missing := required - provided
          count(missing) > 0
          msg := sprintf("you must provide labels: %v", [missing])
        }
EOF
```

It requires about 10 seconds to initialize the Constraint template.

3. Run the following command to create a Constraint from the preceding Constraint template. This Constraint declares that all pods created in a namespace that is added with the name=test-gatekeeper label must be labeled with gatekeeper-test-label.

```
kubectl apply -f - <<EOF
apiVersion: constraints.gatekeeper.sh/v1beta1
kind: K8sRequiredLabels
metadata:
  name: pod-must-have-gatekeeper-test-label
spec:
  match:
    kinds:
      - apiGroups: [""]
        kinds: ["Pod"]
    namespaceSelector:
      matchExpressions:
      - key: name
        operator: In
        values: ["test-gatekeeper"]
  parameters:
    labels: ["gatekeeper-test-label"]
EOF
```

It takes about 10 seconds to initialize the Constraint.

4. Check whether the namespace is constrained.
   - Run the following command to create a pod that is not labeled with `gatekeeper-test-label` in the test-gatekeeper namespace. The test-gatekeeper namespace is added with the `name=test-gatekeeper` label.

     ```
     kubectl -n test-gatekeeper run test-deny --image=nginx --restart=Never
     ```

     Expected output:

     ```
     Error from server ([denied by pod-must-have-gatekeeper-test-label] you must provide labels: {"gatekeeper-test-label"}): admission webhook "validation.gatekeeper.sh" denied the request: [denied by pod-must-have-gatekeeper-test-label] you must provide labels: {"gatekeeper-test-label"}
     ```

     The test-gatekeeper namespace is added with the `name=test-gatekeeper` label. The pod is to be created without the `gatekeeper-test-label` label. Therefore, the creation fails.
   - Run the following command to create a pod that is labeled with `gatekeeper-test-label` in the test-gatekeeper namespace. The test-gatekeeper namespace is added with the `name=test-gatekeeper` label.

     ```
     kubectl -n test-gatekeeper run test-pass -l gatekeeper-test-label=pass --image=nginx --restart=Never
     ```

     Expected output:

     ```
     pod/test-pass created
     ```

     The test-gatekeeper namespace is added with the `name=test-gatekeeper` label. The pod is to be created with the `gatekeeper-test-label` label. Therefore, the creation succeeds.
   - Run the following command to create a pod that is not labeled with `name=test-gatekeeper` in a namespace that is not subject to the Constraint:

     ```
     kubectl -n default run test-deny --image=nginx --restart=Never
     ```

     Expected output:

     ```
     pod/test-deny created
     ```

     The namespace is not subject to the Constraint. Therefore, the creation succeeds. The created pod is not added with the `name=test-gatekeeper` label.

   The preceding steps show that gatekeeper can be used to constrain pod creations in a specific namespace. In this example, the pod to be created in the namespace must be added with the gatekeeper-test-label label.

Release notes

January 2022

Version	Image address	Release date	Description	Impact
v3.7.0.82-gafe4391b-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/gatekeeper:v3.7.0.82-gafe4391b-aliyun	2022-01-14	This image version is in canary release. • OPA Gatekeeper is upgraded to V3.7.0. The gatekeeper component is dependent on OPA Gatekeeper. For more information about OPA Gatekeeper V3.7.0, see Releases V3.7.0. • Support for ARM 64 architecture.	If exceptions occur during the component upgrade, changes to cluster resources may fail. We recommend that you perform the upgrade during off-peak hours.

September 2021

Version	Image address	Release date	Description	Impact
v3.6.0.62-g156146d-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/gatekeeper:v3.6.0.62-g156146d-aliyun	2021-09-20	• gatekeeper V3.6.0.62 and later versions support only ACK clusters of Kubernetes V1.16.9 and later. • OPA Gatekeeper is upgraded to V3.6.0. The gatekeeper component is dependent on OPA Gatekeeper. For more information about OPA Gatekeeper V3.6.0, see Releases V3.6.0.	If exceptions occur during the component upgrade, changes to cluster resources may fail. We recommend that you perform the upgrade during off-peak hours.

March 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v3.3.0.24-8e68abc-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/gatekeeper:v3.3.0.24-8e68abc-aliyun	2021-03-16	• gatekeeper can be installed in registered Kubernetes clusters. • OPA Gatekeeper is upgraded to V3.3.0. The gatekeeper component is dependent on OPA Gatekeeper.	If exceptions occur during the component upgrade, changes to cluster resources may fail. We recommend that you perform the upgrade during off-peak hours.

August 2020

Version	Image address	Release date	Description	Impact
v3.1.0.11-24bab09-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/gatekeeper:v3.1.0.11-24bab09-aliyun	2020-08-20	OPA Gatekeeper is upgraded to V3.1.0-beta.12. The gatekeeper component is dependent on OPA Gatekeeper. **Note** OPA Gatekeeper is an open source project based on which gatekeeper is developed.	If exceptions occur during the component upgrade, changes to cluster resources may fail. We recommend that you perform the upgrade during off-peak hours.

5.6.5. ack-kubernetes-webhook-injector

In some scenarios where fine-grained permission control is required, you may need to dynamically add the IP addresses of pods to specific whitelists of Alibaba Cloud services. You may also need to remove these IP addresses from specific whitelists of Alibaba Cloud services. You can use ack-kubernetes-webhook-injector to perform these operations. This requires you to add annotations to pod configurations. This topic describes the usage notes of ack-kubernetes-webhook-injector and lists the latest changes to ack-kubernetes-webhook-injector.

Introduction

ack-kubernetes-webhook-injector is a Kubernetes component that can be used to add pod IP addresses to or remove pod IP addresses from the whitelists of a variety of Alibaba Cloud services. This frees you from manual operations to do this. The following figure shows the architecture of ack-kubernetes-webhook-injector.

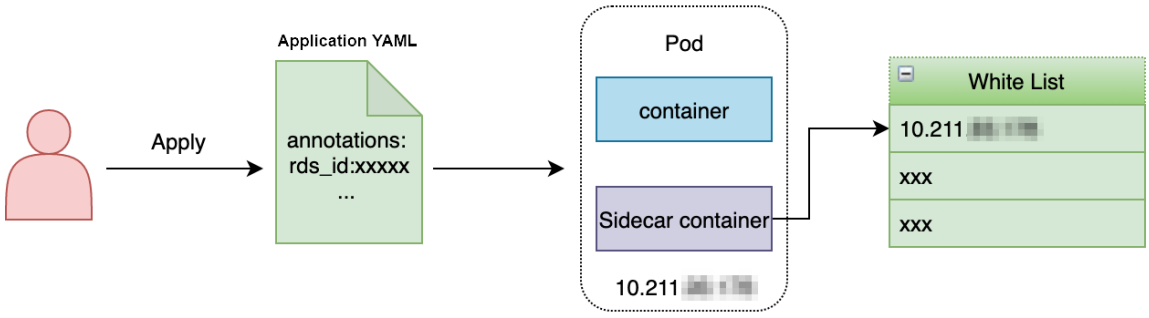

Usage notes

For more information about how to use ack-kubernetes-webhook-injector, see Dynamically add the IP addresses of pods to the whitelists of Alibaba Cloud services.

Release notes

April 2021

Version	Image address	Release date	Description	Impact
v0.0.3-d63ac7e	registry-vpc.cn-hangzhou.aliyuncs.com/acs/k8s-webhook-injector:v0.0.3-d63ac7e	2021-04-12	Access control of Server Load Balancer (SLB) instances is supported.	No impact on workloads

5.6.6. policy-template-controller

policy-template-controller is a key component that is used to manage pod security policies. This topic describes the features, usage notes, and release notes for policy-template-controller.

Introduction

The policy-template-controller component is a Kubernetes controller that is used to manage pod security policies based on clusters and policy instances that are created from policy templates. For more information, see Configure and enforce ACK pod security policies.

Usage notes

For more information about the usage notes for policy-template-controller, see Configure and enforce ACK pod security policies.

Release notes

February 2022

Version	Image address	Release date	Description	Impact
v0.2.0.0-g91ade1a-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/policy-template-controller:v0.2.0.0-g91ade1a-aliyun	2022-02-15	This version is in canary release. • The repetitive synchronization issue is fixed. This reduces the number of requests that are sent to the API server. • The ARM64 architecture is supported.	No impact on workloads

November 2021

Version	Image address	Release date	Description	Impact
v0.1.1.22-gc87e2aa-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/policy-template-controller:v0.1.1.22-gc87e2aa-aliyun	2021-11-16	The policy governance-related features that are described in Configure and enforce ACK pod security policies are supported.	No impact on workloads

5.7. Edge

5.7.1. edge-hub

The edge-hub component is a sidecar that runs as a static pod on each node. This topic introduces the edge-hub component and describes the usage notes and release notes for edge-hub.

Introduction

The edge-hub component serves as a proxy between the components on the edge and the Kubernetes API server on the cloud. The components on the edge include kubelet, kube-proxy, Flannel, and CoreDNS. You can deploy edge-hub on the edge or in the cloud. edge-hub that is deployed on the edge caches data received from the cloud. edge-hub that is deployed on the cloud is named cloud-hub. edge-hub provides the following features:

- Resolves the issue that application pods on the edge cannot use InClusterConfig to access the

Kubernetes API server.

- Supports data caching. Application pods can retrieve cached data from edge-hub when they are restarted even if the edge is disconnected from the cloud. This prevents service interruptions.

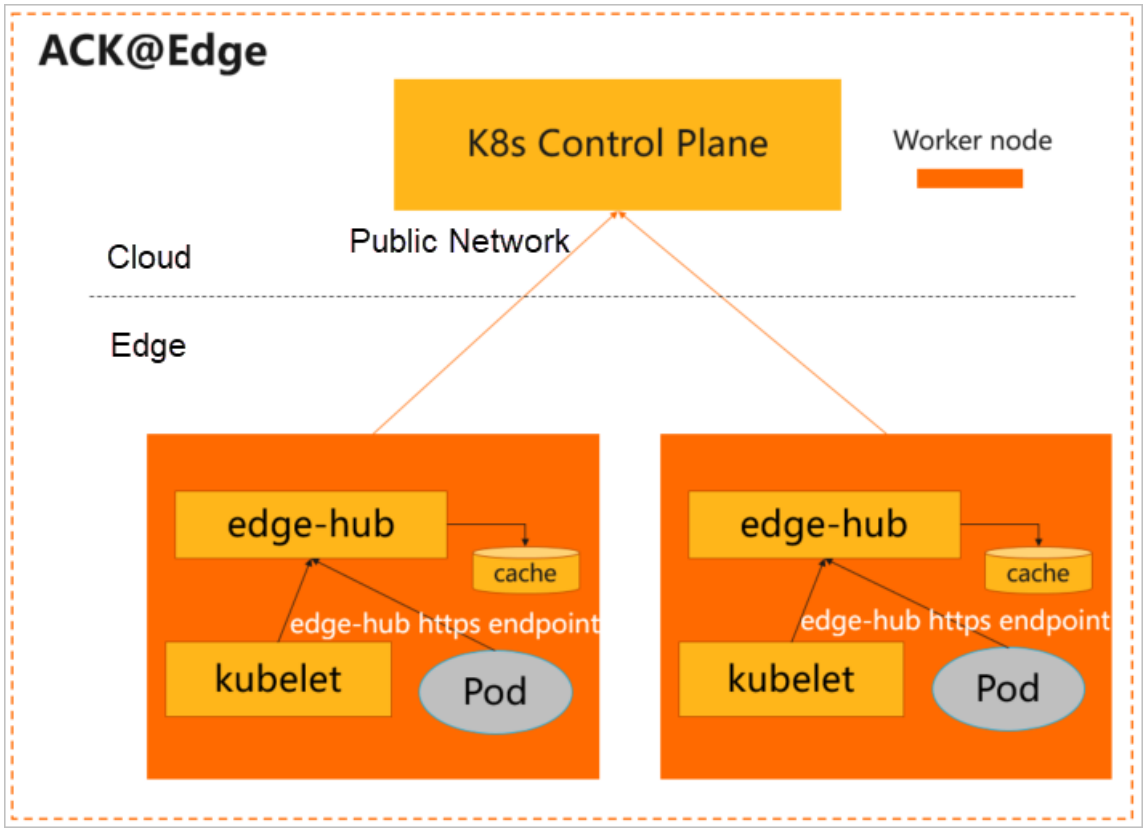

Usage notes

For more information about the features of edge-hub and how to use the features, see Configure node autonomy and Run application pods that use InClusterConfig at the edge without making pod-facing changes.

Release notes

January 2022

Version	Image address	Release date	Description	Impact
			The first edge-hub version v1.20.11-aliyunedge.1 for ACK edge clusters is released. • Health checks are enhanced: The issue that **edge-hub** keeps sending heartbeats after kubelet is stopped is	

Version	Image address	Release date	Description	Impact
v0.10.0	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.10.0	2022-01-27	fixed. • Management of node certificates is enhanced: If the **edge-hub** certificate is not deleted after a node is disconnected from a cluster, the **edge-hub** certificate is automatically updated when the node is connected to another cluster. • Traffic statistics on edge nodes are optimized: You can view traffic statistics about requests on edge nodes by accessing the endpoint of **edge-hub**: http://127.0.0.1:10267/metrics. • Stability is improved: The issue that a data race occurs when edge nodes process a large number of concurrent requests is fixed.	No impact on workloads

November 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.9.5	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.9.5	2021-11-15	The following issue is fixed: Edge requests that do not contain the Accept header are rejected.	No impact on workloads
v0.9.4	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.9.4	2021-11-01	The following issue is fixed: The key usages of HTTPS server certificates are different in different programming languages.	No impact on workloads

October 2021

Version	Image address	Release date	Description	Impact
v0.9.3	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.9.3	2021-10-27	• The following issue is fixed: Large-sized data in data filters is truncated. • Requests from LoadBalancer Services are not filtered.	No impact on workloads
v0.9.2	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.9.2	2021-10-12	• HTTPS servers are supported and port 10268 is added. • The data filtering framework is supported on the edge.	No impact on existing workloads. If you want to use the new features for existing workloads, recreate the pods.

September 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v0.9.1	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.9.1	2021-09-23	Edge nodes that run Windows are supported.	No impact on workloads

July 2021

Version	Image address	Release date	Description	Impact
v0.9.0	registry.cn-hangzhou.aliyuncs.com/acs/edge-hub:v0.9.0	2021-07-12	The first edge-hub version v1.18.8-aliyunedge.1 for ACK edge clusters is released.	No impact on workloads

5.7.2. edge-tunnel

You can use the edge-tunnel component to access edge nodes from the cloud. After you create a Container Service for Kubernetes (ACK) edge cluster, the **edge-tunnel-server** and **edge-tunnel-agent** components are automatically deployed in the cluster to establish tunnels between the cloud and edge nodes. This topic introduces the edge-tunnel component and describes the usage notes and release notes for edge-tunnel.

Introduction

edge-tunnel can establish reverse tunnels, which are commonly used to enable communication between different networks. edge-tunnel is deployed in the client-server architecture. **edge-tunnel-server** is deployed on the cloud and runs as the server. **edge-tunnel-agent** is deployed on edge nodes and runs as the client. edge-tunnel provides the following features:

- edge-tunnel establishes encrypted tunnels over the Internet. The system creates a Server Load Balancer (SLB) instance for the Service that is created by **edge-tunnel-server**. **edge-tunnel-agent** on each node establishes an encrypted tunnel to **edge-tunnel-server** through the SLB instance.
- When components in the cloud, such as kube-apiserver and metrics-server, send requests to port 10250 and port 10255 on edge nodes, edge-tunnel automatically forwards the requests to **edge-tunnel-server**. You do not need to modify the components in the cloud.

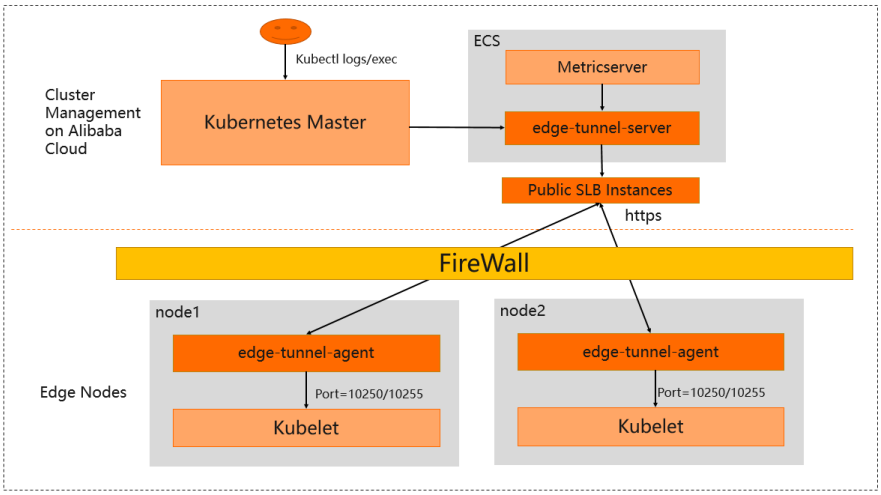

Usage notes

For more information about how to use edge-tunnel, see Cloud-edge tunneling.

Release notes

January 2022

Version	Image address	Release date	Description	Impact
			The first edge-tunnel version for ACK edge clusters of 1.20.11-aliyunedge.1. • Request forwarding is improved: ◦ Requests that are destined for `{nodeName:Port}` can be forwarded from the cloud to edge nodes.	

Version	Image address	Release date	Description	Impact
v0.10.0	**edge-tunnel-server**: registry.cn-hangzhou.aliyuncs.com/acs/edge-tunnel-server:v0.10.0 **edge-tunnel-agent**: registry.cn-hangzhou.aliyuncs.com/acs/edge-tunnel-agent:v0.10.0	2022-01-27	○ Requests that are destined for the `localhost endpoints` on edge nodes can be forwarded from the cloud to edge nodes. This requires you to configure the `localhost-proxy-ports` field in the **edge-tunnel-server-cfg** ConfigMap. • The configuration for access to ports other than 10250 and 10255 is optimized:	No impact on workloads

Version	Image address	Release date	Description	Impact
			○ To configure access to ports other than 10250 and 10255 on an edge node, configure the `http-proxy-ports` field in the **edge-tunnel-server-cfg** ConfigMap if the edge node uses HTTP **endpoints** or configure the `https-proxy-ports` field in the edge-tunnel-server-cfg ConfigMap if the edge node uses HTTPS endpoints. The `dnat-ports-pair` field is retained. However, we recommend that you do not use dnat-ports-pair. • Certificate management is improved for **edge-tunnel-server**. When the IP address of the **edge-tunnel-server-svc** Service is changed, the **tls server** certificate of **edge-tunnel-server** is	

Version	Image address	Release date	Description	Impact
			automatically updated. For example, the certificate is automatically updated when the edge-tunnel-server-svc Service is associated with a new SLB instance.	

5.8. Other SQL statements

5.8.1. ack-arena

The ack-arena component is used to simplify the installation of open source Arena. You can install ack-arena in the Container Service for Kubernetes (ACK) console with a few clicks. This topic introduces ack-arena, describes the usage notes of ack-arena, and lists the latest changes to ack-arena.

Introduction

Arena is a lightweight client that is used to manage Kubernetes-based machine learning tasks. Arena allows you to streamline data preparation, model development, model training, and model prediction throughout a complete lifecycle of machine learning. This improves the work efficiency of data scientists. Arena is also deeply integrated with the basic services of Alibaba Cloud. It supports GPU sharing and Cloud Paralleled File System (CPFS). Arena can run in deep learning frameworks optimized by Alibaba Cloud. This maximizes the performance and utilization of heterogeneous computing resources provided by Alibaba Cloud.

ack-arena is used to simplify the installation of open source Arena. You can install ack-arena in the ACK console with a few clicks. This way, you can deploy Arena in your ACK clusters in an efficient manner.

Usage notes

For more information about how to use ack-arena, see Install Arena.

Release notes

June 2021

Version	Image address	Release date	Description	Impact
0.8.5-264b96a-aliyun	• registry.cn-beijing.aliyuncs.com/acs/arena-binary-installer:0.8.5-264b96a-aliyun • registry.cn-beijing.aliyuncs.com/acs/arena-deploy-manager:0.8.5-264b96a-aliyun	2021-06-08	• The issue that the total number of GPUs is 0 in the output of the arena top node command is fixed. • The issue that role-based access control (RBAC) permissions are not granted to manage CronJobs is fixed. • Inference tasks of the NVIDIA Triton type are supported. • The arena-uninstall command is supported to uninstall Arena.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
0.8.0-ba37c8a-aliyun	• registry.cn-beijing.aliyuncs.com/acs/arena-binary-installer:0.8.0-ba37c8a-aliyun • registry.cn-beijing.aliyuncs.com/acs/arena-deploy-manager:0.8.0-ba37c8a-aliyun	2021-04-06	• The issue that Spark jobs cannot be submitted is fixed. • The issue that the LogViewer URL cannot be obtained when no chief pods are provisioned is fixed. • Arena SDK for Python and Arena SDK for Java are supported. • Inference tasks of the Seldon type can be submitted. • kubeconfig scripts can be generated for multiple tenants to use Arena. • The startup sequence of roles can be customized for TensorFlow training jobs.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
0.7.1-3559f56-aliyun	• registry.cn-beijing.aliyuncs.com/acs/arena-binary-installer:0.7.1-3559f56-aliyun • registry.cn-beijing.aliyuncs.com/acs/arena-deploy-manager:0.7.1-3559f56-aliyun	2021-01-27	The issue that et-operator is not installed in the arena-system namespace is fixed.	No impact on workloads

Version	Image address	Release date	Description	Impact
0.7.0-c6f5800-aliyun	• registry.cn-beijing.aliyuncs.com/acs/arena-binary-installer:0.7.0-c6f5800-aliyun • registry.cn-beijing.aliyuncs.com/acs/arena-deploy-manager:0.7.0-c6f5800-aliyun	2021-01-25	• The issue that multiple jobs are deleted by running the arena serve delete command is fixed. • The output formats of the arena list, arena get, arena serve list, and arena serve get commands are modified. • Arena SDK for Go is supported. • Application Real-Time Monitoring Service (ARMS) Prometheus is supported. • The `-g` option is supported by the arena get command to display GPUs. • The `-c` option is supported by the arena logs command to specify a container.	No impact on workloads

5.8.2. ack-kubernetes-cronhpa-controller

The ack-kubernetes-cronhpa-controller component is used to enable resource scaling by schedule. This topic introduces ack-kubernetes-cronhpa-controller and describes the usage notes and release notes of the component.

Introduction

The ack-kubernetes-cronhpa-controller component is a Kubernetes Horizontal Pod Autoscaler (HPA) controller that you can use to scale the number of pods by schedule. Cron Horizontal Pod Autoscaler (CronHPA) automatically scales the number of pods based on a schedule. You can use CronHPA with Kubernetes objects whose subresources can be scaled. These objects include Deployments and StatefulSets. The subresources must be open source projects on GitHub.

Usage notes

For more information about the usage notes of ack-kubernetes-cronhpa-controller, see CronHPA.

Release notes

May 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.4.0-a2f4954d-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kubernetes-cronhpa-controller:v1.4.0-a2f4954d-aliyun	2021-05-27	The authorization issue related to elastic workloads is fixed.	No impact on workloads

December 2020

Version	Image address	Release date	Description	Impact
v1.4.0-fc4f6060-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/kubernetes-cronhpa-controller:v1.4.0-fc4f6060-aliyun	2020-12-18	• Basic features of CronHPA are supported. • Compatibility with HPA is supported.	No impact on workloads

5.8.3. ack-virtual-node

The ack-virtual-node component enables seamless integration between Kubernetes and Elastic Container Instance. This topic introduces ack-virtual-node, and describes the usage notes and releases notes for ack-virtual-node.

Introduction

The ack-virtual-node component is developed based on the open source project Virtual Kubelet and is extended to support Aliyun Provider. Many improvements are made to the ack-virtual-node component to enable seamless integration between Kubernetes and Elastic Container Instance. This way, Kubernetes clusters are empowered with high elasticity and are no longer limited by the computing capacity of cluster nodes. You can create Elastic Container Instance-based pods as needed to save the hassle of planning the cluster capacity.

The ack-virtual-node component can be used to run pods as serverless container groups on elastic container instances. The pay-as-you-go billing method is supported. This ensures optimal auto scaling capability and avoids unnecessary resource usage. You can use ack-virtual-node to improve the efficiency of application deployment and auto scaling, and minimize the computing cost in multiple scenarios. These scenarios include Job tasks, continuous integration and continuous delivery (CI/CD), Apache Spark big data computing, and online application auto scaling.

For more information about Elastic Container Instance-based pods, see Elastic Container Instance overview.

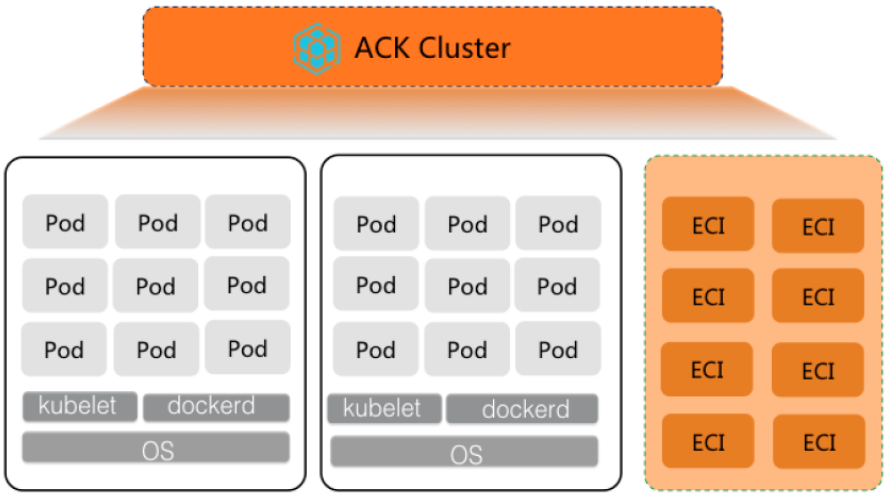

Usage notes

For more information about how to deploy ack-virtual-node from App Catalog, see Deploy the virtual node controller and use it to create Elastic Container Instance-based pods.

Release notes

April 2022

Version	Image address	Release date	Description	Impact
v2.5.3	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.5.3	2022-04-28	• The issue that Metrics data occasionally fail to be obtained is fixed. • The issue that pods occasionally fail to connect to the API server is fixed.	No impact on workloads

Version	Image address	Release date	Description	Impact
v2.5.2	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.5.2	2022-04-20	• The issue that Alibaba Cloud DNS PrivateZone cannot be enabled is fixed. • The issue that the system does not respond to the `exec` command when the cluster runs Kubernetes 1.16 is fixed. • The issue that Metrics data cannot be obtained is fixed. • The issue that the status of the component may be inconsistent with the events of the component is fixed.	No impact on workloads

March 2022

Version	Image address	Release date	Description	Impact
			• The retention period of image caches can be configured. • The requests that are sent to the API server are optimized. • Pod creation is accelerated when you create a large number of pods.	

Version	Image address	Release date	Description	Impact
v2.5.1	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.5.1	2022-03-31	• The latency that occurs when you query Metrics data is reduced. • The issue that memory spikes occur when you concurrently query Metrics data is fixed. • The issue that duplicate Metrics data is returned is fixed. • The issue that the same data is returned from the `/metrics` and `/metrics/cadvisor` endpoints is fixed. • The issue that the SecretRef parameter of the Container Storage Interface (CSI) plug-in does not take effect is fixed. • The system does not respond or even crashes after you run the `logs` or `exec` command is fixed. • The issue that the component cannot run as normal in a private network is fixed.	No impact on workloads

Version	Image address	Release date	Description	Impact
v2.4.3	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.4.3	2022-03-17	The issue that the system does not respond to the `logs` and `exec` commands when the cluster runs Kubernetes 1.16 is fixed.	No impact on workloads

January 2022

Version	Image address	Release date	Description	Impact
v2.4.2	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.4.2	2022-01-11	• The issue of inconsistent status between a cluster and elastic container instances in the cluster is fixed. • The issue that Prometheus Monitoring cannot obtain data is fixed. • The issue that specific annotations in Elastic Container Instance Profile do not take effect is fixed.	No impact on workloads

December 2021

Version	Image address	Release date	Description	Impact
			• Kubernetes 1.22 is supported. • Dynamic expansion of persistent volume claims (PVCs) is supported. • The WaitForFirstCon	

Version	Image address	Release date	Description	Impact
v2.4.0	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.4.0	2021-12-28	sumer volume binding mode is supported. • The efficiency of calling the Elastic Container Instance API is improved. • The issue that the version number of a virtual node is improperly displayed after the cluster is upgraded is fixed. • The issue of status inconsistency between Job pods and the elastic container instances on which the pods run is fixed. • The issue of status inconsistency between pods and the elastic container instances on which the pods run is fixed. • The issue that the SecretRef parameter in the CephFS and FlexVolume configurations does not take effect is fixed. • The issue that the component occasionally stops running is fixed.	No impact on workloads

November 2021

Version	Image address	Release date	Description	Impact
v2.3.0	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.3.0	2021-11-15	• Service account token volume projection is supported. • IPv4/IPv6 dual stack is supported. • Different types of overlay networks are supported. • The issue that the component stops running in specific cases is fixed. • Pod creation is accelerated. • The pod scheduling issue is fixed. • The issue that the reclaim process of elastic container instances is slow is fixed.	No impact on workloads

October 2021

Version	Image address	Release date	Description	Impact
v2.2.0	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.2.0	2021-10-26	Queries for events and status are accelerated.	No impact on workloads

Version	Image address	Release date	Description	Impact
v2.1.0	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.1.0	2021-10-12	**New features:** • Images can be pulled without a password from instances of Container Registry Enterprise Edition. • Secrets can be referenced for PVCs, CephFS, and FlexVolume. • The HugePages feature is supported by nodes. **Fixed issues:** • The issue of Elastic Container Instance resource leakage in specific cases • The issue that the displayed information about CPU workloads is wrong • The issue that nodes change to the NotReady state due to slow restarts • The issue of pod creation failures due to excessively long URLs • The issue that ClusterRoles are deleted in specific cases **Optimizations:** • Pod creation is accelerated. • Image caching is accelerated.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact
v2.0.0.122-bdb884460-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.122-bdb884460-aliyun	2021-05-24	The issue that nodes are repeatedly created when concurrent requests are processed is fixed.	No impact on workloads

Version	Image address	Release date	Description	Impact
v2.0.0.113-650b21c48-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.113-650b21c48-aliyun	2021-05-13	Rolling update of PrivateZone is supported.	No impact on workloads
v2.0.0.110-49c928287-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.110-49c928287-aliyun	2021-05-06	Node selectors can be used to configure pod scheduling. This means that when a pod is configured with the `type=virtual-kubelet` label, the pod is managed by Virtual Kubelet.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
v2.0.0.102-045a06eb4-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.102-045a06eb4-aliyun	2021-04-22	The issue that information about Virtual Kubelet nodes is improperly displayed is fixed.	No impact on workloads
v2.0.0.101-cd8bcfa04-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.101-cd8bcfa04-aliyun	2021-04-16	The calling of the Alibaba Cloud DNS PrivateZone API is optimized. This avoids repeated calls to BindZoneVpc.	No impact on workloads
v2.0.0.93-ef86bbfa1-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.93-ef86bbfa1-aliyun	2021-04-13	The value of eciLogtailMacheineGroupKey is modified to support the log collection feature of environment variables and CustomResourceDefinitions (CRDs).	No impact on workloads

Version	Image address	Release date	Description	Impact
v2.0.0.90-15deb126e-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.90-15deb126e-aliyun	2021-04-06	The feature of Elastic Container Instance Profile is supported.	No impact on workloads

March 2021

Version	Image address	Release date	Description	Impact
v2.0.0.86-9005a977d-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.86-9005a977d-aliyun	2021-03-17	• The */etc/hosts* file is optimized to support the Domain and DNS annotations. • The time required to check ClusterRoleBinding is reduced. • Leader elections for multiple components are optimized.	No impact on workloads

February 2021

Version	Image address	Release date	Description	Impact
v2.0.0.80-f9a46a994-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.80-f9a46a994-aliyun	2021-02-25	The feature of checking the CreatePod state is optimized: A pod can be created only when CreatePod is in the Pending state.	No impact on workloads
v2.0.0.76-6e9e19bd5-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.76-6e9e19bd5-aliyun	2021-02-22	• PProf debugging is supported. • Switches are added to control leader elections. • ClusterRole is updated for EndPointSlice. • Annotations related to Kubernetes versions can be added to pods. • Different webhooks can be configured in related Kubernetes versions.	No impact on workloads

December 2020

Version	Image address	Release date	Description	Impact
v2.0.0.618-7fd50d738-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v2.0.0.618-7fd50d738-aliyun	2020-12-04	• Compatibility with pods is improved. • Multiple leader election logic of the Virtual Kubelet controller is supported.	Make sure that virtual private clouds (VPCs) and security groups are correctly configured for pods, and the pods can access the API server of your cluster. Otherwise, the pods remain in the waiting state.

September 2020

Version	Image address	Release date	Description	Impact
v1.0.0.10-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/virtual-nodes-eci:v1.0.0.10-aliyun	2020-09-21	Auto scaling of clusters is supported. Pods are run as serverless container groups on Elastic Container Instance.	No impact on workloads

March 2020

Version	Image address	Release date	Description	Impact
v1.0.0.2-aliyun	registry-vpc.$RegionId.aliyuncs.com/acs/virtual-nodes-eci:v1.0.0.2-aliyun	2020-03-12	• The virtual-nodes-eci controller can be deployed on StatefulSets. This allows you to change the number of replicated pods to create multiple virtual nodes and more pods. • The names of virtual nodes use the virtual-node-eci-$n format. • Pods are allowed to access ClusterIP Services. • Preemptible instances are supported. • Disks can be mounted by using the Container Storage Interface (CSI) plug-in.	If the virtual-nodes-eci controller is deployed on a Deployment, delete the pods on the virtual-kubelet node and install the component again.

5.8.4. sgx-device-plugin

This topic describes the features of sgx-device-plugin and lists the latest changes to the component.

Introduction

sgx-device-plugin is a Kubernetes device plug-in developed by Container Service for Kubernetes (ACK) and Ant Financial. sgx-device-plugin facilitates the use of Intel (R) Software Guard Extensions (SGX) in containers. Intel (R) SGX is a set of CPU instructions provided by Intel. Intel (R) SGX increases the security of application code and data. This protects your code and data against disclosure and malicious tampering. For more information, see software-guard-extensions.

Features

sgx-device-plugin provides the following features:

- Intel (R) SGX can be used without the need to enable the privilege mode.
- The Enclave Page Cache (EPC) size can be automatically retrieved.
- Declarative EPC resource allocation is supported.

Dependencies

sgx-device-plugin is dependent on the following components and tools:

- TEE-SDK. TEE-SDK enables the compatibility between and Intel SGX and Intel SGX Platform Software (PSW).
- The Kubernetes version must be V1.10 or later.
- The Go version must be V1.10 or later.

FAQ

- Can I deploy sgx-device-plugin in Kubernetes clusters that are deployed outside Alibaba Cloud?

  Yes, sgx-device-plugin can be deployed in all types of Kubernetes clusters. However, you can run sgx-device-plugin only on SGX-enabled nodes.

- Can I use sgx-device-plugin to control the EPC size for SGX-enabled containers?

  No, you cannot use sgx-device-plugin to control the EPC size for SGX-enabled containers. The EPC size limit specified by the alibabacloud.com/sgx_epc_MiB parameter applies only to kube-scheduler. Intel (R) SGX Driver does not support this parameter.

- Is sgx-device-plugin open source?

  Yes, sgx-device-plugin is open source. For more information, see sgx-device-plugin.

Usage notes

By default, sgx-device-plugin is installed in ACK clusters. You can use sgx-device-plugin without extra configurations.

Release notes

April 2021

Version	Image address	Release date	Description	Impact
v1.1.0-bb1f5f9-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sgx-device-plugin:v1.1.0-bb1f5f9-aliyu	2021-04-30	• Intel (R) SGX can be used without the need to enable the privilege mode. • The EPC size can be automatically retrieved. • Declarative EPC resource allocation is supported.	No impact on workloads

5.8.5. Intel SGX AESM

This topic describes Intel (R) Software Guard Extensions (SGX) Architectural Enclave Service Manager (AESM) and lists the latest changes to Intel (R) SGX AESM.

Introduction

Intel (R) SGX AESM is a system component of Intel SGX. This component provides launch support for SGX Enclave, key provisioning, and remote attestation services. In trusted execution environments (TEEs) provided by Container Service for Kubernetes (ACK), Intel (R) SGX AESM runs on DaemonSets in ACK clusters.

Usage notes

By default, Intel (R) SGX AESM is installed in ACK clusters. You can use Intel (R) SGX AESM without extra configurations.

Release notes

April 2021

Version	Image address	Release date	Description	Impact
2.13.100.4-bionic1-d83e54d-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/aesm:2.13.100.4-bionic1-d83e54d-aliyun	2021-04-30	The Intel (R) SGX AESM component is added. The component provides launch support for SGX Enclave, key provisioning, and remote attestation services.	No impact on workloads

5.8.6. sandboxed-container-controller

This topic describes sandboxed-container-controller and lists the latest changes to sandboxed-container-controller.

Introduction

sandboxed-container-controller is a controller component that is provided by the Sandboxed-Container runtime. The component is used to enhance and extend the basic features of sandboxed containers. The following features are supported:

- Custom kernel parameter settings for pods that run sandboxed containers.
- Automatic calculation and configuration for VMs that run sandboxed containers.
- Direct mounting of disks and Apsara File Storage NAS (NAS) file systems to sandboxed containers.

Usage notes

By default, sandboxed-container-controller is installed in ACK cluster. You can use this component without extra configurations.

Release notes

December 2020

Version	Image address	Release date	Description	Impact
v1.1.1-55d545f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sandboxed-container-controller:v1.1.1-55d545f-aliyun	2020-12-22	The following annotations are forcibly overwritten if they are manually added to pod configurations: *securecontainer.alibabacloud.com/cpus* and *securecontainer.alibabacloud.com/memory*.	No impact on workloads

November 2020

Version	Image address	Release date	Description	Impact
v1.1.0-3b3d499-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sandboxed-container-controller:v1.1.0-3b3d499-aliyun	2020-11-26	Custom kernel parameter settings are supported for pods that run sandboxed containers.	No impact on workloads

Version	Image address	Release date	Description	Impact
v1.0.3-e993d8f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sandboxed-container-controller:v1.0.2-8ac82bf-aliyun	2020-11-12	The PodEraseRuntimeclassRunc admission controller is supported. Docker does not support the RuntimeClass feature. Therefore, when `pod.spec.runtimeClassName` is set to `runc`, this parameter is reset to an empty string.	No impact on workloads

August 2020

Version	Image address	Release date	Description	Impact
v1.0.1-8484958-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sandboxed-container-controller:v1.0.1-8484958-aliyun	2020-08-26	Sandboxed-Container 2.0 is supported. The PodQuota admission controller is supported for sandboxed containers. This admission controller can set pod specifications based on the total CPU and memory resources used by sandboxed containers.	No impact on workloads

June 2020

Version	Image address	Release date	Description	Impact
v1.0.0-e408663-aliyun	registry.cn-beijing.aliyuncs.com/acs/sandboxed-container-controller:v1.0.0-e408663-aliyun	2020-06-10	The NAS image address of the init container is changed from a public image address to a private image address.	No impact on workloads

March 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.0-a8b276f-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sandboxed-container-controller:v1.0.0-a8b276f-aliyun	2020-03-26	The feature of directly mounting disks and NAS file systems to sandboxed containers is supported. This provides the same performance as when these volumes are mounted through the host. This feature allows you to prevent performance loss when volumes are mounted over 9pfs.	No impact on workloads

5.8.7. sandboxed-container-helper

This topic describes the features of sandboxed-container-helper and lists the latest changes to sandboxed-container-helper.

Introduction

sandboxed-container-helper performs health checks and O&M operations on sandboxed containers. sandboxed-container-helper provides the following features:

- Provides the Prometheus exporter to collect information about the disk space allocated by Device Mapper. You can deploy ack-arms-prometheus in Container Service for Kubernetes (ACK) clusters to monitor the disk space allocated by Device Mapper and configure alerts. For more information, see Enable ARMS Prometheus.
- Checks and reports unusual events to kube-apiserver, such as storage leaks, container data leaks, and orphaned pods. You can deploy ack-node-problem-detector in ACK clusters to collect and monitor these events. For more information, see Event monitoring.

Usage notes

By default, sandboxed-container-helper is installed in ACK clusters. You can use the component without extra configurations.

Release notes

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.0-7a70086-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/sandboxed-container-helper:v1.0.0-7a70086-aliyun	2020-05-12	New features: • Unusual events, such as container data leaks and orphaned pods, are reported to kube-apiserver. • The monitoring of the disk space allocated by Device Mapper is provided. • Scripts are provided to fix system issues.	No impact on workloads

5.8.8. yurt-app-manager

The yurt-app-manager component provides cell-based management at the edge for edge Kubernetes clusters. This topic describes usage notes and latest changes to the yurt-app-manager component.

Introduction

- In edge computing, edge nodes are classified into groups by zone, region, or other logical attribute such as CPU architecture, Internet service provider (ISP), or cloud service provider.
- Same applications or images may be deployed to different node pools.
- The backend endpoints of Kubernetes-native Services are arbitrarily distributed across nodes. Consequently, when Service requests are distributed to nodes across groups, these requests may fail to reach the nodes or may not be answered promptly.

ACK@Edge provides a solution to solve these issues, as shown in the following figure.

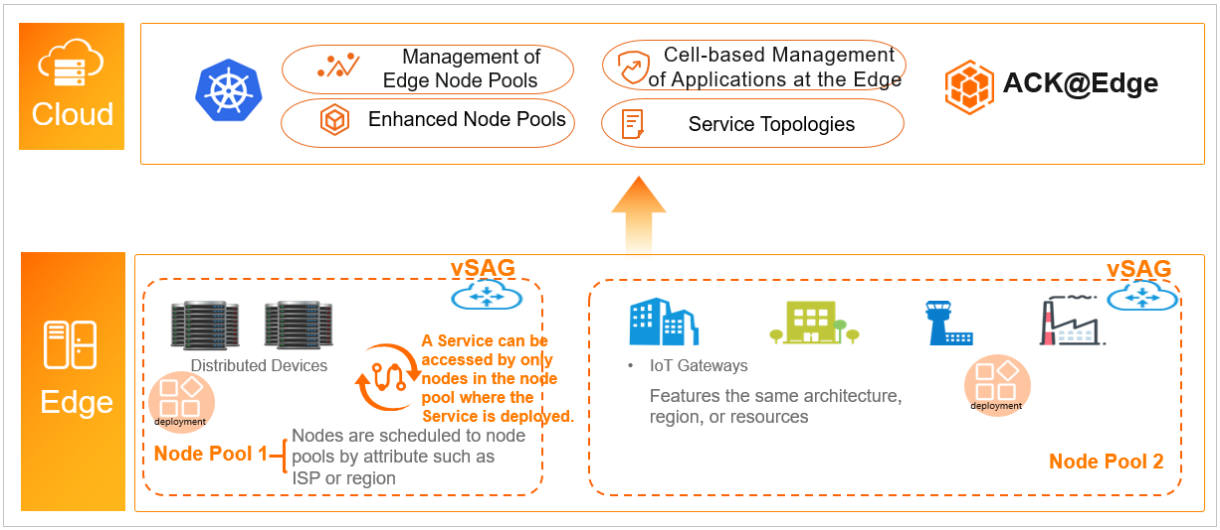

- Node cell: You can create node pools to manage and maintain hosts in different regions.
- Application cell: You can deploy workloads to different node pools. This way, you can manage the number of pods and the image version of containers by node pool.
- Traffic cell: You can configure a Service topology to limit access to Service endpoints. For example, you can expose an application on an edge node to only the current node or other nodes in the same edge node pool.

The yurt-app-manager component supports edge node management (NodePool) and workload management (UnitedDeployment).

Usage notes

The NodePool controller is used to create node pools. The UnitedDeployment controller is used to centrally manage multiple workloads. For more information, see Overview of edge node pools and Use the UnitedDeployment controller to deploy applications.

> **Note** By default, the yurt-app-manager component is installed when an edge Kubernetes cluster is created. For more information about how to upgrade the components of an edge Kubernetes cluster, see Upgrade components in an edge Kubernetes cluster.

Release notes

January 2021

Version	Image address	Release date	Description	Impact
v0.9.11-c2c8cce-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/yurt-app-manager:v0.9.11-c2c8cce-aliyun	2021-01-14	The yurt-app-manager component is released.	No impacts on your workloads.

5.8.9. migrate-controller

The migrate-controller component is developed based on open source Velero. migrate-controller is deeply integrated with the snapshot technologies and Hybrid Backup Recovery (HBR) service of Alibaba Cloud. You can use migrate-controller to back up, restore, and migrate data and application templates. This topic introduces migrate-controller and describes the release notes for migrate-controller.

Introduction

For more information about how to install migrate-controller, see Enable cluster backup.

You can use migrate-controller to back up applications and persistent volumes (PVs) in your Kubernetes clusters. A variety of policies are supported to filter the applications that you want to back up. For example, you can back up applications based on namespaces, labels, and resource types. migrate-controller allows you to create disk snapshots and use HBR to back up files. migrate-controller also allows you to reduce the recovery time objective (RTO) of data restoration by using snapshots.

Usage notes

You can use migrate-controller to back up, restore, and migrate applications. For more information, see Back up and restore applications and Migrate applications across clusters.

Release notes

November 2021

Version	Image address	Release date	Description	Impact
v1.5.8-14de9e1-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.5.8-14de9e1-aliyun	2021-11-28	• migrate-controller can be used in multi-cluster scenarios. • Snapshots of persistent volume claims (PVCs) that are used to provision disk volumes can be created. • The concurrency of snapshot creation requests is improved. • Cross-region data restoration that uses snapshots is supported.	Ongoing backup tasks may be interrupted.

October 2021

Version	Image address	Release date	Description	Impact
v1.5.6-94176c2-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.5.6-94176c2-aliyun	2021-10-12	The methods to back up file systems are optimized. The HBR client can be dynamically started.	Ongoing backup tasks may be interrupted.
v1.5.7-561a15f-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.5.7-561a15f-aliyun	2021-10-15	The logic of internal access to HBR is optimized.	Ongoing backup tasks may be interrupted.

September 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.5.5-d5e168a-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.5.5-d5e168a-aliyun	2021-09-20	• The error message of backup failure is updated. • The issue that namespaces cannot be created when the system has ongoing data restoration tasks is fixed. • The DaemonSet of the backup client of HBR is no longer automatically deployed. If you want to use HBR, you must manually deploy the DaemonSet. • Backslashes (\) are supported to specify the names of subdirectories in Object Storage Service (OSS) buckets. • An error is reported when you back up a file system without having HBR enabled.	Ongoing backup tasks may be interrupted.

Version	Image address	Release date	Description	Impact
v1.5.2-d728be2-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.5.2-d728be2-aliyun	2021-09-10	• Compatibility with earlier versions is supported. • Snapshot creation methods are upgraded. • HBR is used to back up files. • Disk snapshots and snapshots of Apsara File Storage NAS (NAS) file systems are supported to reduce the RTO of data restoration. • The permission control policies of ACK managed clusters are upgraded. • Restic is no longer supported to back up data.	Ongoing backup tasks may be interrupted.

July 2021

Version	Image address	Release date	Description	Impact
v1.4.1-f745291-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.4.1-f745291-aliyun	2021-07-13	The snapshot feature can be enabled based on the Kubernetes version and the Container Storage Interface (CSI) plug-in version.	Ongoing backup tasks may be interrupted.

March 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.3.1-3026c10-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.3.1-3026c10-aliyun	2021-03-31	The logic of component uninstallation is upgraded. Errors about components are fixed.	Ongoing backup tasks may be interrupted.
v1.3.1-665c3ea-aliyun	registry.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.3.1-665c3ea-aliyun	2021-03-31	Volume snapshots are supported in registered clusters.	Ongoing backup tasks may be interrupted.
v1.3.0-9274575-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.0.1.2-h-e616322-aliyun	2021-03-29	Volume snapshots are supported to back up volumes.	Ongoing backup tasks may be interrupted. To upgrade the component, to contact the technical support team.

December 2020

Version	Image address	Release date	Description	Impact
v1.0.1.2-h-e616322-aliyun	registry.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.0.1.2-h-e616322-aliyun	2020-12-24	Velero is upgraded to the open source version 1.5.2 in registered clusters.	Ongoing backup tasks may be interrupted. To upgrade the component, to contact the technical support team.
v1.0.1.2-a-e616322-aliyun	registry-vpc.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.0.1.2-h-e616322-aliyun	2020-12-30	Velero is upgraded to the open source version 1.5.2 in ACK managed clusters and ACK dedicated clusters.	Ongoing backup tasks may be interrupted. To upgrade the component, to contact the technical support team.

September 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.0.1.1-30d319f-aliyun	registry.{{.Region}}.aliyuncs.com/acs/velero-installer:v1.0.1.1-30d319f-aliyun	2020-09-21	Application migration, backup, and restoration are supported in registered clusters and ACK managed clusters.	First version.

5.8.10. aliyun-acr-acceleration-suite

This topic describes the features of the aliyun-acr-acceleration-suite component and lists the latest changes to aliyun-acr-acceleration-suite.

Introduction

The aliyun-acr-acceleration-suite component is a client plug-in that enables on-demand image loading. This accelerates image loading. You can deploy this component on a worker node as a DaemonSet. Container Service for Kubernetes (ACK) collaborates with Container Registry to enable accelerated images. After you deploy aliyun-acr-acceleration-suite in a cluster, the system automatically converts a source image into an accelerated image. When an image pull request is received, the system decompresses data based on your needs without interrupting services. In this case, the system does not have to download or decompress full data. This speeds up application deployment and enables on-demand image loading.

> **Note** Only managed and dedicated Kubernetes clusters of Kubernetes 1.16.9 or later support accelerated images. When you create the cluster, set the container runtime to Docker and select one of the following operating systems: Aliyun Linux 2.1903, CentOS 7.6, CentOS 7.7, CentOS 7.8, and CentOS 7.9.

Usage notes

For more information about aliyun-acr-acceleration-suite, see Load resources of a container image on demand.

Release notes

March 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
0.2.0	registry-vpc.cn-hangzhou.aliyuncs.com/acr-toolkit/aliyun-acr-acceleration-suite:v0.2.0.0-b745758c-aliyun	2021-03-15	• HTTP readiness and HTTP liveness probes are supported for mutating admission webhooks. • Labels can be added to workloads to enable on-demand image loading. The key of the label is `k8s.aliyun.com/image-accelerate-mode`. The value of the label is `on-demand`. **Note** For more information about how to configure on-demand image loading, see Load resources of a container image on demand.	No impact on workloads

November 2020

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
0.1.0	registry.cn-hangzhou.aliyuncs.com/acs/aliyun-acr-acceleration-suite:v0.1.0.0-00e2f5e1-aliyun	2020-11-19	Management of image storage plug-ins, configuration of access to image repositories, and automatic injection of accelerated images to pods are supported.	We recommend that you upgrade during off-peak hours.

5.8.11. managed-kube-proxy-windows

This topic introduces managed-kube-proxy-windows and describes the usage notes and release notes for managed-kube-proxy-windows.

Introduction

managed-kube-proxy-windows is a containerized kube-proxy used by managed Kubernetes clusters. managed-kube-proxy-windows manages the endpoints of Services on Windows nodes, including internal endpoints and external endpoints.

Usage notes

When the system creates a managed Kubernetes cluster, managed-kube-proxy-windows is automatically deployed as a DaemonSet in the cluster. managed-kube-proxy-windows is reliant on the privileged proxy processes on Windows nodes. managed-kube-proxy-windows can be deployed only on Windows nodes that are added with the `windows.alibabacloud.com/deployment-topology=2.0` label. The following section describes the rules for adding the `windows.alibabacloud.com/deployment-topology=2.0` label to Windows nodes:

- If the Kubernetes version of your cluster is 1.20.4 or later, the `windows.alibabacloud.com/deployment-topology=2.0` label is automatically added to Windows nodes when you create the Windows nodes. If the label is not added to the Windows nodes, you can manually add the `windows.alibabacloud.com/deployment-topology=2.0` label to the Windows nodes.
- If the Kubernetes version of your cluster is earlier than 1.20.4, the `windows.alibabacloud.com/deployment-topology=2.0` label is not added to Windows nodes when you create the Window nodes. You cannot manually add the `windows.alibabacloud.com/deployment-topology=2.0` label to the Windows nodes.

If your cluster does not contain Windows nodes or the Windows nodes in your cluster do not have the `windows.alibabacloud.com/deployment-topology=2.0` label, no replicated pods are provisioned for managed-kube-proxy-windows.

If no replicated pods are provisioned for managed-kube-proxy-windows on a Windows node with the `windows.alibabacloud.com/deployment-topology=2.0` label, you must manually install managed-kube-proxy-windows.

1.

2.
3.
4. In the left-side navigation pane, choose **Operations > Add-ons**.
5. On the **Add-ons** page, find managed-kube-proxy-windows and click **Install**.
6. In the message that appears, click **OK**.

Release notes

August 2021

Version	Image address	Release date	Description	Impact
v1.20.4-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.20.4-aliyun.1	2021-08-22	• The base images are updated. The original version numbers are overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads.
v1.18.8-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.18.8-aliyun.1	2021-08-20	• The base images are updated. The original version numbers are overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads.
v1.16.9-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.16.9-aliyun.1	2021-08-20	• The base images are updated. The original version numbers are overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads.

July 2021

Version	Image address	Release date	Description	Impact
v1.20.4-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.20.4-aliyun.1	2021-07-05	• The base images are updated. The original version numbers are overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1999), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1052).	No impact on workloads.
v1.18.8-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.18.8-aliyun.1	2021-07-05	• The base images are updated. The original version numbers are overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1999), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1052).	No impact on workloads.
v1.16.9-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.16.9-aliyun.1	2021-07-05	• The base images are updated. The original version numbers are overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1999), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1052).	No impact on workloads.

June 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.20.4-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.20.4-aliyun.1	2021-06-17	• The base images are updated. The original version numbers are overwritten. • The following Windows systems are supported: Windows Server version 1809 (10.0.17763.1935), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.985).	No impact on workloads.
v1.18.8-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.18.8-aliyun.1	2021-06-17	• The base images are updated. The original version numbers are overwritten. • The following Windows systems are supported: Windows Server version 1809 (10.0.17763.1935), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.985).	No impact on workloads.
v1.16.9-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.16.9-aliyun.1	2021-06-17	• The base images are updated. The original version numbers are overwritten. • The following Windows systems are supported: Windows Server version 1809 (10.0.17763.1935), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.985).	No impact on workloads.

April 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.20.4-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.20.4-aliyun.1	2021-04-23	• Rancher Wins is automatically installed on a Windows node with the `windows.alibabacloud.com/deployment-topology=2.0` label. You can deploy kube-proxy 1.20.4-aliyun.1 on a Windows node by using Rancher Wins. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763), Windows Server version 1909 (10.0.18363.1440), and Windows Server version 2004 (10.0.19041.867).	No impact on workloads.
v1.18.8-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.18.8-aliyun.1	2021-04-23	• Rancher Wins is automatically installed on a Windows node with the `windows.alibabacloud.com/deployment-topology=2.0` label. You can deploy kube-proxy 1.18.8-aliyun.1 on a Windows node by using Rancher Wins. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763), Windows Server version 1909 (10.0.18363.1440), and Windows Server version 2004 (10.0.19041.867).	No impact on workloads.

Version	Image address	Release date	Description	Impact
v1.16.9-aliyun.1	registry-vpc.${region}.aliyuncs.com/acs/kube-proxy-windows:v1.16.9-aliyun.1	2021-04-23	• Rancher Wins is automatically installed on a Windows node with the `windows.alibabacloud.com/deployment-topology=2.0` label. You can deploy kube-proxy v1.16.9-aliyun.1 on a Windows node by using Rancher Wins. The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763), Windows Server version 1909 (10.0.18363.1440), and Windows Server 2004 (10.0.19041.867). • The following Windows systems are supported: Windows Server Release 1809 (10.0.17763), Windows Server Release 1909 (10.0.18363.1440), and Windows Server Release 2004 (10.0.19041.867).	No impact on workloads.

5.8.12. resource-controller

This topic lists the latest changes to resource-controller.

Introduction

resource-controller is the core component that is used to dynamically schedule pods. To enable topology-aware CPU scheduling for a professional Kubernetes cluster, you must install resource-controller in the cluster.

Usage notes

For more information about the usage notes of resource-controller, see the following topics:

- Manage system components
- Topology-aware CPU scheduling

Release notes

June 2021

Version	Image address	Release date	Description	Impact

Version	Image address	Release date	Description	Impact
v1.2.2-0ac97de0-aliyun	• registry.cn-hangzhou.aliyuncs.com/acs/node-resource-agent:v21.3.9.0-adecd8a-aliyun • registry.cn-hangzhou.aliyuncs.com/acs/resource-controller:v1.2.2-0ac97de0-aliyun	2021-06-21	Memory Bandwidth Allocation (MBA) control is supported. Issues related to the change memory feature and containerd are fixed. client-go is upgraded.	No impact on workloads

May 2021

Version	Image address	Release date	Description	Impact
v1.2.1-d1e280f-aliyun	• registry.cn-hangzhou.aliyuncs.com/acs/node-resource-agent:v21.3.9.0-adecd8a-aliyun • registry.cn-hangzhou.aliyuncs.com/acs/resource-controller:v1.2.1-d1e280f-aliyun	2021-05-21	L3 control is supported. Topology-aware scheduling is compatible with AMD products.	No impact on workloads

April 2021

Version	Image address	Release date	Description	Impact
v1.2.0-ec8a979-aliyun	• registry.cn-hangzhou.aliyuncs.com/acs/node-resource-agent:v21.3.9.0-adecd8a-aliyun • registry.cn-hangzhou.aliyuncs.com/acs/resource-controller:v1.2.0-ec8a979-aliyun	2021-04-20	resources-controller is compatible with containerd. Access to multiple containers and their metadata is supported. The container log can be collected. The resource water marks of tasks can be adjusted in real time.	No impact on workloads

March 2021

Version	Image address	Release date	Description	Impact
v1.1.0-e7388cb	• registry.cn-hangzhou.aliyuncs.com/acs/node-resource-agent:v21.3.9.0-adecd8a-aliyun • registry.cn-hangzhou.aliyuncs.com/acs/resource-controller:v1.1.0-e7388cb	2021-03-09	node-agent is released. node-agent allows you to add the `label` field to the ConfigMaps that store information about non-uniform memory access (NUMA) nodes.	No impact on workloads

January 2021

Version	Image address	Release date	Description	Impact
v1.1.0-e7388cb	• registry.cn-hangzhou.aliyuncs.com/acs/node-resource-agent:121ffbe • registry.cn-hangzhou.aliyuncs.com/acs/resource-controller:v1.1.0-e7388cb	2021-01-26	Distributed deployment of resource-controller is supported. Custom values of the cpuset parameter are supported.	No impact on workloads

5.8.13. directx-device-plugin-windows

This topic introduces directx-device-plugin-windows and describes the usage notes and release notes for directx-device-plugin-windows.

Introduction

directx-device-plugin-windows is a DirectX device plug-in for Container Service for Kubernetes (ACK) clusters. directx-device-plugin-windows enables GPU acceleration based on DirectX for Windows containers that are deployed on GPU-accelerated virtualization instances. For more information about GPU-accelerated virtualization instances, see vGPU-accelerated instance families.

directx-device-plugin-windows supports only ACK clusters that use Kubernetes 1.20.4 or later.

Usage notes

By default, directx-device-plugin-windows is deployed as a DaemonSet. directx-device-plugin-windows is reliant on the privileged proxy processes of Windows nodes. directx-device-plugin-windows can be installed only on Windows nodes that are added with the `windows.alibabacloud.com/deployment-topology=2.0` and the `windows.alibabacloud.com/directx-supported=true` labels.

The following section describes the rules for adding the `windows.alibabacloud.com/deployment-topology=2.0` label to Windows nodes:

- If the Kubernetes version of your cluster is 1.20.4 or later, the `windows.alibabacloud.com/deployment-topology=2.0` label is automatically added to Windows nodes when the Windows nodes are created. If the Windows nodes are not added with the `windows.alibabacloud.com/deployment-topology=2.0` label, you can manually add the label to the Windows nodes.
- If the Kubernetes version of your cluster is earlier than 1.20.4, the `windows.alibabacloud.com/deployment-topology=2.0` label is not added to Windows nodes when the Windows nodes are created. You cannot manually add the `windows.alibabacloud.com/deployment-topology=2.0` label to the Windows nodes.

The following section describes the rules for adding the `windows.alibabacloud.com/directx-supported=true` label to Windows nodes:

When you create a Windows node pool, the instances that you select must be GPU-accelerated instances. This way, you can add the `windows.alibabacloud.com/directx-supported=true` label to the Windows nodes. If the instances that you are select are GPU-accelerated instances when you create the Windows node pool, the `windows.alibabacloud.com/directx-supported=true` label is automatically added to the Windows nodes. If the Windows nodes are not added with the `windows.alibabacloud.com/directx-supported=true` label, you can manually add the label to the Windows nodes.

directx-device-plugin-windows can only be manually installed. For more information, see Enable GPU acceleration for DirectX in Windows containers.

Release notes

August 2021

Version	Image address	Release date	Description	Impact
v1.0.0	registry-vpc.${region}.aliyuncs.com/acs/directx-device-plugin-windows:v1.0.0	2021-08-19	• The base images are updated. The original version number is overwritten. • The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.2114), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1165).	No impact on workloads.

July 2021

Version	Image address	Release date	Description	Impact
v1.0.0	registry-vpc.${region}.aliyuncs.com/acs/directx-device-plugin-windows:v1.0.0	2021-07-05	The following Windows operating systems are supported: Windows Server version 1809 (10.0.17763.1999), Windows Server version 1909 (10.0.18363.1556), and Windows Server version 2004 (10.0.19041.1052).	No impact on workloads.

5.8.14. ack-cluster-agent

ack-cluster-agent is a component that is deployed in an external cluster after you register the cluster to Container Service for Kubernetes (ACK). ack-cluster-agent is used to establish a channel for communication between the external cluster and the components of the ACK control plane. This topic introduces ack-cluster-agent, and describes the usage notes and release notes for ack-cluster-agent.

Introduction

ack-cluster-agent is the ACK registration agent that runs as a Deployment in your external Kubernetes cluster. ack-cluster-agent receives requests from ACK Stub and forwards them to the Kubernetes API server of your external cluster. ack-cluster-agent also receives responses from the API server and forwards them to ACK Stub.

Usage notes

For more information about how to use ack-cluster-agent, see Overview.

Release notes

March 2022

Version	Image address	Release date	Description	Impact
v1.13.1.58-gf538528-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/ack-cluster-agent:v1.13.1.58-gf538528-aliyun	2022-03-16	A field is added to indicate the sources of clusters.	No impact on workloads

December 2021

Version	Image address	Release date	Description	Impact
v1.13.1.52-ge2a50ae-aliyun	registry.cn-hangzhou.aliyuncs.com/acs/ack-cluster-agent:v1.13.1.52-ge2a50ae-aliyun	2021-12-22	The issue that information about the providers of some instance types in external clusters cannot be obtained is fixed.	No impact on workloads

5.8.15. ack-slo-manager

ack-slo-manager is a key component used by Container Service for Kubernetes (ACK) clusters to support service level objective (SLO)-aware workload scheduling. ack-slo-manager improves resource utilization and ensures the performance of your applications. This topic introduces ack-slo-manager and provides usage notes and release notes for ack-slo-manager.

Introduction

ack-slo-manager is a key component used by ACK clusters to support SLO-aware workload scheduling. SLO-aware workload scheduling is a method to schedule different types of workloads to the same node. The workloads run at different resource demand levels to meet SLOs. This improves resource utilization and ensures the performance of the workloads. For more information about the use scenarios, refer to the following topics:

- CPU Burst
- Topology-aware CPU scheduling
- Load-aware pod scheduling
- Dynamic resource overcommitment
- Memory QoS
- Obtain recommendations on resource specifications

Usage notes

To install ack-slo-manager, go to the **Marketplace** page in the ACK console.

1.
2.
3. On the **Marketplace** page, click the **App Catalog** tab. Then, enter *ack-slo-manager* in the search box on the tab and click **ack-slo-manager**.
4. On the **ack-slo-manager** page, click **Deploy**.
5. On the **Basic Information** wizard page, select a cluster and a namespace, and then click **Next**.
6. On the **Parameters** wizard page, configure the required parameters and click **OK**.

   After you deploy ack-slo-manager, click **Clusters** in the left-side navigation pane. On the **Clusters** page, click the name of the cluster or click **Details** in the **Actions** column. In the left-side navigation pane of the cluster details page, choose **Applications > Helm**. On the **Helm** page, if the **Status** column of ack-slo-manager displays **Deployed**, ack-slo-manager is deployed.

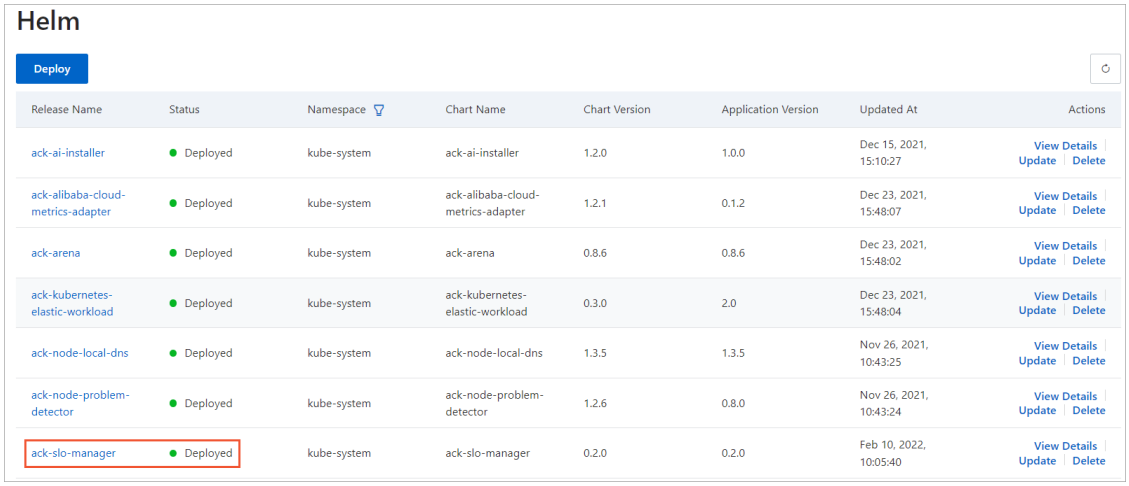

Release notes

February 2022

Version	Image address	Release date	Description	Impact
v0.3.0	registry.cn-hangzhou.aliyuncs.com/acs/ack-slo-manager:v0.3.0	2022-02-25	• Recommendations on resource specifications are provided. • The Memory Quality of Service (QoS) feature is supported. • Dynamic resource overcommitment is supported. • Internal API operations are optimized.	N/A

December 2021

Version	Image address	Release date	Description	Impact
v0.2.0	registry.cn-hangzhou.aliyuncs.com/acs/ack-slo-manager:v0.2.0	2021-12-10	• CPU Burst is supported. • Topology-aware CPU scheduling is supported.	N/A

September 2021

Version	Image address	Release date	Description	Impact
v0.1.1	registry.cn-hangzhou.aliyuncs.com/acs/ack-slo-manager:v0.1.1-c2ccefa	2021-09-02	Internal API operations are optimized.	N/A

July 2021

Version	Image address	Release date	Description	Impact
v0.1.0	registry.cn-hangzhou.aliyuncs.com/acs/ack-slo-manager:v0.1.0-09766de	2021-07-08	Workload-aware scheduling is supported.	N/A