

Enterprise Distributed Application Service (EDAS)

Product Introduction

Product Introduction

What is EDAS

Enterprise Distributed Application Service (EDAS) is a PaaS platform developed by Alibaba Middleware team to provide highly available and distributed Internet architecture solutions for enterprises.

EDAS offers a rich set of robust features, which have withstood the tough tests of Alibaba's key e-commerce platforms during the past few years. Fully utilizing Alibaba Cloud's existing resource and service management systems, EDAS provides functions such as distributed service frameworks, service governance, centralized configuration management, distributed tracing, capacity planning, high availability, and data-driven operations.

With EDAS, you can easily build large-scale distributed systems for your enterprise to host your applications and enable a microservice architecture to address increasingly demanding business needs.

Features

As a key part of the Alibaba Group distributed service architecture, EDAS delivers the capabilities for full life cycle application management and comprehensive O&M mechanisms.

Apache Tomcat container

As the basic container for running applications on EDAS, EDAS Container is fully compatible with Apache Tomcat. EDAS Container integrates with Alibaba middleware technology stack to provide significant enhancement in container startup, monitoring, stability, and performance.

Application-centric PaaS platform

Application management and O&M

In the EDAS console, you can manage the full life cycle of applications, such as creating, deploying, starting, stopping, scaling up, scaling down, and deleting applications. Leveraging Alibaba's rich experience in operating and maintaining large-scale clusters, you can easily operate and maintain applications deployed on thousands of instances.

Application deployment

You can use WAR/JAR packages or Docker images to publish applications.

Auto scaling

EDAS supports scaling in or out capabilities for applications both manually and automatically. With real-time monitoring of CPU, memory, and workload, you can complete scale-in or scale-out in seconds.

Primary and sub-accounts

EDAS allows you to create primary and sub-accounts according to the organization of your departments, teams, and projects. Additionally, ECS resources are organized based on accounts so that you can easily allocate resources based on different accounts.

Roles and permissions

The maintenance of an application typically involves its developers, DevOps engineers, and machine resource managers. Since different roles need to perform different management activities on the application, EDAS provides a role and permission control mechanism that allows you to define roles and manage permissions.

Namespace

Namespace is designed to logically isolate ECS instances within the same region. You can deploy applications in a specific namespace to prevent service calls from applications in other namespaces.

Rich distributed services

Dubbo/RPC

To meet the demand for massive transformation of Alibaba's e-commerce platforms to distributed architectures since 2007, Alibaba developed the distributed service frameworks of HSF (High-Speed Service Framework) and Dubbo. Built on a high-performance network communication framework,

HSF is a distributed service framework for enterprise Internet architectures, providing various proven features such as service publishing, registration, calling, routing, authentication, rate limiting and throttling, service degradation, and distributed tracing.

Spring Cloud

EDAS provides complete development environment for both Spring Cloud and Spring Boot. The applications you developed in Spring Cloud can be easily deployed on EDAS without modifying any codes.

Distributed configuration management

EDAS provides efficient distributed configuration management that allows centralized management of all configuration information across your distributed system. Furthermore, EDAS enables you to modify the configuration information in the EDAS console, and notify all the machines of the modification within seconds.

Distributed task scheduling

With its PaaS distributed task scheduling service, SchedulerX, EDAS allows you to configure any periodically scheduled tasks, be it standalone or distributed. It also allows you to manage the task running periods and query the running history of the tasks. It applies to task scheduling scenarios such as migrating historical data at two o' clock every morning, triggering a task every five minutes, or sending a monthly report on the first day of each month.

Transaction service

Global Transaction Service (GTS) is a high-performance, highly reliable middleware that ensures transaction consistency for distributed transaction systems. By using EDAS with GTS, you can easily implement distributed database transactions, multi-database transactions, messaging transactions, transactions spanning multiple services, and any combination of these transactions with a rich set of policies, ensuring both high usability and performance.

Maintenance and service governance

Service authentication

HSF is designed to ensure the reliability and security of every distributed call. Strict authentication is implemented in every phase, from service registration and subscription to service call.

Rate limiting and throttling

You can set a number of rate limiting rules for each application to control service traffic and ensure

proper service functionality. EDAS supports the configuration of volume control rules by both QPS and thread to ensure stability at traffic peaks.

Graceful service degradation

Each application may call many external services. In this case, you can configure degradation rules for these services to pinpoint and block out deteriorating services. This helps ensure that your applications run properly and prevents compromising your application's performance through its dependency on deteriorating services. EDAS allows you to configure degradation rules by response time, which effectively blocks deteriorating services at traffic peaks.

Automatic load test

The automatic load test tool is designed for daily use. With 100% usage of actual traffic in production environments, you can implement an authentic online stress test by controlling the weight for each service with no compromise on stability.

Comprehensive monitoring and digitized operations

Distributed tracing

EDAS EagleEye provides analysis of every service call, message sending, and database access occurred within the distributed system to help you identify system bottlenecks and risks precisely.

Service monitoring

EDAS monitors service calls made by your application in terms of QPS, response time, and error rate of your services.

Tenant-level service report

EDAS shows the service performance data of all applications at tenant level, including service call volumes in the recent 24 hours, average response time (RT) of service calls, and service call error numbers.

IaaS monitoring

EDAS allows you to monitor the servers in terms of basic metrics such as CPU, memory, network, and disk utilization as well as system workload.

Benefits

EDAS supports more than 99% of Alibaba's large-scale application systems, including all the key online systems that manage transactions, products, stores, logistics, and customer reviews, and provides highly stable and reliable services.

High reliability

- Used and tested within Alibaba for about ten years.
- Supports the stable operation of all Alibaba's key applications.
- Tested and proven during the past Double 11 shopping festivals.
- Provides sound authentication mechanism that ensures every single service call is secure and reliable.

Full-fledged features

- Supports full life cycle management of applications with a sound PaaS platform.
- Provides a set of service governance solutions to enable effective management of distributed services.
- Provides a comprehensive application diagnostics system to help you easily identify the root cause of issues.
- Provides online stress test and capacity planning capabilities to help you easily access online operation performance metrics and real-time server workloads.
- Supports automatic scaling to help you easily cope with unexpected traffic peaks.

In-depth monitoring

- Provides in-depth and global metrics reporting.
- Performs all-round monitoring and diagnosis.
- Tracks every single distributed call.
- Conducts dependency analysis of every system bottleneck.

Openness

- Contributed multiple Internet middleware components to open-source community.
- Donated top-grade Apache project, which is widely acclaimed in the industry.
- Easily replaceable with open-source software without technology hard bundle.

Scenarios

Publish and manage applications

Application deployment and management can be complicated in complex cloud environments. For locally developed applications, you need to allocate appropriate servers, and log on each server to deploy them. It is also required to restart and scale up the applications as your business keeps growing. The increasing number of servers poses a great challenge to the maintenance staff.

For this scenario, EDAS provides a visual friendly platform to deploy and manage applications, which allows you to easily manage the full life cycle of applications in the web console regardless of the scale of the cluster.

Build distributed systems

How to ensure reliable service calls between systems in a distributed architecture is always a headache when you transform a centralized system to a distributed system. For example, you have to nail down a lot of technical details in network communication and the design of serialization protocols.

EDAS provides a high-performance RPC framework, which allows you to build highly available distributed systems and take into consideration the technical details such as distributed service discovery, service routing, service calling, and service security easily.

Analyze system runtime status

After you develop an application and deploy it to the production environment, you often need to monitor the runtime status of the application, such as CPU usage, machine workload, memory usage, and network traffic. However, this type of basic monitoring cannot meet all business needs. For example, you can neither identify the bottleneck when the system runs slow nor troubleshoot when an error occurs upon opening a page.

To address these challenges, EDAS provides a series of digital operation components, which allows you to precisely monitor and trace every single component or service in the distributed system and pinpoint the bottleneck quickly.

Terms

This document defines and explains the proper names and terms related to EDAS.

Ali-Tomcat

Ali-Tomcat is a container that EDAS relies on to run services. It integrates service publishing, subscription, service call tracing, and other core functions. You must publish applications in this container in both development and runtime environments.

Dubbo

Dubbo is a distributed service framework that provides high-performance and transparent RPC remote service calls. Dubbo is the core framework of Alibaba's SOA service governance solution. Each day, over 2,000 services support over 3 billion access requests. This solution is widely used in various member sites of Alibaba Group.

EDAS

EDAS, short for Enterprise Distributed Application Service, incorporates the tested and proven distributed products of the Alibaba middleware team as its core components. It provides highly available distributed solutions oriented to the enterprise-level cloud computing market and is a core product in Alibaba's enterprise Internet architecture solution. EDAS makes full use of Alibaba Cloud's resource management and service systems and introduces a whole set of mature distributed products from the middleware team. With EDAS, enterprise customers can easily build large-scale distributed application service systems.

EDAS Agent

EDAS Agent is a Daemon program installed on your ECS instances to implement communication between an EDAS cluster and the applications deployed on the ECS instances. It's used for application management, status reporting, and information retrieval. It also serves as the communication channel between EDAS console and your applications.

EDAS RPC service

The EDAS RPC service provides support for the Dubbo framework. An application developed in the Dubbo framework and deployed with the WAR package can be published seamlessly with EDAS, and can be managed with service governance and data operation functions.

EDAS application life cycle

As the basic management unit in EDAS, an application usually contains multiple machines. EDAS provides a comprehensive application life cycle management solution, covering the entire process from application publishing to operation, including application creation, deployment, startup, rollback, scaling, disabling and removal, and so on.

EDAS EagleEye monitoring system

The EDAS EagleEye monitoring system analyzes distributed system calls, message sending, and database access to identify system bottlenecks and risks precisely.

EDAS application instance quota

This quota sets the maximum number of instances for all applications held by the primary account and its sub-accounts. When the billing account allocates a quota to each primary account, the sum of the quotas of all primary accounts bound to the billing account cannot be greater than the total application instance quota of the billing account.

IaaS monitoring

As one of EDAS monitoring functions, it monitors many aspects of the application runtime status, including CPU, memory, load, network, and disk.

Distributed tracing

EDAS EagleEye analyzes every service call, sent message, and database access by the distributed system to help you precisely identify the bottlenecks and risks of the system.

Distributed configuration management

This EDAS function centrally manages distributed system configuration information in the EDAS Console. It can add, modify, and delete configurations in real time and can push configuration updates globally or within a specific application.

Distributed task scheduling

Distributed task scheduling allows you to configure any standalone or distributed tasks for periodic scheduling. You can also manage the periods and query the history of the tasks. This service is suitable for task scheduling scenarios, for example, migrating historical data at 02:00 every morning, triggering a task every five minutes, or sending a monthly report on the first day of each month.

Service authentication

High-Speed Service Framework (HSF) is designed to ensure the reliability and security of each distributed service call. Rigid service authentication is applied in every phase, from service registration to service subscription and service calling.

Service degradation

Contrary to rate limiting and throttling, service degradation pinpoints and blocks poor services that your application calls. This feature ensures the stable operation of your application and prevents the functionality of your application from being compromised by dependency on poor services. EDAS allows you to configure degradation rules based on response time, preventing your application from depending on poor services during traffic peaks.

Service call monitoring

This is another EDAS monitoring function. It can fully monitor the service calls made by your application in terms of the service QPS, response time, and error rate.

Rate limiting and throttling

EDAS allows you to configure throttling rules for the services provided by each application to ensure

service stability. Throttling rules can be configured based on QPS and threads to ensure the system's best operation stability during traffic peaks.

Automated load testing

The automated load testing tool helps you make performance testing a part of your daily routines. All the traffic for automated testing is the actual data in the production environment. By controlling the weights of each service, you can perform authentic online load testing without undermining stability.

Release notes

Version	Release Date	Description
2.15.0	2017-09-30	Added virtual region support, and multi-environment isolation support.
2.14.1	2017-08-29	Fixed and optimized a few application deployment and publishing issues.
2.14.0	2017-08-23	Added application deployment and publishing function with batch mode support.

2.15.0

New features

PaaS

- **Virtual region** allows you to create virtual regions in existing Alibaba Cloud regions, so a single tenant can have multiple isolated environments (for example, development, testing, and production) in a single region.
- **Component Center page** allows you to view all the sub-services of the EDAS platform. Here, application diagnostic components have been linked to other modules in the system menu.
- **New Overview page of the Console** displays four metrics: application count, machine count, service count, and the number of applications published in the past

seven days. In addition, it shows three recent alarms from all the applications under the current account, and the service call trendline.

- **EDAS container version 3.3.2**, the latest version, provides all the functions that previous versions support, such as traffic greyscale, JAR package deployment, and HSF RESTful service calls.
- Docker applications can be developed in Spring Cloud/Boot framework and packed and run as FatJar.
- The application publish change details page is optimized with health check and optimized HSF service disabling.

Monitoring & service governance

- **Slow SQL statement parsing** is added to service tracing. When slow database access is detected, you can view the executed SQL statements.
- **Optimized rate limiting and service degradation** provide improved performance and allow you to customize the redirect page for rate throttling.

2.14.1

New features

PaaS

- **Application status change prompts** shows a global message “This application is performing a XXX change. Click for more information.” at the top of each page in a changing application.
- The change details page automatically populates the current change type and the deployment package name and version.
- The change details page shows the application group. If the action does not apply to the entire group, then the specific machines are added to the “Change Target” field.
- Change records are color coded to indicate the status of the changes. Successful changes are green, failed red, terminated yellow, and running black.
- Fixed a bug that the WAR package download failure caused by a full machine disk goes undetected during application deployment.
- Fixed a bug that the “Description” field in the publishing list is truncated.
- Fixed a bug that the “conversion successful” message is displayed despite the 503 error during the conversion of a Docker cluster node.
- Fixed a bug that the JVM xms parameter can be set to be greater than 70% of the memory size for Tomcat in a Docker application.

2.14.0

New features

PaaS

Application publishing order is a new way to deploy and publish an application. In this way, you can view the success or failure status of each stage in the publishing process. You can also stop a publishing when necessary.

Application batch publishing deploys and publishes applications in batches. You can define the machines for each batch and suspend the publishing between batches, so you can decide whether to continue the publishing or stop it and roll back the process.