Apusic 金蝶天燕

生成时间：2026-05-19

User Manual

Apusic Application Server V10

版权声明

本文档所涉及的软件著作权、版权等知识产权已依法进行了注册，由金蝶天燕云计算股份有限公司合法拥有。受《中华人民共和国著作权法》《计算机软件保护条例》《知识产权保护条例》和相关国际版权条约、法律、法规以及其它知识产权法律和条约的保护。未经授权许可，不得非法使用。

免责声明

本文档包含的版权信息由金蝶天燕云计算股份有限公司合法拥有，受法律的保护，金蝶天燕云计算股份有限公司对本文档可能涉及到的非金蝶天燕云计算股份有限公司的信息不承担任何责任。在法律允许的范围内，您可以查阅并仅能够在《中华人民共和国著作权法》规定的合法范围内复制和打印本文档。任何单位和个人未经金蝶天燕云计算股份有限公司书面授权许可，不得使用、修改、再发布本文档的任何部分和内容，否则将被视为侵权，金蝶天燕云计算股份有限公司有依法追究其责任的权利。

本文档如有更新，不另行通知。对本文档中的问题您可向金蝶天燕云计算股份有限公司告知或查询。未经本公司明确授予的任何权利均予保留。

商标声明

Apusic是深圳市金蝶天燕云计算股份有限公司向中华人民共和国国家商标局申请注册的注册商标，注册商标专用权由金蝶天燕合法拥有，受法律保护。未经金蝶天燕的书面许可，任何单位及个人不得以任何方式或理由对该商标的任何部分进行使用、复制、修改、传播、抄录或与其它产品捆绑使用销售。凡侵犯金蝶天燕商标权的，金蝶天燕将依法追究其法律责任。本文档提及的其他所有商标或注册商标，由各自的所有人拥有。

目录

1 Preface

This document is the instructions for using the Kingdee Apusic Application Server V10, commonly known as AAS V10. It provides detailed information on installing and configuring the Kingdee Apusic Application Server.

1.1 object-oriented

This manual is mainly aimed at developers who use the Kingdee Apusic Application Servers for application development, as well as relevant management and operation personnel.

1.2 release-noted

Date	Version	Product	Description
December 2023	V10E01F01	AAS V10	Update license authorization authentication method

2 Overview

The Kingdee Apusic application server provides a simple and fast development and operation platform for complex applications, and provides features such as easy scalability, scalability, and high security for distributed enterprise-level applications. The following will introduce the architecture of the Apusic application server and the services and functions it provides, to demonstrate the support capabilities of the Apusic application server for developing large-scale application systems.

2.1 Basic Concepts

Before using application servers correctly to deploy and manage applications, it is necessary to first understand the following basic concepts.

1.Apusic Application Server

A server is the physical deployment unit of an application server. Intuitively speaking, it is a physical installation of an application server on a user's machine.

2.domain

A domain provides a common authentication and administration point for a collection of zero or more server instances. The administration domain encompasses several manageable resources, including instances, clusters, and their individual resources. A manageable resource, such as a server instance, may belong to only one domain.

3.instance

An instance in Apusic Application Server has its own Java EE configuration, Java EE resources, application deployment areas, and server configuration settings.For many users, one server instance meets their needs. However, depending upon your environment, you might want to create additional server instances. For example, in a development environment you can use different server instances to test different Apusic Application Server configurations, or to compare and test different application deployments. Because you can easily add or delete a server instance, you can use them to create temporary "sandbox" areas to experiment with while developing.

4.node

Each node corresponds to a host equipped with Apusic Application Server software, and a corresponding node must exist on the host where the application server instance is located. Node configuration information includes the host name and the location where the application server is installed on the host.

5.cluster

A server cluster is a collection of logical entities consisting of multiple server instances that share the same applications, resources, and configuration information. Regardless of whether the server instance is on the same host or different hosts, the application server can manage all instances in the cluster as a unit in the management control center.

2.2 Basic Directory Values

There are basic directory value of the Apusic Application Server:

Variable	Discription	Default Value
JAVA_HOME	Java environment variables	
APUSIC_HOME	The installed directory of Apusic Application Server	
DOMAIN_HOME	The domain directory of Apusic Application Server	APUSIC_HOME/domains/[domainname]

2.3 Basic Manager Value

There are basic manager value of the Apusic Application Server:

Name	Default Value
Domain name	mydomain
asadmin command-line utility	APUSIC_HOME/bin
Configuration file	DOMAIN_HOME/config
Log file	DOMAIN_HOME/logs

2.4 Basic Ports

There are basic ports of the Apusic Application Server:

Functional Module	Port
Management console port	6848
HTTP port	6888
HTTP SSL port	6887
IIOP SSL port	6838
IIOP MUTUALAUTH port	6839
JMS port	6876
IIOP port	6837
JMX port	6886
OSGI SHELL port	6866
JAVA DEBUGGER port	8000

2.5 Directories and Files

There are directories of the Apusic Application Server :

Name	Description
AAS-V10.tar.gz	Installation package for Apusic Application Server V10
ApusicAS/	Default directory for Apusic Application Server V10
aas/	Default installation directory for Apusic Application Server V10
aas/bin	The command-line utility of Apusic Application Server V10
aas/config	The configurations of Apusic Application Server V10
aas/docs	The files of Apusic Application Server V10
aas/domains	The domains of Apusic Application Server V10
aas/jmods	The jmods file of Apusic Application Server V10
aas/lib	The library files of Apusic Application Server V10
aas/modules	The module files of Apusic Application Server V10
aas/osgi	The osgi files of Apusic Application Server V10
aas/templates	The template file of Apusic Application Server V10
aas/license.xml	The license file of Apusic Application Server V10
bin/	Directory for storing command line
install/	Directory for storing external application packages
javadb/	Directory for storing javadb
mq/	Directory for storing MQ application
samples/	Directory for storing samples
tools/	Directory for storing tool installation packages

3 Setup

There are several ways to set up Apusic Application Server for running on different platforms.

3.1 Java Environment

The java environment supported by Apusic Application Server contains the following:

• Oracle JDK 8+

• Open JDK 8+

• IBM JDK 8+

3.2 Operating System Environment

The Apusic Application Server is compatible with various commonly used browsers.

System Components	System Requirements
JAVA Environment	JDK 1.8 and above versions
Memory size	2GB+
Hard disk space size	10GB+
Browsers	Firefox，Chrome，360，Microsoft Edge..

3.3 Install

3.3.1 Windows

Installing Apusic Application Server on Windows can be done easily using the Windows installer. Its interface and functionality is similar to other wizard based installers,just follow the instructions to operate.

If you get the zip package like AAS-V10.zip or AAS-V10.tar.gz,simply unzip to the specified directory to complete the installation.

If you get the GUI package like AAS-V10.exe，execute scrips and follow instructions.

3.3.2 Unix deamon

If you get the zip package like AAS-V10.zip or AAS-V10.tar.gz,simply unzip to the specified directory to complete the installation.

If you get the GUI package like AAS-V10.bin，execute scrips and follow instructions.

3.4 Uninstall

If you want to uninstall Apusic Application Server,ensure important information is backed up,then delete install directory finish uninstall operation.

3.5 Start domain

When installing Apusic Application Server V10, it comes with the default domain `mydomain`.You can start domain by the `start-domain` subcommand.

```
asadmin start-domain
```

Please set the password of administrators when you start domain at first time.

3.6 Stop domain

You can stop domain by the `stop-domain` subcommand.

```
asadmin stop-domain
```

4 Console User Guide

The Apusic Application Server management console provides an interface based unified management configuration platform for configuration, management, and monitoring, supporting local and remote access to the management console.

When the domain is in a RUNNING state, open a Web browser and enter the domain's URL. Use the Administration Port value you entered when creating the domain. For example:

```
https://localhost:6848
```

4.1 Administrators

The Apusic Application Server has three default administrator's roles,system administrator,security administrator and audit administrator.The default administrator role has corresponding default users with different management permissions,the user list is as follows:

Role Name	Description	Default User
sysadmin	The system administrator is primarily responsible for managing application deployment, resource management, and configuration management, among other functions. This role ensures that systems are properly set up, configured, and maintained to meet the needs of the organization. The system administrator is responsible for deploying new applications and updates, configuring system settings and parameters to optimize performance and security. They also monitor system performance and troubleshoot issues as needed to ensure the smooth operation of the systems.	admin
security	The security is primarily responsible for managing information related to role management, user management, password policies, and other security-related aspects. This role ensures that access controls are properly implemented and maintained, user accounts are managed securely, and password policies adhere to best practices to minimize the risk of unauthorized access or data breaches.	secure
auditor	The auditor primarily manages logs and operational audit information. This role is responsible for reviewing and analyzing system logs, transaction logs, and other relevant data to identify potential security incidents, unauthorized access attempts, or any other irregularities that may indicate a breach of security policies or procedures. By monitoring and auditing these activities, the security auditor helps to ensure the integrity, confidentiality, and availability of the system and its data.	audit

4.2 Domain

4.2.1 Domain Attributes

Use the Domain Attributes page to set advanced domain attributes for the Apusic Application Server.

The Domain Attributes page contains the following options.

- Load console after domain administration server startup

  Determines when to load the Administration Console:Usage BasedLoads the Administration Console after the Domain Administration Server (DAS) if the console has been used in the last 24 hours or if other server instances exist aside from the DAS. This is the default.AlwaysAlways loads the Administration Console after the DAS. The console is loaded after the main Apusic Application Server startup sequence has completed.NeverNever loads the Administration Console after the DAS. The console is loaded when access to it is requested.

- Check Instance State

  If this option is selected,enabled instance check. When there is an instance downtime, it will automatically restart.Disabled by default.

- Check Instance Time

  The time interval for instance downtime detection, if set to 60 seconds, indicates that it will be detected every 60 seconds.The default value is 30 seconds.

- Create Conflict Report

If this option is selected, enable the function of creating a class conflict detection report. When class conflict information is detected in the Class Conflict section, click Download Test Report to download a .txt format detection report.

- Application Root

  The full path of the directory where deployed applications reside. The default value is *domain-dir* `/applications` .

- Log Root

  Identifies where the server log files are kept. The default value is *domain-dir* `/logs` .

- Locale

  If this field is left blank, the default locale of the host will be used.A locale is an identifier that specifies a particular combination of language and region. For example,the locale for Chinese is `zh-CN` , the locale for American English is `en_US` , and for Japanese it is `ja_JP` .After the server is set to display in English, the client also needs to set the preference to use English. For example, the Firefox browser needs to set the value of `intl.accept_language` to `en-US` in `about:config` at the top.

4.2.2 Applications Configuration

Use the Applications Configuration page to set properties that help ensure that changes to deployed applications are detected and that the modified classes are reloaded.

The Applications Configuration page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Reload

  If this option is enabled, modified classes are reloaded dynamically. The server periodically checks for changes in the deployment directory and redeploys the application, automatically and dynamically, with the changes. This is useful in a development environment, because it allows code changes to be tested quickly. In a production environment, however, dynamic reloading might degrade performance. In addition, whenever a reload is done, the sessions at that transit time become invalid. The client must restart the session. This option is enabled by default.

- Reload Poll Interval

  Identifies how often the server should check for code changes in the deployed applications and modules. The default value is 2.

- Admin Session Timeout

  Identifies the number of minutes of inactivity before the Administration Console times out and you have to log in again. The default value is 60.

- Deploy Recover Enabled

  If this option is selected,place the uninstalled application in the recover folder.Disabled by default.

- RecoverDir

  Set the directory of uninstalled application.The default value is `${com.apusic.aas.instanceRoot}/recover` .

- RecoverFileExpireDays

  Set the validity period of recover documents,0 represents permanent retention.The default value is 180 days.

- Auto Deploy

  If this option is enabled, applications in the auto-deploy directory are deployed automatically. This option is enabled by default.

- Auto Deploy Poll Interval

  Specifies how often the server should check the auto-deploy directory for application or module changes. Changing the poll interval does not affect the time taken for deploying an application or module. The default value is 2.

- Auto Deploy Retry Timeout

  Specifies the number of seconds that a partially copied file can remain unchanged in size before an error occurs. The default value is 4.If a file is copied slowly into the auto-deploy directory, the file can appear before the entire file is copied. As a result, the attempt to automatically deploy the application fails. If an

attempt to automatically deploy an application fails for this reason, the Apusic Application Server tries to deploy the application again. If the size of the partially copied file remains unchanged for the specified period of time, an error occurs.

- Auto Deploy Directory

  Identifies the directory to monitor for automatic deployment of applications. The default value is *domain-dir* `/autodeploy`.

- XML Validation

  Specifies the type of XML validation to be performed on standard and Apusic Application Server deployment descriptors. The type may be any of the following:FullIf XML validation fails, deployment fails. This value is the default.ParsingXML validation errors are reported, but deployment occurs.NoneNo XML validation is performed.

- Verifier

  If this option is enabled, the verifier is run before autodeployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the file. Verification of large applications is often time-consuming. This option is disabled by default.

- Precompile

  If this option is enabled, JavaServer Pages (JSP) files are precompiled. If this option is disabled, JSP files are compiled at runtime when they are first accessed. This option is disabled by default.

- Additional Properties

  Additional properties for applications. The Apusic Application Server does not define any additional properties for applications.

4.2.3 Update Administrator Password

Use the Update Administrator Password page to change the password for the current administrator account.

Note:

If secure administration is enabled, you cannot change the password to an empty value.

The Update Administrator Password page contains the following options.

- User ID

  The user ID of the current administrator. This is a read-only field.

- Group List

  The group to which the administrator belongs. This a read-only field. The value is `asadmin`.

- New Password

  The new password for the administrator.

- Confirm New Password

  The new password for the administrator.

4.2.4 Password Aliases

Use the Password Aliases page to create and manage password aliases.

The Password Aliases page displays a list of password aliases in the domain. For each alias, the following information is displayed:

- Name

  The name that was assigned to the password alias when it was created. Clicking the name opens the Edit Password Alias page for the alias.

The Password Aliases table also contains the following options.

- New

  Button to create a password alias.

- Delete

Button to delete one or more selected password aliases.

4.2.4.1 New Password Alias

Use the New Password Alias page to create a new password alias for use in a password file or the domain configuration file.

The New Password Alias page contains the following options:

- Alias Name

  The name of the alias. This name must be unique across all password aliases in the domain.

- Password

  The password that the password alias stores in an encrypted form.

- Confirm Password

  The password that the password alias stores in an encrypted form. The Password and Confirm Password values are compared to ensure the password was entered without typing mistakes.

4.2.4.2 Edit Password Alias

Use the Edit Password Alias page to change the password associated with an existing password alias.

The Edit Password Alias page contains the following options:

- Alias Name

  The name of the alias. This option is not editable.

- Password

  The password that the password alias stores in an encrypted form.

- Confirm Password

  The password that the password alias stores in an encrypted form. The Password and Confirm Password values are compared to ensure the password was entered without typing mistakes.

4.2.5 Configuration backup

Use the Applications Configuration page to backup configuration information for clusters or instances.

The backup file is stored in the `${DOMAIN_HOME}/config/config-backup` directory.

- Backup

  Button to backup the configuration of instance or clsuter.

- Reduction

  Button to Restore the backup configuration file.

- Download

  Button to dowload the backup configuration file

4.2.6 Class conflict

Use the Class conflict page to perform class conflict detection.

Upload war/jar/ear packages for class conflict detection. The uploaded package files will be compared with the classes in Apusic Application Server, and any conflicts will be displayed in the Class conflict page.

Three methods can be used for detection:

- Uploading jar/war/ear packages to the server
- Accessing local jar/war/ear packages from Apusic Server
- Selecting to deploy applications.

If "Create Conflict Report" is enabled in the Domain Attributes, a "Download Test Report" will be displayed after detection. You can click to download. A detection report in `.txt` format will be downloaded. You can also view it under `${DOMAIN_HOME}/mydomain/conflict-reports/`

4.3 Server(Admin Server)

4.3.1 General Information

The Apusic Application Server creates one application server instance, called `server` at the time of installation. You can delete the server instance and create a new instance with a different name if you prefer.

Each Apusic Application Server instance has its own Java configuration, Java resources, application deployment areas, and server configuration settings. Changes to one application server instance have no effect on other application server instances. You can have one application server instance within one administrative domain.

For many users, one application server instance meets their needs. However, depending upon your environment, you might want to create one or more additional application server instances. For example, in a development environment you can use different application server instances to test different Apusic Application Server configurations, or to compare and test different application deployments. Because you can easily add or delete an application server instance, you can use them to create temporary "sandbox" areas to experiment with while developing.

In addition, for each application server instance you can also create virtual servers. Within a single installed application server instance you can offer companies or individuals domain names, IP Addresses, and some administration capabilities. For the users, it is almost as if they have their own web server, without the hardware and basic server maintenance. These virtual servers do not span application server instances. For more information about virtual servers, see JVM General Settings).

In operational deployments, for many purposes you can use virtual servers instead of multiple application server instances. However, if virtual servers do not meet your needs, you can also use multiple application server instances.

An Apusic Application Server instance is not started automatically. Once you start an instance, the instance runs until you stop it. When you stop an application server instance, it stops accepting new connections, then waits for all outstanding connections to complete. If your machine crashes or is taken offline, the server quits and any requests it was servicing may be lost.

Use the General Information page to verify Apusic Application Server settings and to view Java Virtual Machine data.

The General Information page contains the following information.

- Stop

  Click the Stop button to stop the Apusic Application Server.

- Restart

  Click the Restart button to restart the Apusic Application Server.

- View Log Files

  Click the View Log Files button to view log files for a Apusic Application Server instance or cluster.

- Rotate Log

  Click the Rotate Log button to rotate the log file for the Admin Server (named `server`).

- Recover Transaction

  Click the Recover Transaction button to recover transactions for the Admin Server (named `server`) on the Recover Transactions page.

- Secure Administration

  Click the Secure Administration button to enable or disable secure administration on the Secure Administration page.

- View Access Log

  Click the View Access Log Files button to view access log files for a Apusic Application Server instance or cluster.

- Name

  The name of the current server.

- Status

  The current status of the server instance. The server can be stopped, started, or running.

- Uptime

  The number of hours and minutes that the server instance has been continuously running. This is read-only.

- HTTP Load Balancer

  The name of the load balancer, if it is enabled for the server instance.

- JVM

  If you click JVM Report, a separate window opens and displays reports on the Java Virtual Machine, including a summary report, memory management and garbage collection report, class loading report, and a current thread dump.

- Configuration Directory

  The directory on the host machine that contains configuration files, such as `domain.xml` .

- Installed Version

  The product version identifier.

- Secure Administration

  The current status of secure administration. Possible values are Enabled and Not Enabled.

- Debug

  Identifies whether debugging is enabled. If debugging is enabled, the port number is also displayed.

- Up Time

  The running time of Apusic Application Server.

- License Type

  The type of the license for the Apusic Application Server.

- License expire date

  The expire date of the license file for the Apusic Application Server.

- License create date

  The start date of the license file for the Apusic Application Server.

- License version

  The license version of Apusic Application Server.

- License File

  Display the license of Apusci Application Server. If you need to update, copy the content of the license file to the input box, replace the current license, and click "Update License File" to take effect in real time. If the license of the node needs to be replaced, replace it in the corresponding location in Node page. After enabling email verification, a reminder message will be sent to the email address 7 days before the license expires. The reminder will only stop sending after the license is replaced.

- HTTP Port(s)

  The currently configured ports for HTTP requests and responses.

- IIOP Port(s)

  The currently configured ports for IIOP requests and responses.

4.3.2 Resources

Use the Resources page to enable, disable, or create resources associated with the administration server.

4.3.3 Properties

Use the System Properties page to view and manage Apusic Application Server properties.

A system property defines a common value for a setting at the server level. You can refer to a system property in an Admin Console text field by enclosing it in a dollar sign and curly braces. For example, to refer to a property named *prop-name*, you would use the syntax `${` *prop-name* `}`.

For each property, the System Properties page contains the following information.

- Name

  The name of the property.

- Value

  The value of the property.

- Description

  An optional description of the property.

A number of system properties are predefined but are not shown on this page. For details, see Predefined System Properties.

4.3.4 Predefined System Properties

The following system properties are predefined for the Apusic Application Server.

- `com.sun.aas.installRoot`

  The directory where the Apusic Application Server is installed.

- `com.sun.aas.instanceRoot`

  The top-level directory for an Apusic Application Server domain.

- `com.sun.aas.hostName`

  The name of the host (machine).

- `com.sun.aas.javaRoot`

  The installation directory for the Java runtime.

- `com.sun.aas.imqLib`

  The library directory for the Oracle Message Queue software.

- `com.sun.aas.domainName`

  The name of the domain. This property is not used in the default configuration, but can be used to customize configuration. The default value is `domain1`.

4.3.5 Monitoring

Use the Monitoring page to view monitoring data for server.

For a description of the monitored properties, see the Monitoring module.

4.3.6 Batch

4.3.6.1 Batch Job Executions

Apusic Application Server provides a batch runtime for the scheduling and execution of batch jobs. Batch jobs are typically long-running, bulk-oriented tasks that contain a series of steps. Batch applications submit jobs to the batch runtime and provide instructions about how and when to execute the steps. The batch runtime processes the steps and stores information about jobs in a job repository. In Apusic Application Server, the job repository is a database.

Use the Batch Job Executions page to view details about recorded executions of batch jobs.

The Batch Job Executions page contains the following information.

- Instance Name

- Cluster Name

  The name of the Apusic Application cluster or server instance to which the configuration applies.

- Batch Jobs

  A list of recorded batch job executions for the instance or cluster. For each execution, the following information is displayed:Execution IDThe ID assigned to the execution of the batch job. Clicking the execution ID opens the Batch Job Execution Details page for the execution. A new execution is created the first time a job is started and every time the existing execution is restarted.Job NameThe name of the job.Batch StatusThe status of the execution as set by the batch runtime.Exit StatusThe status of the execution as set by the Job XML for the job or by the batch application. By default, the exit status and the batch status are the same unless the exit status is explicitly overridden.Instance IDThe ID assigned to the instance of the batch job. An instance of a batch job can have multiple executions.Start TimeThe time the execution started.End TimeThe time the execution finished.

4.3.6.2 Batch Job Execution Steps

The Batch Job Execution Steps page displays details about steps in a specific batch job execution. The following information is displayed.

- Instance Name
- Cluster Name

  The name of the Apusic Application cluster or server instance to which the configuration applies.

- Job Name

  The name of the batch job.

- Execution ID

  The ID assigned to the execution of the batch job. A new execution is created the first time a job is started and every time the existing execution is restarted.

- Job Steps

  The steps in a specific batch job execution. For each step, the following information is displayed:Step NameThe name of the step.Batch StatusThe status of the step as set by the batch runtime.Exit StatusThe status of the step as set by the Job XML for the job or by the batch application. By default, the exit status and the batch status are the same unless the exit status is explicitly overridden.Start TimeThe time the step started.End TimeThe time the step finished.Step MetricsMetrics for the step, listing the number of items read, written, committed, and so on.

4.3.6.3 Batch Job Execution Details

The Batch Job Execution Details page displays details about a specific batch job execution. The following information is displayed.

- Instance Name
- Cluster Name

  The name of the Apusic Application cluster or server instance to which the configuration applies.

- Job Name

  The name of the batch job.

- Execution ID

  The ID assigned to the execution of the batch job. A new execution is created the first time a job is started and every time the existing execution is restarted.

- Step Count

  The number of steps in the batch job execution.

- Batch Status

  The status of the batch job execution as set by the batch runtime.

- Exit Status

  The status of the batch job execution as set by the Job XML for the job or by the batch application. By default, the exit status and the batch status are the same unless the exit status is explicitly overridden.

- Start Time

  The time the batch job execution started.

- End Time

The time the batch job execution finished.

- Job Parameters

  The properties passed to the batch runtime for the batch job execution, listed as name/value pairs.

4.3.6.4 Batch Runtime Configuration

Use the Batch Runtime Configuration page to view and change the configuration of the batch runtime. The batch runtime uses a managed executor service and a data source to execute batch jobs. The managed executor service provides threads to jobs, and the data source stores job information. Batch runtime configuration data is stored in the `config` element in `domain.xml`.

The Batch Runtime Configuration page contains the following information.

- Instance Name
- Cluster Name

  The name of the Apusic Application cluster or server instance to which the configuration applies. This is a read-only field.

- Executor Service Lookup Name

  The JNDI lookup name of the managed executor service that provides threads to batch jobs. The default managed executor service is `concurrent/__defaultManagedExecutorService`.The managed executor service can be changed after a batch job has been submitted to the batch runtime without impacting execution of the job.

- Data Source Lookup Name

  The JNDI lookup name of the data source that stores information about current and past batch jobs. The default data source is `jdbc/__TimerPool`.Do not change the data source after the first batch job has been submitted to the batch runtime for execution. If the data source must be changed, stop and restart the domain and then make the change before any jobs are started or restarted. However, once the data source has been changed, information stored in the previous data source becomes inaccessible.

4.4 Certificate Management

The module used to manage certificates. After entering the management console, select Server Management Server, and then select the Certificate Management tab.

4.4.1 Concept

Certificate: A certificate is a medium for a public key, which also contains information about the certificate holder, such as name and address. The file extension for a certificate is usually crt or cert.

Certificate store: A repository for storing certificates. The most common certificate store formats are JKS, which is only used in Java programs, and pfx (or p12). In addition to certificates, a certificate store can also contain private keys. In SSL configurations, sometimes in addition to configuring a certificate store, it is also necessary to configure a trust certificate store for storing trusted third-party certificates during mutual authentication.

Certificate request file: When applying for a certificate from a CA, sometimes the CA will require a certificate request file csr file, which contains the applicant's personal information and can be opened using tools such as Notepad.

One-way authentication (or server-side authentication): This refers to the process where the client verifies the server's identity using the server's SSL certificate. In this scenario, only the server presents its SSL certificate to the client for verification.

Two-way authentication (or mutual authentication): This refers to the process where both the client and the server authenticate each other's identity using SSL certificates. In this scenario, the server presents its SSL certificate to the client for verification, and the client also presents its SSL certificate to the server for verification. This ensures mutual trust between the two parties.

4.4.2 Keystore Entities

A keystore is a repository that stores keys and certificates, and each individual item stored in it can be referred to as an "entity" or "entry". So, if you are referring to a specific item (like a certificate or a key pair) stored in a keystore, you can call it a "keystore entity" or "keystore entry".

Note:

- The default alias of the keystore entity used by the control platform is `kaas`, and the default alias of the keystore entity used by the `sec-admin-listener` in `default-config` is `aas-instance`. Currently, it is not possible to configure the SSL certificate for the control platform, namely `sec-admin-listener`.
- If you want to configure the ShangMi certificate, refer to the ShangMi configuration.

For each keystore entity, the following information is provided.

- Alias

  The alias of storing this entity.

- Keystore

  The keystore file of this entity.

- Details

  Detailed description of the keystore repository.

The Keystore Entities table also contains the following options.

- Add

  Button to add a keystore entity.

- Delete

  Button to delete a keystore entity.Before deleting, please ensure that the certificate has not been referenced.

4.4.2.1 Keystore entity reference

You can add a Listener when creating a new entity, indicating that the listener references the keystore entity. At the same time, set the "Certificate Nickname" in Configurations -> [server-config] -> Network Config -> Protocols. Only one alias for the keystore entity can be set, and if it is set during creation, it will replace the existing alias.

After setting the "Certificate Nickname" and saving it, click https to view the certificate information when accessing the browser.

Note:

The protocol needs to have "security" enabled to use HTTPS access, and the certificate needs to be valid.

4.4.2.2 Add Keystore Entities

Use the Add Keystore Entities page to create a new keystore entity.

The Add Keystore Entities page contains the following options.

- Instance name

  The instance of this keystore entity.Currently, only Server is allowed to be configured.

- Alias

  Specify the alias for storing this entity. If there is no specific description in the certificate or keystore file, you can name it freely; the format includes letters, numbers, horizontal lines, or underscores, and it must be unique and required.

- File

  Upload certificate or keystore file, supports PEM and JKS formats, required.

- Password

  The password for uploading the keystore. It is required to enter the password if it is available, but not required if it is not available.

- Listener

  Add entries to the listener. An empty value indicates that the entry is added to the instance library. It is recommended to fill in `http-listener-2`

4.4.3 Truststore Entities

A truststore is a repository that stores certificates that are trusted by a system or application. These certificates are used to verify the authenticity of remote parties (such as SSL/TLS servers) during secure communications. Each individual certificate or other trusted material stored in the truststore can be referred to as an "entity" or "entry".

For each Truststore entity, the following information is provided.

- Alias

The alias of storing this entity.

- Truststore

  The truststore file of this entity.

- Details

  Detailed description of the truststore repository.

The Truststore Entities table also contains the following options.

- Add

  Button to add a truststore entity.

- Delete

  Button to delete a truststore entity.Before deleting, please ensure that the certificate has not been referenced.

4.4.3.1 Add Truststore Entities

Use the Add Truststore Entities page to create a new truststore entity.

The Add Truststore Entities page contains the following options.

- Alias

  Specifies the alias for storing this entity. If the certificate or keystore file does not specify a specific name, you can name it arbitrarily; the format includes letters, numbers, horizontal lines, or underscores, and it must be unique and required.

- File

  Upload certificate or keystore file, imported certificate, supports PEM, DER, JKS formats, required item.

- Password

  The password for uploading the keystore. If there is a password, it needs to be entered. If there is no password, it can be left blank.

- Listener

  Add entries to the listener. Leave blank to add to the instance library. It is recommended to fill in `http-listener-2`

4.4.3.2 Truststore Entity Reference

You can add a Listener when creating a new instance, indicating that the listener references the truststore entity. At the same time, set the Trust store to `${com.apusic.aas.instanceRoot}/config/cacerts.jks` in the Configurations- [Server-config] - Network Config]- Protocols settings.

Note:

The protocol needs to enable "security" to use HTTPS access, and the certificate can only take effect

4.5 Applications

The Applications page displays a list of applications that are deployed on the Apusic Application Server. You can view and manage the applications that are already deployed, and you can deploy more applications.

For each application, the following information is provided.

- Name

  The application name.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Enabled or Status

  A check mark if the application is enabled, or an X if the application is disabled, if only the default server instance, `server`, exists.If clusters or other standalone server instances exist, select Enabled on Targets to view the targets on which the application is deployed.

- Engines

  The types of containers used by the application. Container types can be any of the following: `web``webservices``ejb``connector``appclient``weld` (the container for Contexts and Dependency Injection for the Java EE Platform applications)

- Action

  Links to the actions you can perform on the component after deploying it: Redeploy/Reload/backup for all components, Launch for web applications and application clients, and Download Client Stubs for application clients.

The Deployed Applications table also contains the following options.

- Deploy

  Button to deploy an application.

- Undeploy

  Button to undeploy one or more selected applications.

- Enable

  Button to enable one or more selected applications, if only the default server instance, `server`, exists.

- Disable

  Button to disable one or more selected applications, if only the default server instance, `server`, exists.

- Filter

  Drop-down list that filters applications by engine.

- Recover

  Button to the page of uninstalled applications.

- Backup List

  Button to application backup file display page.

4.5.1 Deploy Applications or modules

Use the Deploy Applications or Modules page to deploy an application.Follow the deployment wizard to operate.

The Deploy Applications or Modules page contains the following options for all applications.

4.5.1.1 Step1

- Location

  The location of the archive for the application that you are deploying.

  The following options specify from where the archive is accessible and whether the archive is a file or a directory.

  Packaged File to Be Uploaded to the Server

  The archive is in a file that resides on or is accessible from the client machine.

  The client machine is the host on which you are viewing the Administration Console through a browser.

  Local Packaged File or Directory That Is Accessible From the Apusic Application Server

  The archive is a file that resides on the server machine, or is an unpackaged application in an exploded directory.

  The server machine is the host that is running the Apusic Application Server domain administration server.

After selecting the file, you need to click the "Upload" button.

- Temporary location

The temporary location is used to temporarily store uploaded application files.

4.5.1.2 Step2

- Type

  The type of the application. Available choices are:

  - Web Application
  - Enterprise Application
  - EBA Application
  - Application Client
  - Connector Module
  - EJB Jar
  - Other

- Targets

  Clusters and standalone instances to which to deploy the application or module. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.This option is displayed only if clusters or standalone instances have been created in the domain.

Additional options for the various application types are described in the following sections.

- Additional Options for a Web Application
- Additional Options for an Enterprise Application
- Additional Options for an Application Client
- Additional Options for a Connector Module
- Additional Options for an EJB Jar

4.5.1.2.1 ADDITIONAL OPTIONS FOR A WEB APPLICATION

If the application type is Web Application, the following options appear.

- Context Root

  The path to the application. In the URL of the web application, the context root immediately follows the port number (`http://` *host* `:` *port* `/` *context-root* `/...`). The context root must start with a forward slash, for example, `/hello` .

- Application Name

  The name of the application.The name should be the unique.

- Virtual Servers

  The virtual servers associated with this application.The Virtual Servers option appears if only the default server instance, `server` , exists. If clusters or other standalone server instances exist, you can select virtual servers after deployment. Go to the Edit Application page, select the Target tab, and select Manage Virtual Servers for the desired target.

- Status

  If this option is selected, the application is enabled. This option is selected by default.

- Implict CDI

  Implicit discovery of CDI beans.This option is selected by default.

- Delegate

  If this option is selected,the parent class loader for priority,if not,the child class loader for loading.

- Secret Level

The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Availability

  If the Enabled checkbox is selected, high-availability is enabled for web sessions and for stateful session bean (SFSB) checkpointing and potentially passivation. If set to false (default) all web session saving and SFSB checkpointing is disabled for the specified application, web application, or EJB module. If set to true, the specified application or module is enabled for high-availability. Set this option to true only if high availability is configured and enabled at higher levels, such as the server and container levels.This option appears if clusters or standalone server instances aside from the default server instance (`server`) exist.

- Precompile JSPs

  If this option is selected, JavaServer Pages (JSP) files are precompiled. If this option is disabled, JSP files are compiled at runtime when they are first accessed. This option is disabled by default.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages must be installed from the Update Tool or a warning is logged and this option is ignored.

- Force Redeploy

  If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Keep State

  This option controls whether web sessions, SFSB instances, and persistently created EJB timers are retained between redeployments.This option is disabled by default. This option is supported only on the default server instance, named `server` . It is not supported and ignored for any other target.Some changes to an application between redeployments prevent this feature from working properly. For example, do not change the set of instance variables in the SFSB bean class.For web applications, this feature is applicable only if in the `apusic-web-app.xml` file the `persistence-type` attribute of the `session-manager` element is `file` .For stateful session bean instances, the persistence type without high availability is set in the server (the SFSB Persistence Type option) and must be set to `file` , which is the default and recommended value.If any active web session, SFSB instance, or EJB timer fails to be preserved or restored, *none* of these will be available when the redeployment is complete. However, the redeployment continues and a warning is logged.To preserve active state data, Apusic Application Server serializes the data and saves it in memory. To restore the data, the class loader of the newly redeployed application deserializes the data that was previously saved.

- Preserve Application Scoped Resources

  If checked, preserves any application-scoped resources and restores them during redeployment. This option is not checked by default.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Libraries

  A comma-separated list of library JAR files specific to this module or application. The paths may be absolute or relative. A relative path is relative to *domain-dir* `/lib/applibs` . If the path is absolute, the path must be accessible to the domain administration server (DAS), which means it must be under *domain-dir*. The libraries are made available to the application in the order in which they are specified.

- Share Libraries

  If the application requires the use of a shared class library, you can choose the name of the shared class library, but you need to first set the shared library in the shared library module.

- Share Libraries Priority

  Set the loading priority of shared library classes.Application first,classes directory first,share library first.Default application priority.

- Hot Loading

If this option is selected, the application enables hot loading functionality.You need to set hot loading delay time as well.Disabled by default.

- Hot Loading Delay Time

  Detected delayed loading time after application modification,default to 60 seconds.

- Session Storage

  If the application wants to configure session caching,you can choose the name of the session storage,but you need to first set the session manager in the Session Storage module.

- Enable AAS Built-in Components

  Apusic Application Server has built-in components,sometimes it conflicts with the application.Including JPA,JSF,CDI,Bean Validation,JSONP,RESTful,Web Service.It should be decided whether to enable built-in components based on the actual situation.Default enabled JPA,CDI,Web Service.

- Description

  A description of the application.

4.5.1.2.2 ADDITIONAL OPTIONS FOR AN ENTERPRISE APPLICATION

If the application type is Enterprise Application, the following options appear.

- Application Name

  The name of the application.

- Virtual Servers

  The virtual servers associated with this application.The Virtual Servers option appears if only the default server instance, `server`, exists. If clusters or other standalone server instances exist, you can select virtual servers after deployment. Go to the Edit Application page, select the Target tab, and select Manage Virtual Servers for the desired target.

- Deploy Module

  This attribute will be displayed when the enterprise application has web/ejb or other modules; Deployment modules can be selected as needed.Default select all.

- Status

  If this option is selected, the application is enabled. This option is selected by default.

- Delegate

  If this option is selected,the parent class loader for priority,if not,the child class loader for loading.

- Secret Level

The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Availability

  If the Enabled checkbox is selected, high-availability is enabled for web sessions and for stateful session bean (SFSB) checkpointing and potentially passivation. If set to false (default) all web session saving and SFSB checkpointing is disabled for the specified application, web application, or EJB module. If set to true, the specified application or module is enabled for high-availability. Set this option to true only if high availability is configured and enabled at higher levels, such as the server and container levels.This option appears if clusters or standalone server instances aside from the default server instance (`server`) exist.

- Java Web Start

  If this option is selected, Java Web Start access is permitted for an application client module. This option is disabled by default.

- Precompile JSPs

  If this option is selected, JavaServer Pages (JSP) files are precompiled. If this option is disabled, JSP files are compiled at runtime when they are first accessed. This option is disabled by default.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages

must be installed from the Update Tool or a warning is logged and this option is ignored.

- Compatibility

  If checked, uses Apusic Application Server v2 JAR visibility requirements for applications instead of the stricter Java EE 6 requirements implemented in Apusic Application Server v3 releases, including 4.0. This option is not checked by default.The Java EE 6 platform specification imposes stricter requirements than Java EE 5 did on which JAR files can be visible to various modules within an EAR file. In particular, application clients must not have access to EJB JAR files or other JAR files in the EAR file unless references use the standard Java SE mechanisms (extensions, for example) or the Java EE library-directory mechanism. Checking this box removes these Java EE 6 restrictions.

- Force Redeploy

  If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Keep State

  This option controls whether web sessions, SFSB instances, and persistently created EJB timers are retained between redeployments.This option is disabled by default. This option is supported only on the default server instance, named `server`. It is not supported and ignored for any other target.Some changes to an application between redeployments prevent this feature from working properly. For example, do not change the set of instance variables in the SFSB bean class.For web applications, this feature is applicable only if in the `apusic-web-app.xml` file the `persistence-type` attribute of the `session-manager` element is `file`.For stateful session bean instances, the persistence type without high availability is set in the server (the SFSB Persistence Type option) and must be set to `file`, which is the default and recommended value.If any active web session, SFSB instance, or EJB timer fails to be preserved or restored, *none* of these will be available when the redeployment is complete. However, the redeployment continues and a warning is logged.To preserve active state data, Apusic Application Server serializes the data and saves it in memory. To restore the data, the class loader of the newly redeployed application deserializes the data that was previously saved.

- Preserve Application Scoped Resources

  If checked, preserves any application-scoped resources and restores them during redeployment. This option is not checked by default.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Libraries

  A comma-separated list of library JAR files specific to this application. Specifies an absolute or relative path. A relative path is relative to *domain-dir* `/lib/applibs`. If the path is absolute, the path must be accessible to the domain administration server (DAS), which means it must be under *domain-dir*. The libraries are made available to the application in the order in which they are specified.

- Description

  A description of the application.

4.5.1.2.3 ADDITIONAL OPTIONS FOR AN ENTERPRISE BUSINESS APPLICATION

If the application type is Enterprise Business Application, the following options appear.

- Application Name

  The name of the application.

- Virtual Servers

  The virtual servers associated with this application.The Virtual Servers option appears if only the default server instance, `server`, exists. If clusters or other standalone server instances exist, you can select virtual servers after deployment. Go to the Edit Application page, select the Target tab, and select Manage Virtual Servers for the desired target.

- Status

  If this option is selected, the application is enabled. This option is selected by default.

- Secret Level

The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Force Redeploy

If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Deployment Order The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Libraries

  A comma-separated list of library JAR files specific to this application. Specifies an absolute or relative path. A relative path is relative to *domain-dir*`/lib/applibs`. If the path is absolute, the path must be accessible to the domain administration server (DAS), which means it must be under *domain-dir*. The libraries are made available to the application in the order in which they are specified.

- Description

  A description of the application.

4.5.1.2.4 ADDITIONAL OPTIONS FOR AN APPLICATION CLIENT

If the application type is Application Client, the following options appear.

- Application Name

  The name of the application.

- Java Web Start

  If this option is selected, Java Web Start access is permitted for an application client module. This option is disabled by default.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages must be installed from the Update Tool or a warning is logged and this option is ignored.

- Force Redeploy

  If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Description

  A description of the application.

4.5.1.2.5 ADDITIONAL OPTIONS FOR A CONNECTOR MODULE

If the application type is Connector Module, the following options appear.

- Application Name

  The name of the application.

- Status

  If this option is selected, the application is enabled. This option is selected by default.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages must be installed from the Update Tool or a warning is logged and this option is ignored.

- Force Redeploy

  If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Preserve Application Scoped Resources

  If checked, preserves any application-scoped resources and restores them during redeployment. This option is not checked by default.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Description

  A description of the application.

4.5.1.2.6 ADDITIONAL OPTIONS FOR AN EJB JAR

If the application type is EJB Jar, the following options appear.

- Application Name

  The name of the application.

- Status

  If this option is selected, the application is enabled. This option is selected by default.

- Availability

  If the Enabled checkbox is selected, high-availability is enabled for web sessions and for stateful session bean (SFSB) checkpointing and potentially passivation. If set to false (default) all web session saving and SFSB checkpointing is disabled for the specified application, web application, or EJB module. If set to true, the specified application or module is enabled for high-availability. Set this option to true only if high availability is configured and enabled at higher levels, such as the server and container levels.This option appears if clusters or standalone server instances aside from the default server instance (`server`) exist.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages must be installed from the Update Tool or a warning is logged and this option is ignored.

- Compatibility

  If checked, uses Apusic Application Server v2 JAR visibility requirements for applications instead of the stricter Java EE 6 requirements implemented in Apusic Application Server v3 releases, including 4.0. This option is not checked by default.The Java EE 6 platform specification imposes stricter requirements than Java EE 5 did on which JAR files can be visible to various modules within an EAR file. In particular, application clients must not have access to EJB JAR files or other JAR files in the EAR file unless references use the standard Java SE mechanisms (extensions, for example) or the Java EE library-directory mechanism. Checking this box removes these Java EE 6 restrictions.

- Force Redeploy

  If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Keep State

  This option controls whether web sessions, SFSB instances, and persistently created EJB timers are retained between redeployments.This option is disabled by default. This option is supported only on the default server instance, named `server` . It is not supported and ignored for any other target.Some changes to an

application between redeployments prevent this feature from working properly. For example, do not change the set of instance variables in the SFSB bean class.For web applications, this feature is applicable only if in the `apusic-web-app.xml` file the `persistence-type` attribute of the `session-manager` element is `file` .For stateful session bean instances, the persistence type without high availability is set in the server (the SFSB Persistence Type option) and must be set to `file` , which is the default and recommended value.If any active web session, SFSB instance, or EJB timer fails to be preserved or restored, *none* of these will be available when the redeployment is complete. However, the redeployment continues and a warning is logged.To preserve active state data, Apusic Application Server serializes the data and saves it in memory. To restore the data, the class loader of the newly redeployed application deserializes the data that was previously saved.

- Preserve Application Scoped Resources

  If checked, preserves any application-scoped resources and restores them during redeployment. This option is not checked by default.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Libraries

  A comma-separated list of library JAR files specific to this module or application. Specifies an absolute or relative path. A relative path is relative to *domain-dir* `/lib/applibs` . If the path is absolute, the path must be accessible to the domain administration server (DAS), which means it must be under *domain-dir*. The libraries are made available to the application in the order in which they are specified.

- Description

  A description of the application.

4.5.1.2.7 ADDITIONAL OPTIONS FOR APPLICATIONS OF TYPE OTHER

If the application type is Other, the following options appear.

- Application Name

  The name of the application.

- Virtual Servers

  The virtual servers associated with this application.The Virtual Servers option appears if only the default server instance, `server` , exists. If clusters or other standalone server instances exist, you can select virtual servers after deployment. Go to the Edit Application page, select the Target tab, and select Manage Virtual Servers for the desired target.

- Status

  If this option is selected, the application is enabled. This option is selected by default.

- Availability

  If the Enabled checkbox is selected, high-availability is enabled for web sessions and for stateful session bean (SFSB) checkpointing and potentially passivation. If set to false (default) all web session saving and SFSB checkpointing is disabled for the specified application, web application, or EJB module. If set to true, the specified application or module is enabled for high-availability. Set this option to true only if high availability is configured and enabled at higher levels, such as the server and container levels.This option appears if clusters or standalone server instances aside from the default server instance (`server`) exist.

- OSGI Type

  If this option is selected, it specifies a hybrid OSGi/Java-EE module. This option only appears if the Type selected is Other.

- Precompile JSPs

  If this option is selected, JavaServer Pages (JSP) files are precompiled. If this option is disabled, JSP files are compiled at runtime when they are first accessed. This option is disabled by default.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages must be installed from the Update Tool or a warning is logged and this option is ignored.

- Force Redeploy

  If this option is selected, the application is redeployed if it is already deployed. If this option is not selected, an attempt to deploy an application that is already deployed results in an error. This option is disabled by default.

- Keep State

  This option controls whether web sessions, SFSB instances, and persistently created EJB timers are retained between redeployments.This option is disabled by default. This option is supported only on the default server instance, named `server`. It is not supported and ignored for any other target.Some changes to an application between redeployments prevent this feature from working properly. For example, do not change the set of instance variables in the SFSB bean class.For web applications, this feature is applicable only if in the `apusic-web-app.xml` file the `persistence-type` attribute of the `session-manager` element is `file`.For stateful session bean instances, the persistence type without high availability is set in the server (the SFSB Persistence Type option) and must be set to `file`, which is the default and recommended value.If any active web session, SFSB instance, or EJB timer fails to be preserved or restored, *none* of these will be available when the redeployment is complete. However, the redeployment continues and a warning is logged.To preserve active state data, Apusic Application Server serializes the data and saves it in memory. To restore the data, the class loader of the newly redeployed application deserializes the data that was previously saved.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Libraries

  A comma-separated list of library JAR files specific to this module or application. Specifies an absolute or relative path. A relative path is relative to *domain-dir*`/lib/applibs`. If the path is absolute, the path must be accessible to the domain administration server (DAS), which means it must be under *domain-dir*. The libraries are made available to the application in the order in which they are specified.

- Description

  A description of the application.

4.5.1.3 Step3

Specifies the target to which you are deploying.Deploys the component to the server instance or cluster, `server` and is the default value.

4.5.1.4 Step4

Confirm deployment properties.

4.5.2 Edit Application

Click on the application name,enter the application information editing page.Use the Edit Application page to modify an existing application.

The Edit Application page contains some or all of the following options, depending on the application type.

- Name

  Read-only field that displays the name of the application that you are editing.

- Status

  If the Enabled checkbox is selected, the application is enabled. This option is selected by default.

- Context Root

  For a web application, specifies the path to the application. In the URL of the web application, the context root immediately follows the port number (`http://` *host* `:` *port* `/` *context-root* `/...`). The context root must start with a forward slash, for example, `/hello`.

- Availability

  If the Enabled checkbox is selected, high-availability is enabled for web sessions and for stateful session bean (SFSB) checkpointing and potentially passivation. If set to false (default) all web session saving and SFSB checkpointing is disabled for the specified application, web application, or EJB module. If set to true, the specified application or module is enabled for high-availability. Set this option to true only if high availability is configured and enabled at higher levels, such as the server and container levels.This option appears if clusters or standalone server instances aside from the default server instance (`server`) exist.

- Virtual Servers

  The virtual servers associated with this application.The Virtual Servers option appears if only the default server instance, `server`, exists. If clusters or other standalone server instances exist, you can select virtual servers after deployment. Go to the Edit Application page, select the Target tab, and select Manage Virtual Servers for the desired target.

- Java Web Start

  (For some application types) If the Enabled checkbox is selected, the application uses Java Web Start software. Java Web Start provides a browser-independent way to deploy Java applications to run in a dedicated Java Virtual Machine.

- Description

  A short description of the application.

- Location

  The location of the deployed application. This is a read-only field.

- Deployment Order

  The deployment order of the application.Applications with a lower number are loaded first at server startup. An application with a deployment order of 102 is loaded before an application with a deployment order of 110. If a deployment order is not specified at the time an application is deployed, the default deployment order of 100 is assigned. If two applications have the same deployment order, the application that was deployed first is loaded first. Specifying a deployment order is useful if the application has dependencies and must be loaded in a certain order.

- Libraries

  The library JAR files required by the application. This is a read-only field.

- Modules and Components

  Table that displays the names of the application's modules and their engines. For a web service endpoint, you can select View Endpoint to display the Web Service Endpoint Information page. For an application client, you can select Launch to display the Application Client Launch page or Download Client Stubs to download the client stubs.

4.5.3 Module Descriptors

Use the Module Descriptors page to view module descriptors for an application.

The Module Descriptors page contains the following options.

- Deployment Descriptors

  A table displaying the application name, subcomponent name (if applicable), and descriptor file name.

Use the Deployment Descriptor page to view the deployment descriptor for a module.

The Deployment Descriptor page contains the following information.

- Application Name

  The name of the application.

- Module Name

  The name of the module that contains the deployment descriptor.

- Descriptor File Name

  The name of the deployment descriptor file.

4.5.4 Application Client Launch

Use the Application Client Launch page to launch a deployed application client. The client must be a web client that can be run with Java Web Start.

The Launch Application Client page contains the following options.

- Application

The name of the application that contains the application client.

- Module

  The name of the module that contains the application client.

- Links

  The URL at which the client is deployed.

- Arguments

  The arguments to be passed to the application

4.5.5 Application Targets

Use the Application Targets page to view target clusters and standalone server instances on which the application can be enabled.

The Application Targets page contains the following information.

- Target Name

  The name of the cluster or standalone server instance.

- Enabled

  Displays `true` if the application is enabled on the target, or `false` if the application is disabled.

- Virtual Servers

  To associate the application with specific virtual servers on a target, select Manage Virtual Servers for that target.

The Targets table also contains the following options.

- Manage Targets

  Button to manage application targets.

- More Actions

  Drop-down list of the following actions.EnableAction to enable the application on the selected target.DisableAction to disable the application on the selected target.

4.5.5.1 Manage Targets

Use the Manage Targets page to change the target clusters and standalone server instances on which an application can be enabled. The application can be enabled only on targets in the Selected Targets column.

The Manage Targets page contains the following information.

- Available Targets

  The clusters and instances on which the application is not deployed.

- Selected Targets

  The clusters and instances on which the application is deployed.

- Add

  Button to move the selected target from Available Targets to Selected Targets.

- Add All

  Button to move all Available Targets to Selected Targets.

- Remove

  Button to move the selected target from Selected Targets to Available Targets.

- Remove All

Button to move all Selected Targets to Available Targets.

4.5.5.2 Virtual Server Targets

Use the Virtual Server Targets page to change the target virtual servers with which an application or module is associated.

The Virtual Server Targets page contains the following information.

- Available Targets

  The virtual servers with which an application or module is not associated.

- Selected Targets

  The virtual servers with which an application or module is associated.

- Add

  Button to move the selected target from Available Targets to Selected Targets.

- Add All

  Button to move all Available Targets to Selected Targets.

- Remove

  Button to move the selected target from Selected Targets to Available Targets.

- Remove All

  Button to move all Selected Targets to Available Targets.

4.5.6 Redeploy Applications or Module

Click "Redeploy" button,you can redepoly the application or module.Use the Redeploy Applications or Modules page to redeploy a previously deployed application.

The Redeploy Applications or Modules page contains the following options.

- Location

  The location of the new archive for the application that you are redeploying.The following options specify from where the archive is accessible and whether the archive is a file or a directory.Packaged File to Be Uploaded to the ServerThe archive is in a file that resides on or is accessible from the client machine.The client machine is the host on which you are viewing the Administration Console through a browser.Local Packaged File or Directory That Is Accessible From the Apusic Application ServerThe archive is a file that resides on the server machine, or is an unpackaged application in an exploded directory.The server machine is the host that is running the Apusic Application Server domain administration server.

- Application Name

  Read-only field displaying the name of the application that you are redeploying.

- Availability

  If the Enabled checkbox is selected, high-availability is enabled for web sessions and for stateful session bean (SFSB) checkpointing and potentially passivation. If set to false (default) all web session saving and SFSB checkpointing is disabled for the specified application, web application, or EJB module. If set to true, the specified application or module is enabled for high-availability. Set this option to true only if high availability is configured and enabled at higher levels, such as the server and container levels.This option appears if clusters or standalone server instances aside from the default server instance (`server`) exist.

- Precompile JSPs

  If this option is selected, JavaServer Pages (JSP) files are precompiled. If this option is disabled, JSP files are compiled at runtime when they are first accessed. This option is disabled by default.

- Run Verifier

  If this option is selected, deployment descriptors are verified before deployment. If verification fails, deployment is not performed. The verifier examines the structure and content of the deployment descriptors. Verification of large applications is often time-consuming. This option is disabled by default.Verifier packages must be installed from the Update Tool or a warning is logged and this option is ignored.

- Java Web Start

If this option is selected, Java Web Start access is permitted for an application client module. This option is disabled by default.

- Keep State

  This option controls whether web sessions, SFSB instances, and persistently created EJB timers are retained between redeployments.This option is disabled by default. This option is supported only on the default server instance, named `server`. It is not supported and ignored for any other target.Some changes to an application between redeployments prevent this feature from working properly. For example, do not change the set of instance variables in the SFSB bean class.For web applications, this feature is applicable only if in the `apusic-web-app.xml` file the `persistence-type` attribute of the `session-manager` element is `file`.For stateful session bean instances, the persistence type without high availability is set in the server (the SFSB Persistence Type option) and must be set to `file`, which is the default and recommended value.If any active web session, SFSB instance, or EJB timer fails to be preserved or restored, *none* of these will be available when the redeployment is complete. However, the redeployment continues and a warning is logged.To preserve active state data, Apusic Application Server serializes the data and saves it in memory. To restore the data, the class loader of the newly redeployed application deserializes the data that was previously saved.

- Preserve Application Scoped Resources

  If checked, preserves any application-scoped resources and restores them during redeployment. This option is not checked by default.

4.5.7 Deploy Multiple Versions Application

The application server supports deploying multiple versions of the same application and allows multiple versions to be launched simultaneously.

When deploying an application, setting the 'Context Root' to be the same indicates deploying different versions of the same application.

When multiple versions of an application are launched simultaneously, if the browser is already accessing the application, when deploying other versions again, the browser is still accessing the original version; If you close the browser and then reopen it, when you access the application again, you will be accessing the latest version.

If multiple versions of an application are running simultaneously and the browser is not in a state where the application has already been accessed, accessing the application will access the latest deployed version.

If multiple versions of the application are running simultaneously and the "Deployment Order" is set, restarting the instance before accessing the application will prioritize accessing the version with the lower number of "Deployment Order".

Note:The multi version deployment feature of the application is limited to web applications.

4.5.8 Manager Class Library

The Class library management page lists the libraries of the application, which can be added or deleted.

Click 'New' to add a library to the application. If it is a class file with the same name, uploading it will overwrite the file with the original name.

After adding, the application needs to be reloaded for it to take effect.

4.5.9 Additional Update

The application allows incremental updates, which means supplementing and updating the application. When using it, please note to keep the content path of the update file consistent with the application's requirements.

For example, if you need to update the jar files under WEB-INF/lib, you should ensure the directory structure is the same before packaging them into a zip file and uploading.

Enter the "Additional Update" section of the application, click "New," and upload the zip file.

After confirmation, the zip file will be uploaded, and the application will reload and update the corresponding files in the zip in real-time. Reloading the application will make the class files effective.

After uploading, an "addition/" directory will be generated under ${DOMAIN_HOME}/mydomain/, showing the incremental update records.

Note: If the updated file already exists, it will be overwritten. Please use it carefully.

4.5.10 Visualization of Apusic application configuration file

The application allows setting Apusic configuration information, such as class loading strategy, etc., on the management and control platform.

Setting the class loading policy needs to be set when the application is deployed, and cannot be set in the editing page. At this time, the application in domain.xml has an attribute < property name = "delegate" value = "true" > </property >. Indicates that when the application has a class conflict with AAS, public classes will be loaded first.

To set other items of Apusic configuration, you can set them in "Additional Properties" of [Application]. You need to add the prefix "apusic-". If it is set in the editing page, you need to reload the application to take effect. Common configuration items are:

`prefer-app-pkg` For specific classes, applications are loaded preferentially (to avoid causing other exceptions). The values are generally specific package names or class names, and the delimiter is ';'.

`include-pkg` Load the classes in AAS V10 first, with the delimiter ';'.

`preferJars` Set the preferred loading method of the jar (which can be a folder or a jar, separated by a comma or a colon), for example, set the value to `WEB-INF/preferLib,WEB-INF/extra/aas.jar` .

`serfault-chartset` Set the default encoding format.

`userMyfaces` Generally, it is not set. Use it when there is a conflict with the JSF framework MyFaces.

`ignore-resources` Exclude resource scanning (only scan applications and ignore resources in AAS V10).

`ignore-class-prefix` Exclude class scanning and set the list of ignored classes (first set delegate to false, give priority to loading application classes. If the WebAppClassLoader fails to load, no longer delegate to the parent class loader of AAS V10 for loading).

`ignore-parents-files` Exclude the function switch to avoid conflicts caused by the SPI mechanism scanning classes related to AAS V10, and combine with delegate="false".

`ignore-resource-cache` Exclude the cache switch to avoid conflicts caused by the spi mechanism scanning classes related to AAS V10, and combine it with delegate="false".

`ignore-resource-cache-suffix` Exclude the cache by the suffix (file name) to avoid conflicts caused by the spi mechanism scanning classes related to AAS V10. Combine it with delegate="false".

`ignore-resource-cache-patterns` Exclude the cache through regular expressions to avoid conflicts caused by the spi mechanism scanning classes related to AAS V10, and combine with delegate="false".

`alternatedocroot_*` To add a virtual path, see the Application Configuration Virtual Path section for details.

`cookie-samesite` Configure the cross - domain configuration of SameSite for cookies; the values can be None (no restrictions), Strict (strict, completely prohibits third - parties from obtaining cookies. When cross - site, cookies will not be sent under any circumstances; only when the URL of the current web page is the same as the request target will the cookie be carried), Lax (prevents cross - site access, prohibits obtaining cookies in most cases, unless it is a GET request (links, pre - loads, GET forms) navigating to the target URL).

`BuiltinWebserviceEnabled` Whether to enable the AAS built - in webservice component. A value of true means enable, and a value of false means disable.

`BuiltinCDIEnabled` Whether to enable the AAS built - in CDI component in the configuration. A value of true means enabled, and a value of false means disabled.

`BuiltinJPAEnabled` Whether to enable the AAS built - in JPA component in the configuration. A value of true indicates enabling, and a value of false indicates disabling.

`BuiltinJSFEnabled` Configure whether to enable the AAS built - in JSF component. A value of true means enable, and a value of false means disable.

`BuiltinBeanValidationEnabled` Whether to enable the AAS built - in BeanValidation component. A value of true indicates enabling, and a value of false indicates disabling.

`BuiltinJSONPEnabled` Whether to enable the built - in JSONP component of AAS in the configuration. A value of true indicates enabling, and a value of false indicates disabling.

`BuiltinRESTFULEnabled` Whether to enable the built - in RESTFul component of AAS in the configuration. A value of true means enabled, and a value of false means disabled.

For example:

There is a version conflict of Jersey (the customer application also has a different version of Jersey. It loads the Jersey components in AAS V10 through the SPI mechanism, and the version inconsistency causes an exception). The configuration for resolution is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<aas-web-app error-url="">
    <class-loader delegate="false">
```

```
    <property name="ignoreParentsFiles" value="true"></property>
    <property name="ignoreResources" value="META-
INF/services/org.glassfish.jersey.server.spi.ComponentProvider;META-
INF/services/org.glassfish.jersey.servlet.internal.spi.ServletContainerProvider"></property>
    <property name="ignoreClassPrefix" value="org.glassfish.jersey"></property>
    </class-loader>
</aas-web-app>
```

4.5.11 Configure the global apsic-web.xml

When there is a class conflict between the application and the application server, or the application needs to configure properties, you can add `apusic-web.xml` to the WEB_INF. At the same time, you can configure the global `apusic-web.xml`, named `default-apusic-web.xml`, and put it in the `${DOMAIN_HOME}/mydomain/config/` directory. Priority is given to loading `apusic-web.xml` in the application, and then loading `default-apusic-web.xml`.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE aas-web-app PUBLIC "-//Apusic.com//DTD Apusic Application Server 10.1 Servlet 3.0//EN"
"http://apusic.com/dtds/apusic-web-app_3_0-1.dtd">
<aas-web-app error-url="">
    <class-loader delegate="false"/>
</aas-web-app>
```

.

4.5.12 Hot Loading Application

The application allows hot reloading, which means that when jar or class files in the application are updated, the application does not need to be manually reloaded, and the class information will be automatically reloaded.

To enable hot reloading for a specific application:

Enter the General Information page of the application and tick the checkbox.

- Application Hot Reloading

  Checking this box enables the application's hot reloading function.Disabled by default.

- Application Hot Reloading Delay Time:

  The delay time after detecting modifications to the application, default is 60 seconds.

Click "Save".

After updating jar or class files in the application, if modifications are detected, the log will print an update prompt.

Global Enabling Method: Add `-Dapusic.class.reload.enabled=true` to the JVM options; add `-Dapusic.class.reload.delay=60` to indicate that the update will start after 60 seconds upon detecting modifications to the application. Adding or modifying JVM options requires restarting the instance for the changes to take effect.

Note: If the updated file already exists, it will be overwritten. Please use it carefully.

4.5.13 Enables JSP Hot Loading

Applications deployed to AAS turn off JSP hot loading by default, which can be enabled by configuring aas-web.xml or apusic-web.xml.

```
<?xml version="1.0" encoding="UTF-8"?>
<aas-web-app>
  <context-root>/hello</context-root>
  <jsp-config>
    <property name="development" value="true" />
    <property name="modificationTestInterval" value="0" />
  </jsp-config>
</aas-web-app>
```

`development` : JSP hot - load switch, with values of true/false

`modificationTestInterval` : Detect the JSP update interval

4.5.14 Backup Application or Module

Click "Backup"button,you can backup the application or module.After backup, you can view the backup file information on the backup list page.

The backed up files are stored in `${com.apusic.aas.instanceRoot}/backup_app` by default.If you want to specify the directory for storing backup files,go to the Domain module of the control console,and then enter the application configuration page settings.

- Application Backup Directory Backup file storage directory,stored in `${com.apusic.aas.instanceRoot}/backup_app` by default.Supports relative and absolute directories,ensure that the directory has permissions set.If the application is deployed on an instance of another node, a directory with the corresponding name will be created for the server location of that node.

The Application Backup List page contains the following information.

- Application Name

  The name of the application that you have backed up.

- Backup Time

  The time of the application that you have backed up.

- File Size

  The size of the application that you have backed up.

- Target

  The target of the application that you have backed up.

- Restore

  Selected application's name,click the button of "Restore",you can recover the application if you want.If the application is restored, it will forcibly overwrite the current application with the same name.

- Delete

  Selected application's name,click the button of "Delete",you can delete the application if you want.

4.5.15 Recover Application or Module

Use the Recover Application or Module page to manager application or module files which undeployed.

The recover files are stored in `${com.apusic.aas.instanceRoot}/recover` by default.If you want to specify the directory for storing recover files,go to the Domain module of the control console,and then enter the application configuration page settings.

- Application Recover Directory Recover file storage directory,stored in `${com.apusic.aas.instanceRoot}/recover` by default.Supports relative and absolute directories,ensure that the directory has permissions set.If the application is deployed on an instance of another node, a directory with the corresponding name will be created for the server location of that node.

The Recover Application or Module List page contains the following information.

- Instance

  View all current recycle bin applications for this instance, `server` and is the default value.

- Application Name

  The name of the application that you have undeployed.

- Undeployed Time

  The time of the application that you have undeployed.

- File Size

  The size of the application that you have backed up.

- Restore

  Selected application's name,click the button of "Restore",you can recover the application if you want.After restoration, the "Recover Application or Module List" page no longer displays the application file, and the application recover directory synchronously moves the application file; After restoration, the current files of the application will be forcibly overwritten.

- File Size

  The size of the application that you have backed up.

4.5.16 Prevent Application File Tampering

It supports the configuration to protect specified files or directories in business applications with the anti-tampering function. When files or directories protected by the anti-tampering function are overwritten, it can be detected and restored in a timely manner.

Support configuration The files or directories specified in the business application are protected by the tamper-proof function. When the files or directories protected by the tamper-proof function are overwritten, they can be found and restored in time.

Enter the [Application Management] page, click "Tamper" in the application list, and enter the "Edit Tamper" page.

The page can configure the anti-tampering policy for application files and save it. The configuration takes effect immediately.

- Name

The name of the application that you have undeployed.

- Switch

Whether to enable the anti-tampering function.

- WhiteList

Configure the tamper-proof file, click "WhiteListDir", and support the configuration folder, file name, file prefix, and file suffix. For example, .js means to protect all .js files in the application directory; multiple are separated by English commas. When clicking "WhiteListPattern", regular expression input is supported, such as (.) */hello/WEB-INF /(.)*$。 empty means all files are tamper-proof.

- BlackList

Configure exclusion files, that is, the configured files are not protected by the anti-tampering policy. Click "BlackList", and you can configure folders, file names, file prefixes, and file suffixes; separate multiple items with commas.

When clicking "BlackListPattern", input in the form of regular expressions is supported. If it is empty, it means that all files are protected by the anti-tampering policy.

- Interval

Set the time interval for detecting application files.

4.5.17 Support separating different jar dependencies from multiple folders

Support the ear application to load the class libraries in the application root directory and the first - level directory, improve maintainability, avoid class conflicts, and contribute to modular development.

Create a new META-INF/as-application.xml in the application directory, set the compatibility value to was, and the content is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
 <aas-application>
 <compatibility>was</compatibility>
 </aas-application>
```

4.5.18 Ignore the Annotation Scan of Jar

Add the file `${com.apusic.aas.instanceRoot}/config/ignoreJars` to skip the annotation scan of the jar (multiple need to wrap) to speed up deployment.

4.5.19 Destroy the Thread Pool and Timer Created by the Application

The application server supports destroying the timer and thread pool created by the application when the application is stopped to prevent memory leakage. clearReferencesStopThreads can be used to stop Thread and thread pool threads through configuration parameters; clearReferencesStopTimerThreads is used to stop Timer threads; the default value is false. Three configuration methods are supported:

1. Add additional properties to the control deployment application page

   Add in "Additional Properties" in the application.

   ```
   apusic-clear-references-stop-threads, the value is true
   apusic-clear-references-stop-timer-threads, the value is true
   ```

   Note: If the parameter is added at deployment time, the application will be stopped and the generated thread pool or timer will be destroyed. If it is added in the edit page, the application parameters will need to be "reloaded" to take effect; for the generated thread pool or timer, the instance will need to be restarted to shut down.

2. Application configuration META-INF/context.xml file

   Add in context.xml:

   ```
   <Context clearReferencesStopThreads="true" clearReferencesStopTimerThreads="true"></Context>
   ```

3. Server configuration aas-web.xml or apusic-web.xml

   The content of the file is as follows:

   ```
   <?xml version="1.0" encoding="UTF-8"?>
   <!DOCTYPE aas-web-app PUBLIC "-//Apusic.com//DTD Apusic Application Server 10.1 Servlet 3.0//EN"
   "http://apusic.com/dtds/apusic-web-app_3_0-1.dtd">
   <aas-web-app error-url="">
     <class-loader delegate="false"/>
   <property name="clearReferencesStopThreads" value="true">
   <property name="clearReferencesStopTimerThreads " value="true">
   </aas-web-app>
   ```

4.6 Resources

4.6.1 JDBC

Java DataBase Connectivity (JDBC) technology provides connectivity to relational databases. In the Apusic Application Server, you configure database connectivity by adding JDBC resources (data sources) to the domain. A data source is a Java Platform, Enterprise Edition (Java EE platform) standard method of configuring connectivity to a database. Each data source contains a pool of database connections. Applications look up the data source on the Java Naming and Directory Interface (JNDI) API tree or in the local application context and then retrieve a database connection with the getConnection method. Data sources and their connection pools provide connection management processes that help keep your system running.

4.6.1.1 JDBC Connection Pools

A JDBC connection pool contains a group of JDBC connections that are created when the connection pool is registered. Connection pools use a JDBC driver to create physical database connections. Your application borrows a connection from the pool, uses it, then returns it to the pool by closing it. To create a JDBC resource, specify the associated connection pool. Multiple JDBC resources can specify a single connection pool.

A JDBC connection pool is a group of reusable connections for a particular database. Because creating each new physical connection is time-consuming, the server maintains a pool of available connections to increase performance. When an application requests a connection, it obtains one from the pool. When an application closes a connection, the connection is returned to the pool.

When creating a connection pool with the Administration Console, you are defining the aspects of a connection to a specific database.

Use the JDBC Connection Pools page to configure JDBC connection pools.

For each connection pool, the following information is provided.

- Pool Name

  The name of the connection pool.

- Resource Type

The resource type of the connection pool.

- Classname

  The name of the vendor-specific class that implements the `javax.sql.XADataSource`, `java.sql.ConnectionPoolDataSource`, `javax.sql.DataSource`, or `java.sql.Driver` API.

- Description

  A short description of the connection pool.

The Pools table also contains the following options.

- New

  Button to create a new connection pool.

- Delete

  Button to delete one or more selected connection pools.

4.6.1.1.1 NEW JDBC CONNECTION POOL

New JDBC Connection Pool (Step 1 of 2)

Use the New JDBC Connection Pool (Step 1 of 2) page to create a new JDBC connection pool.

The New JDBC Connection Pool (Step 1 of 2) page contains the following options.

- Pool Name

  The name for the connection pool.Naming needs to be unique.

- Resource Type

  Available resource types are `javax.sql.XADataSource` (global transactions), `java.sql.ConnectionPoolDataSource` (local transactions, possible performance improvements), `javax.sql.DataSource` (local transactions only), and `java.sql.Driver`.

- Database Driver Vendor

  The name of the vendor of your database product.

- Introspect

  If this option is enabled, the database driver is introspected based on the Resource Type and Database Driver Vendor, and any discovered datasource or driver class names are displayed on the New JDBC Connection Pool (Step 2 of 2)page when you click the Next button.For introspection to work, the JDBC driver must be available to Apusic Application Server. See To Integrate a JDBC Driver.

New JDBC Connection Pool (Step 2 of 2)

Use the New JDBC Connection Pool (Step 2 of 2) page to create a new JDBC connection pool.

The New JDBC Connection Pool (Step 2 of 2) page contains the following options.

- Datasource Classname

  The vendor-specific class name that implements the `DataSource`, `ConnectionPoolDataSource`, or `XADataSource` API. If you enabled the Introspect option and selected one of the datasource resource types, this field is filled in automatically with the class name associated with the resource type and database vendor you chose. If you selected the `java.sql.Driver` resource type, this field is disabled.

- Driver Classname

  The vendor-supplied JDBC driver class name. If you enabled the Introspect option and selected the `java.sql.Driver` resource type, this field is filled in automatically with the class name associated with the driver and database vendor you chose. If you selected one of the datasource resource types, this field is disabled.

- Secret Level

  The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- URL

  Specifies the URL for this connection pool. Although this is not a standard property, it is commonly used.

- User Name

  Specifies the user name for connecting to the database.

- Password

  Specifies the password for connecting to the database.

- Ping

  If this option is selected, the pool will be pinged automatically during pool creation or reconfiguration to identify and warn of any erroneous attribute values. This option is disabled by default.

- Description

  A short description for the connection pool.

- Initial and Minimum Pool Size

  The minimum number of connections in the pool. This value also determines the number of connections placed in the pool when the pool is first created or when Apusic Application Server starts. The default value is 8.

- Maximum Pool Size

  The maximum number of connections in the pool. The default value is 32.

- Pool Resize Quantity

  When the pool scales up and scales down toward the maximum and minimum pool size respectively, it is resized in batches. This value determines the number of connections in the batch. Making this value too large delays connection creation and recycling; making it too small will be less efficient. The default value is 2.

- Idle Timeout

  The maximum amount of time, in seconds, that a connection can remain idle in the pool. After this time expires, the connection is removed from the pool. The default value is 300.

- Max Wait Time

  The maximum amount of time, in milliseconds, that the application requesting a connection will wait before getting a connection timeout. The default value is 60,000.

- Max Age

  Maximum lifetime of a connection since creation.0 represents not verifying idle connections.

- Non Transactional Connections

  If this option is enabled, non-transactional connections are returned.

- Transaction Isolation

  If a transaction isolation level is specified, connections in this pool operate at the specified level. Otherwise, the connections operate with default isolation levels provided by the JDBC driver. By default, this option is not specified.

- Isolation Level

  If the Guaranteed checkbox is selected, all connections taken from the pool have the same isolation level. For example, if the isolation level for the connection is changed programmatically (with `con.setTransactionIsolation`) when last used, then this mechanism changes the status back to the specified isolation level. This option is enabled by default.This setting is only applicable if a transaction isolation level has been specified.

- Add JDBC Driver Libraries

  Upload the driver file according to the database type and version. The driver file must be uploaded when connecting to the database for the first time. If the database driver file has already been copied to the ${APUSIC_HOME}/domains/mydomain/lib/ext directory, there is no need to upload it again.

- Additional Properties

Additional properties for the JDBC connection pool.

4.6.1.1.2 EDIT CONNECTION POOL

Use the Edit Connection Pool page to edit the general properties of a JDBC connection pool.

The Edit Connection Pool page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Flush

  Button to reinitialize the connections in the connection pool. Clicking this button destroys any existing connections, recreates connections established for the pool, and restores the pool to its initial and minimum pool size.

- Ping

  Button to verify that the connection pool is usable. If an error message appears when you click this button, check to see if the database server has been started.

- Pool Name

  The name of the connection pool. This is a read-only field. You can only specify the pool name when you create a JDBC connection pool.

- Resource Type

  Choices include javax.sql.XADataSource (global transactions), java.sql.ConnectionPoolDataSource (local transactions, possible performance improvements), javax.sql.DataSource (local transactions only), and `java.sql.Driver`.

- Secret Level

  The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Datasource Classname

  The vendor-specific class name that implements the data source.

- Driver Classname

  The vendor-supplied JDBC driver class name. If you selected the java.sql.Driver resource type, this field is usually filled in automatically with the class name associated with the driver and database vendor you chose. If you selected one of the datasource resource types, this field is disabled.

- Ping

  If this option is selected, the pool will be pinged automatically during pool creation or reconfiguration to identify and warn of any erroneous attribute values. This option is disabled by default.

- Deployment Order

  The deployment order of the connection pool.

- Description

  A short description of the connection pool.

- Initial and Minimum Pool Size

  The minimum number of connections in the pool. This value also determines the number of connections placed in the pool when the pool is first created or when Apusic Application Server starts. The default value is 8.

- Maximum Pool Size

  The maximum number of connections in the pool. The default value is 32.

- Pool Resize Quantity

  When the pool scales up and scales down toward the maximum and minimum pool size respectively, it is resized in batches. This value determines the number of connections in the batch. Making this value too large delays connection creation and recycling; making it too small will be less efficient. The default value is 2.

- Idle Timeout

The maximum time, in seconds, that a connection can remain idle in the pool. After this time expires, the connection is removed from the pool. The default value is 300.

- Max Wait Time

  The maximum time, in milliseconds, that an application requesting a connection will wait before getting a connection timeout. The default value is 60,000.

- Max Age

  Maximum lifetime of a connection since creation.0 represents not verifying idle connections.

- Non Transactional Connections

  If this option is enabled, non-transactional connections are returned. By default, this option is not specified.

- Transaction Isolation

  If a transaction isolation level is specified, connections in this pool operate at the specified level. Otherwise, the connections operate with default isolation levels provided by the JDBC driver. By default, this option is not specified.

- Isolation Level

  If the Guaranteed checkbox is selected, then all connections taken from the pool have the same isolation level. For example, if the isolation level for the connection is changed programmatically (with `con.setTransactionIsolation`) when last used, then this mechanism changes the status back to the specified isolation level.This field is only applicable if a transaction isolation level has been specified.

4.6.1.1.3 EDIT CONNECTION POOL ADVANCED ATTRIBUTES

Use the Edit Connection Pool Advanced Attributes page to specify attributes that help diagnose connection leaks and improve ease-of-use.

The Edit Connection Pool Advanced Attributes page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Pool Name

  The name of the connection pool. The name is a read-only field. You can only specify the connection pool name when you create a new JDBC connection pool.

- Statement Timeout

  The length of time, in seconds, after which abnormally long running queries will be terminated. Apusic Application Server will set `QueryTimeout` on the statements created. A value of -1 means that the option is disabled. The default value is -1.

- Statement Cache Size

  The number of statements to be cached using the LRU (Least Recently Used) caching mechanism. A value of 0 disables statement caching. The default value is 0.

- Init SQL

  An SQL string to be executed whenever a connection is created (not reused) in the pool. Execution of this string initializes the state of the connection.

- SQL Trace Listeners

  A comma-separated list of listener implementation class names, which specifies that SQL statements executed by applications will be traced. The listener classes must implement the `org.Apusic Application.api.jdbc.SQLTraceListener` interface. Use of listeners enables easy filtering of log messages for the SQL statements and helps administrators analyze the statements.

- Wrap JDBC Objects

  If this option is enabled, applications will get wrapped JDBC objects for Statement, PreparedStatement, CallableStatement, ResultSet, DatabaseMetaData. This option is enabled by default.

- Pooling

  Enables connection pooling for the pool. This option is enabled by default.

- Validate At Most Once

Specifies the time interval in seconds between successive requests to validate a connection at most once. Setting this attribute to an appropriate value minimizes the number of validation requests by a connection. Default value is 0, which means that the attribute is not enabled.

- Connection Leak Timeout

  The amount of time, in seconds, to trace connection leaks in a connection pool. Use this field to detect potential connection leaks by the application. A connection that is not returned back to the pool by the application within the specified period is assumed to be potentially leaking, and a stack trace of the caller is logged. A value of 0 means that connection leak tracing is disabled. The default value is 0.Use this field along with Leak Reclaim to avoid potential connection leaks from the application.

- Connection Leak Reclaim

  If this option is enabled, leaked connections are restored to the pool after leak connection tracing is complete. This option is disabled by default.

- Statement Leak Timeout

  The amount of time, in seconds, to trace statement leaks in a connection pool. Use this field to detect potential statement leaks by the application. A statement that is not returned back to the pool by the application within the specified period is assumed to be potentially leaking, and a stack trace of the caller is logged. A value of 0 means that statement leak tracing is disabled. The default value is 0.Use this field along with Statement Leak Reclaim to avoid potential connection leaks from the application.

- Statement Leak Reclaim

  If this option is enabled, leaked statements are restored to the pool after leak statement tracing is complete. This option is disabled by default.

- Creation Retry Attempts

  The number of attempts that will be made if there is a failure in creating a new connection. A value of 0 means that no attempts will be made to create the connection again. The default value is 0.

- Retry Interval

  The interval of time, in seconds, between two attempts to create a connection. The default value is 10. This attribute is used only if the value of Creation Retry Attempts is greater than 0.

- Lazy Association

  If this option is enabled, a physical connection is associated with a logical collection only when the connection is used. Connections are disassociated when the transaction is completed and a component method ends, which helps reuse of the physical connections. This option is disabled by default.

- Lazy Connection Enlistment

  If this option is enabled, a resource is enlisted to the transaction only when it is actually used in a method. This option is disabled by default.

- Associate with Thread

  If this option is enabled, an association is made between a connection and a thread such that when the same thread is in need of a connection, it can reuse the connection already associated with that thread, thereby not incurring the overhead of getting a connection from the pool. This option is disabled by default.

- Match Connections

  If this option is enabled, a connection that is selected from the pool should be matched with connections that have certain credentials. This option must be enabled if the connection pool is used by applications that have multiple user credentials. This option can safely be disabled if the administrator knows that the connections in the pool will always be homogeneous and hence a connection picked from the pool does not need to be matched by the resource adapter. This option is disabled by default.

- Max Connection Usage

  The number of times that a connection should be reused by the pool. Once a connection is reused for the specified number of times, it will be closed. By limiting the number of times a connection can be reused, you can avoid statement leaks if an application does not close statements. A value of 0 means that this attribute is disabled. The default value is 0.

- Connection Validation

  If this option is enabled, connection validation is required, allowing the server to reconnect in case of failure.

- Validation Method

Apusic Application Server can validate database connections in the following ways. `auto-commit` and `metadata` Apusic Application Server validates a connection by calling the `con.getAutoCommit` or `con.getMetaData` method.**Note:**Many JDBC drivers cache the results of these calls. As a result, using these calls might not always provide reliable validations. Check with the driver vendor to determine whether these calls are cached or not. `custom-validation` Apusic Application Server uses a user-defined validation mechanism specified by the custom implementation class in the Validation Classname field. `table` The application queries the database table that is specified. The table must exist and be accessible, but it does not require any rows. Do not use an existing table that has a large number of rows, or a table that is already frequently accessed.

- Table Name

  Name of the database table for validation. This field is applicable only if the Validation Method was set to `table`. You can either select the name from the drop-down list or type it.

- Validation Class Name

  The custom validation implementation class name. This field is applicable only if the Validation Method was set to `custom-validation`. The class name provided must be accessible to Apusic Application Server. The specified class must implement the `org.Apusic Application.api.jdbc.ConnectionValidation` interface. You can either select the name from the drop-down list or type it.

- On Any Failure

  If this option is enabled, the Apusic Application Server closes all connections in the pool and reestablishes them when any connection fails. If the option is disabled, individual connections are reestablished only when they are used. This option is disabled by default.

- Allow Non Component Callers

  If this option is enabled, the pool can be used by non-component callers such as servlet filters and third-party persistence managers. Connections obtained by non-component callers are not automatically closed at the end of a transaction by the container. They must be explicitly closed by the caller. This option is disabled by default.

4.6.1.1.4 EDIT CONNECTION POOL PROPERTIES

Use the Edit Connection Pool Properties page to modify additional properties of an existing JDBC connection pool.

- Additional Properties

  Additional properties for the JDBC connection pool.

Most JDBC drivers allow use of standard property lists to specify the user, password, and other resource configuration information. Although properties are optional with respect to the Apusic Application Server, some properties might be necessary for most databases.

Changing JDBC driver properties requires a server restart.

The following standard properties are common to many JDBC vendors. For information about the properties available for your JDBC driver, consult the vendor's documentation.

- `user`

  Specifies the user name for connecting to the database.

- `password`

  Specifies the password for connecting to the database.

- `databaseName`

  Specifies the database for this connection pool.

- `serverName`

  Specifies the database server for this connection pool.

- `port`

  Specifies the port on which the database server listens for requests.

- `networkProtocol`

  Specifies the communication protocol.

- `roleName`

  Specifies the initial SQL role name.

- `datasourceName`

  Specifies an underlying `XADataSource` , or a `ConnectionPoolDataSource` if connection pooling is done.

- `description`

  Specifies a text description.

- `url`

  Specifies the URL for this connection pool. Although this is not a standard property, it is commonly used.

4.6.1.1.5 DELETE JDBC CONNECTION POOL

Select one or more JDBC connection pools, click the "Delete" button, and confirm the deletion to delete the connection pool information. When deleting a JDBC connection pool, all JDBC resources in the connection pool will be deleted synchronously. Please use this operation with caution.

4.6.1.2 JDBC Resources

A Java DataBase Connectivity (JDBC) resource (data source) provides applications with the means of connecting to a database. Typically, the administrator creates a JDBC resource for each database accessed by the applications deployed in a domain; however, more than one JDBC resource can be created for a database.

Applications get a database connection from a connection pool by looking up a data source on the Java Naming and Directory Interface (JNDI) API tree and then requesting a connection. The connection pool associated with the data source provides the connection to the application.

Use the JDBC Resources page to configure JDBC resources.

For each resource, the following information is provided.

- JNDI Name

  A unique name that identifies the JDBC resource.

- Logical JNDI Name

  The logical JNDI name for the resource.

- Enabled

  (This column is displayed if only the default server instance, `server` , exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Connection Pool

  The JDBC connection pool associated with the resource.

- Description

  A short description of the resource.

The Resources table also contains the following options.

- New

  Button to create a new resource.

- Delete

  Button to delete one or more selected resources.

- Enable

Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

4.6.1.2.1 NEW JDBC RESOURCE

Use the New JDBC Resource page to create a new JDBC resource.

The New JDBC Resource page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Pool Name

  The connection pool associated with the JDBC resource.

- Secret Level

  The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Description

  A short description of the JDBC resource.

- Status

  If this option is enabled, the resource is available at runtime. This option is enabled by default.

- Additional Properties

  Additional properties for the JDBC resource. Apusic Application Server does not define any additional properties for JDBC resources.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.1.2.2 EDIT JDBC RESOURCE

Use the Edit JDBC Resource page to modify an existing JDBC resource.

The Edit JDBC Resource page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource.

- Logical JNDI Name

  The logical JNDI name for the resource. It is a Java EE standard name by which the resource can be accessed. The logical JNDI name is a read-only field.

- Pool Name

  The connection pool associated with the JDBC resource.

- Secret Level

  The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Deployment Order

  The deployment order of the JDBC resource.

- Description

  A short description of the JDBC resource.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

  Additional properties for the JDBC resource. Apusic Application Server does not define any additional properties for JDBC resources.

4.6.1.2.3 DELETE JDBC RESOURCE

Select one or more JDBC resources, click the "Delete" button, and upon confirmation, the JDBC resource information will be deleted. Deleting a JDBC resource will affect the usage of the applications that refer to this JDBC resource, so please proceed with caution.

4.6.1.3 JDBC MultiResources

The JDBC MultiResource provides a method for applications to connect to multiple databases. By setting up a multi-resource, a JDBC resource can connect to multiple connection pools simultaneously, achieving load balancing and failover for database requests.

Applications get a database connection from a connection pool by looking up a data source on the Java Naming and Directory Interface (JNDI) API tree and then requesting a connection. The connection pool associated with the data source provides the connection to the application.

Use the JDBC MultiResources page to configure JDBC multi-resources.

For each resource, the following information is provided.

- JNDI Name

  A unique name that identifies the JDBC resource.

- Logical JNDI Name

  The logical JNDI name for the resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Connection Pool

  The JDBC connection pool associated with the resource.

- Description

  A short description of the resource.

The Resources table also contains the following options.

- New

  Button to create a new resource.

- Delete

  Button to delete one or more selected resources.

- Enable

Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

4.6.1.3.1 NEW JDBC MULTIRESOURCE

Use the New JDBC MultiResource page to create a new JDBC multi-resource.

The New JDBC MultiResource page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Secret Level

  The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Algo Type

  This refers to the method of obtaining a connection next time after obtaining a connection using getConnection() of the data source used this time. There are two options:

  - Load balancing: Randomly select a non-current pool from the connection pool to obtain a connection next time.
  - Failover: Each time a connection is obtained, it is always obtained from the currently used pool until there is a problem with the currently used connection pool, at which point connections are obtained from other connection pools as a failover.
  - FailOverIfBusy: Optional. If the number of connections in the currently used connection pool reaches the maximum value, connections will be obtained from other pools. The maximum number of connections is set in the JDBC connection pool.

- Connection Pool Name

  The connection pool associated with the JDBC resource,choose one or more.The resource type of JDBC connection pool needs to be consistent to be used normally in multiple data source resources.

- TestFrequency

  The interval for testing connections, with a default value of 5 seconds.

- SQL Statement

  SQL Statement for Connection Testing. When multiple data source resources return connections, they do not check the availability of the connections before returning them. Instead, they adopt a periodic polling method to detect each connection pool. If a connection pool is found to be unavailable, it will be removed from the rotation when trying to acquire a connection. Due to the periodic polling, it cannot be guaranteed that the obtained connection is definitely available. Therefore, it is necessary for applications to catch SQLException and retry to obtain a new available connection. As a result, the high availability here is not transparent to applications.

- Status

  If this option is enabled, the resource is available at runtime. This option is enabled by default.

- Additional Properties

  Additional properties for the JDBC multi-resource. Apusic Application Server does not define any additional properties for JDBC multi-resources.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

Note:

1. Changing the algorithm type, test SQL statement, or test interval does not require restarting the application for the changes to take effect.
2. However, when making changes to the connection pools within a multi-data source configuration, such as adding, removing, or modifying pools, it is necessary to re-acquire the data source object from JNDI or restart the application.
3. The JNDI name must be unique and cannot be duplicated with the JNDI name of a JDBC resource.

4.6.1.3.2 EDIT JDBC MULTIRESOURCE

Use the Edit JDBC MultiResource page to modify an existing JDBC multi-resource.

The Edit JDBC MultiResource page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource.

- Secret Level

  The secret level of the application. SECRET/CONFIDENTIAL/ TOPSECRET.

- Algo Type

  This refers to the method of obtaining a connection next time after obtaining a connection using getConnection() of the data source used this time. There are two options:

  - Load balancing: Randomly select a non-current pool from the connection pool to obtain a connection next time.
  - Failover: Each time a connection is obtained, it is always obtained from the currently used pool until there is a problem with the currently used connection pool, at which point connections are obtained from other connection pools as a failover.
  - FailOverIfBusy: Optional. If the number of connections in the currently used connection pool reaches the maximum value, connections will be obtained from other pools. The maximum number of connections is set in the JDBC connection pool.

- Connection Pool Name

  The connection pool associated with the JDBC resource,choose one or more.The resource type of JDBC connection pool needs to be consistent to be used normally in multiple data source resources.

- TestFrequency

  The interval for testing connections, with a default value of 5 seconds.

- SQL Statement

  SQL Statement for Connection Testing. When multiple data source resources return connections, they do not check the availability of the connections before returning them. Instead, they adopt a periodic polling method to detect each connection pool. If a connection pool is found to be unavailable, it will be removed from the rotation when trying to acquire a connection. Due to the periodic polling, it cannot be guaranteed that the obtained connection is definitely available. Therefore, it is necessary for applications to catch SQLException and retry to obtain a new available connection. As a result, the high availability here is not transparent to applications.

- Status

  If this option is enabled, the resource is available at runtime. This option is enabled by default.

- Additional Properties

  Additional properties for the JDBC multi-resource. Apusic Application Server does not define any additional properties for JDBC multi-resources.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.1.3.3 DELETE JDBC MULTIRESOURCE

Select one or more JDBC multi-resources, click the "Delete" button, and upon confirmation, the JDBC multi-resource information will be deleted. Deleting a JDBC multi-resource will affect the usage of the applications that refer to this JDBC multi-resource, so please proceed with caution.

4.6.2 JMS Resources

4.6.2.1 JMS Connection Factories

Use the JMS Connection Factories page to configure JMS connection factory resources.

For each resource, the following information is provided.

- JNDI Name

  The JNDI name that identifies the connection factory.

- Logical JNDI Name

  The logical JNDI name for the connection factory.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Resource Type

  The resource type of the connection factory.

- Description

  A description of the connection factory.

The Connection Factories table also contains the following options.

- New

  Button to create a new connection factory.

- Delete

  Button to delete one or more selected connection factories.

- Enable

  Button to enable one or more selected connection factories.

- Disable

  Button to disable one or more selected connection factories.

4.6.2.1.1 NEW JMS CONNECTION FACTORY

Use the New JMS Connection Factory page to create a new JMS connection factory.

The New JMS Connection Factory page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Resource Type

  The type of the connection factory. Available choices are `javax.jms.ConnectionFactory`, `javax.jms.QueueConnectionFactory`, and `javax.jms.TopicConnectionFactory`.

- Description

  A description of the connection factory.

- Status

  The status of the connection factory. The connection factory can be enabled or disabled. This option is enabled by default.

- Initial and Minimum Pool Size

  The minimum number of connections in the pool for the connection factory. The default value is 8.

- Maximum Pool Size

  The maximum number of connections in the pool for the connection factory. The default value is 32.

- Pool Resize Quantity

  The number of connections to be removed when pool idle timeout expires. The default value is 2.

- Idle Timeout

  The maximum time in seconds that a connection can remain idle in the pool. The default value is 300.

- Max Wait Time

  The amount of time the application requesting a connection will wait before getting a connection timeout. The default value is 60000.

- On Any Failure

  If the Close All Connections checkbox is selected, Apusic Application Server will close all connections in the pool and reestablish them if a single connection fails. This option is disabled by default.

- Transaction Support

  The chosen transaction support overrides the transaction support attribute in the resource adapter associated with this connection pool in a downward compatible way. In other words, it can support a lower transaction level than that specified in the resource adapter, or the same transaction level as that specified in resource adapter, but it cannot specify a higher level.The value may be any of the following:XATransactionThe resource can be used for transactions that involve the use of more than one resource within a transaction scope. This value is the default for a JMS connection factory.For example, transactions may involve this resource plus a JDBC resource, a connector resource, or another JMS connection factory resource. This value offers the most flexibility. A resource that is configured as XATransaction will participate in two-phase commit operations.LocalTransactionThe resource can be used either for transactions that involve only one resource within the transaction scope or as the last agent in a distributed transaction that involves more than one XA resource.This value offers significantly better performance. A resource that is configured as LocalTransaction will not be used in two-phase commit operations.NoTransactionThe resource can never participate in transactions. This setting is of limited use in JMS applications.

- Connection Validation

  If the Required checkbox is selected, connections are validated before being given to the application. If a resource's validation fails, it is destroyed, and a new resource is created and returned. This option is disabled by default.

- Additional Properties

  Additional properties for the connection factory.For information on available properties, see Properties Specific to JMS Connection Factories.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.2.1.2 EDIT JMS CONNECTION FACTORY

Use the Edit JMS Connection Factory page to modify the settings for a JMS connection factory.

The Edit JMS Connection Factory page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource

- Logical JNDI Name

The logical JNDI name for the resource. It is a Java EE standard name by which the resource can be accessed. The logical JNDI name is a read-only field.

- Resource Type

  The type of the connection factory. The resource type is a read-only field. You can only specify the resource type when you create a new JMS connection factory.

- Description

  A description of the connection factory.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Initial and Minimum Pool Size

  The minimum number of connections in the pool for the connection factory. The default value is 8.

- Maximum Pool Size

  The maximum number of connections in the pool. The default value is 32.

- Pool Resize Quantity

  The number of connections to be removed when pool idle timeout expires. The default value is 2.

- Idle Timeout

  The maximum time in seconds that a connection can remain idle in the pool. The default value is 300.

- Max Wait Time

  The amount of time the application requesting a connection will wait before getting a connection timeout. The default value is 60000.

- On Any Failure

  If the Close All Connections checkbox is selected, Apusic Application Server will close all connections in the pool and reestablish them if a single connection fails. This option is disabled by default.

- Transaction Support

  The value may be any of the following:XATransactionThe resource can be used for transactions that involve the use of more than one resource within a transaction scope. This value is the default for a JMS connection factory.For example, transactions may involve this resource plus a JDBC resource, a connector resource, or another JMS connection factory resource. This value offers the most flexibility. A resource that is configured as XATransaction will participate in two-phase commit operations.LocalTransactionThe resource can be used either for transactions that involve only one resource within the transaction scope or as the last agent in a distributed transaction that involves more than one XA resource.This value offers significantly better performance. A resource that is configured as LocalTransaction will not be used in two-phase commit operations.NoTransactionThe resource can never participate in transactions. This setting is of limited use in JMS applications.

- Connection Validation

  If the Required checkbox is selected, connections are validated before being given to the application. If a resource's validation fails, it is destroyed, and a new resource is created and returned. This option is disabled by default.

- Additional Properties

  Additional properties for the connection factory.For information on available properties, see Properties Specific to JMS Connection Factories.

4.6.2.1.3 PROPERTIES SPECIFIC TO JMS CONNECTION FACTORIES

The following additional properties are available for a JMS connection factory.

- `ClientId`

  Specifies a client ID for a connection factory that will be used by a durable subscriber.

- `AddressList`

Specifies the names (and, optionally, port numbers) of a message broker instance or instances with which applications will communicate. Each address in the list specifies the host name (and, optionally, host port and connection service) for the connection. For example, the value might be `earth` or `earth:7677`. Specify the port number if the message broker is running on a port other than the default (7676). If the property setting specifies multiple hosts and ports in a clustered environment, the first available host on the list is used unless the `AddressListBehavior` property is set to `RANDOM`.For details, see the *Message Queue Developer's Guide for Java Clients*.The default value is the local host and default port number (7676). The client will attempt a connection to a broker on port 7676 of the local host.

- `UserName`

  The user name for the connection factory.The default value is `guest`.

- `Password`

  The password for the connection factory.The default value is `guest`.

- `ReconnectEnabled`

  If set to `true`, specifies that the client runtime attempts to reconnect to a message server (or the list of addresses in the `AddressList`) when a connection is lost.The default value is `true`.

- `ReconnectAttempts`

  Specifies the number of attempts to connect (or reconnect) for each address in the `AddressList` before the client runtime tries the next address in the list. A value of -1 indicates that the number of reconnect attempts is unlimited (the client runtime attempts to connect to the first address until it succeeds).The default value is 3.

- `ReconnectInterval`

  Specifies the interval in milliseconds between reconnect attempts. This applies for attempts on each address in the `AddressList` and for successive addresses in the list. If the interval is too short, the broker does not have time to recover. If it is too long, the reconnect might represent an unacceptable delay.The default value is 30000.

- `AddressListBehavior`

  Specifies whether connection attempts are in the order of addresses in the `AddressList` attribute (`PRIORITY`) or in a random order (`RANDOM`). `RANDOM` means that the reconnect chooses a random address from the `AddressList`. If many clients are likely to attempt a connection using the same connection factory, this value prevents them from all being connected to the same address. `PRIORITY` means that the reconnect always tries to connect to the first server address in the `AddressList` and uses another address only if the first broker is not available.The default value is `RANDOM`.

- `AddressListIterations`

  Specifies the number of times the client runtime iterates through the `AddressList` in an effort to establish (or reestablish) a connection). A value of -1 indicates that the number of attempts is unlimited.The default value is 3. The maximum value is 2147483647.

4.6.2.1.4 DELETE JMS CONNECTION FACTORIES

Select one or more JMS connection factory click the "Delete" button, and upon confirmation, the JMS connection factory information will be deleted. Deleting a JMS connection factory will affect the usage of the applications that refer to this JMS connection factory, so please proceed with caution.

4.6.2.2 Destination Resources

4.6.2.2.1 JMS DESTINATION RESOURCES

Use the JMS Destination Resources page to configure JMS destination resources.

For each destination resource, the following information is provided.

- JNDI Name

  A unique name that identifies the destination resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

(This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Resource Type

  The resource type of the destination resource.

- Description

  A description of the destination resource.

The Destination Resources table also contains the following options.

- New

  Button to create a new destination resource.

- Delete

  Button to delete one or more selected destination resources.

- Enable

  Button to enable one or more selected destination resources.

- Disable

  Button to disable one or more selected destination resources.

4.6.2.2.2 NEW JMS DESTINATION RESOURCE

Use the New JMS Destination Resource page to create a new JMS destination resource.

The New JMS Destination Resource page contains the following options.

- JNDI Name

  A unique name that identifies the destination resource.It is a recommended practice to use the naming subcontext prefix `jms/` for JMS resources. For example: `jms/Queue`.

- Physical Destination Name

  The name of the physical destination to be associated with this resource.Apusic Application Server creates the physical destination automatically when it is needed and deletes it when you delete the destination resource.

- Resource Type

  The type of the destination resource. Available choices are `javax.jms.Topic` and `javax.jms.Queue`.

- Description

  A description of the destination resource.

- Status

  The status of the destination resource. The destination resource can be enabled or disabled. This option is enabled by default.

- Additional Properties

  Additional properties for the destination resource. Apusic Application Server does not define any additional properties for JMS destination resources.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.2.2.3 EDIT JMS DESTINATION RESOURCE

Use the Edit JMS Destination Resource page to modify the settings for a JMS destination resource.

The Edit JMS Destination Resource page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique name that identifies the destination resource. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new JMS destination resource.

- Physical Destination Name

  The name of the physical destination associated with this resource.Apusic Application Server creates the physical destination automatically when it is needed and deletes it when you delete the destination resource.

- Resource Type

  The type of the destination resource. Available choices are `javax.jms.Topic` and `javax.jms.Queue`.

- Deployment Order

  The deployment order of the destination resource.

- Description

  A description of the destination resource.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

  Additional properties for the destination resource. Apusic Application Server does not define any additional properties for JMS destination resources.

4.6.2.2.4 DELETE JMS DESTINATION RESOURCES

Select one or more JMS destination resource click the "Delete" button, and upon confirmation, the JMS destination resource information will be deleted. Deleting a JMS destination resource will affect the usage of the applications that refer to this JMS destination resource, so please proceed with caution.

4.6.3 ShareLib

The Apusic Application Server supports class sharing, allowing multiple applications to share resources.

The Resources table also contains the following options.

- New

  Button to create a new share library.

- Delete

  Button to delete one or more selected share library resources.

4.6.3.0.1 NEW SHARELIB

Use the New ShareLib page to create a new share lib.If you want to upload the jar files,go to the Edit ShareLib page.

The New ShareLib page contains the following options.

- ShareLibName

  A unique name that identifies the share lib.

4.6.3.0.2 EDIT SHARELIB

Use the Edit ShareLib page to manager share lib resource.

- Upload

Upload class files.

After uploading the class library, a class library folder will be created under ${DOMAIN_HOME}/mydomain/lib/shared.

Note:

When deploying an application, the classes within the application itself are loaded with priority. In cases where there is a conflict between classes in the application and classes in the shared class library, by default, only the classes from the application are loaded.

This ensures that the application's specific dependencies and requirements are met, and that any potential conflicts or inconsistencies caused by duplicate classes are avoided. If there is a need to use a specific version of a class from the shared library, explicit configurations or adjustments to the classloader hierarchy may be necessary.

4.6.4 JavaMail Sessions

Apusic Application Server includes the JavaMail API. JavaMail provides access from Java applications to Internet Message Access Protocol (IMAP) and Simple Mail Transfer Protocol (SMTP) capable mail servers on your network or the Internet.

Use the JavaMail Sessions page to configure JavaMail sessions.

For each JavaMail session, the following information is provided.

- JNDI Name

  A unique name that identifies the JavaMail session.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Description

  A description of the JavaMail session.

The Sessions table also contains the following options.

- New

  Button to create a new JavaMail session.

- Delete

  Button to delete one or more selected JavaMail sessions.

- Enable

  Button to enable one or more selected JavaMail sessions.

- Disable

  Button to disable one or more selected JavaMail sessions.

4.6.4.1 New JavaMail Session

Use the New JavaMail Session page to create a JavaMail session resource.

The New JavaMail Session page contains the following options.

- JNDI Name

  A unique name that identifies the mail session. Use the naming sub-context prefix `mail/` for JavaMail resources. For example: `mail/MySession`. The name must contain only alphanumeric, underscore, dash, or dot characters.

- Mail Host

The host name of the default mail server. The connect methods of the Store and Transport objects use this value if a protocol-specific host property is not supplied. The name must be resolvable to an actual host name.

- Default User

  The user name to provide when connecting to a mail server. The connect methods of the Store and Transport objects use this value if a protocol-specific username property is not supplied. The user name must contain only alphanumeric, underscore, dash, or dot characters.

- Default Sender Address

  The e-mail address of the default user. The format is *username* `@` *host.domain*.

- Description

  A description of this mail session.

- Status

  The status of the mail session. The mail session can be enabled or disabled. The default value is Enabled.

- Store Protocol

  The storage protocol service, which connects to a mail server, retrieves messages, and saves messages in one or more folders. Allowed values are `imap`, `pop3`, `imaps`, and `pop3s`. The default value is `imap`.

- Store Protocol Class

  The service provider implementation class for storage. Allowed values are: `com.sun.mail.imap.IMAPStore``com.sun.mail.pop3.POP3Store``com.sun.mail.imap.IMAPSSLStore``com.sun.mail.pop3.POP3SSLStore` The default value is `com.sun.mail.imap.IMAPStore`.

- Transport Protocol

  The transport protocol service, which sends messages. Allowed values are `smtp` and `smtps`. The default value is `smtp`.

- Transport Protocol Class

  The service provider implementation class for transport. Allowed values are: `com.sun.mail.smtp.SMTPTransport``com.sun.mail.smtp.SMTPSSLTransport` The default value is `com.sun.mail.smtp.SMTPTransport`.

- Debug

  If this option is selected, debugging for this resource is enabled. If the JavaMail log level is set to `FINE` or finer, the debugging output is generated and is included in the system log file.

- Additional Properties

  Additional properties for the JavaMail session. For a list of the available properties, see the JavaMail API documentation (`http://java.sun.com/products/javamail/javadocs/index.html`). Additional properties must begin with `mail-` and must use hyphens, not periods, as separators. For example, `mail-debug` is correct, but `mail.debug` is not. The Apusic Application Server back end converts the hyphens into the periods expected by the JavaMail API.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.4.2 Edit JavaMail Session

Use the Edit JavaMail Session page to edit a JavaMail session.

The Edit JavaMail Session page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

A unique name that identifies the JavaMail session. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new JavaMail session.

- Mail Host

  The host name of the default mail server. The connect methods of the Store and Transport objects use this value if a protocol-specific host property is not supplied. The name must be resolvable to an actual host name.

- Default User

  The user name to provide when connecting to a mail server. The connect methods of the Store and Transport objects use this value if a protocol-specific username property is not supplied. The user name must contain only alphanumeric, underscore, dash, or dot characters.

- Default Sender Address

  The e-mail address of the default user. The format is *username*`@`*host.domain*.

- Deployment Order

  The deployment order of the JavaMail session resource.

- Description

  A description of this mail session.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Store Protocol

  The storage protocol service, which connects to a mail server, retrieves messages, and saves messages in one or more folders. Allowed values are `imap`, `pop3`, `imaps`, and `pop3s`. The default value is `imap`.

- Store Protocol Class

  The service provider implementation class for storage. Allowed values are:`com.sun.mail.imap.IMAPStore``com.sun.mail.pop3.POP3Store``com.sun.mail.imap.IMAPSSLStore``com.sun.mail.pop3.POP3SSLStore`The default value is `com.sun.mail.imap.IMAPStore`.

- Transport Protocol

  The transport protocol service, which sends messages. Allowed values are `smtp` and `smtps`. The default value is `smtp`.

- Transport Protocol Class

  The service provider implementation class for transport. Allowed values are:`com.sun.mail.smtp.SMTPTransport``com.sun.mail.smtp.SMTPSSLTransport`The default value is `com.sun.mail.smtp.SMTPTransport`.

- Debug

  If this option is selected, debugging for this resource is enabled. If the JavaMail log level is set to `FINE` or finer, the debugging output is generated and is included in the system log file.

- Additional Properties

  Additional properties for the JavaMail session. For a list of the available properties, see the JavaMail API documentation (`http://java.sun.com/products/javamail/javadocs/index.html`). Additional properties must begin with `mail-` and must use hyphens, not periods, as separators. For example, `mail-debug` is correct, but `mail.debug` is not. The Apusic Application Server back end converts the hyphens into the periods expected by the JavaMail API.

4.6.5 Resource Adapter Configs

Use the Resource Adapter Configs page to manage connector (resource adapter) configurations.

The resource adapter config provides the configuration information for a resource adapter. With the Administration Console, you create a resource adapter config after the resource adapter is deployed. In this case, the resource adapter is restarted with the new configuration. To create a resource adapter config before you deploy the resource adapter, use the `asadmin create-resource-adapter-config` command.

Before you can create a resource adapter config, you must create a thread pool for the resource adapter to use.

For each resource adapter config, the following information is provided.

- Name

  The name of the resource adapter.

The Configs table also contains the following options.

- New

  Button to create a new resource adapter config.

- Delete

  Button to delete one or more selected resource adapter configs.

4.6.5.1 New Resource Adapter Config

Use the New Resource Adapter Config page to create a new resource adapter configuration.

The New Resource Adapter Config page contains the following options.

- Resource Adapter Name

  The name of the resource adapter you are configuring.

- Thread Pool ID

  The thread pool to be used for the resource adapter configuration.

- Additional Properties

  Additional properties for the resource adapter. Available properties are the names of setter methods of the class referenced by the `resourceadapter-class` element in the `ra.xml` file, which defines the class name of the resource adapter JavaBeans component. Any modifications to properties defined here override the default values present in `ra.xml`.

4.6.5.2 Edit Resource Adapter Config

Use the Edit Resource Adapter Config page to modify the settings of a resource adapter configuration.

The Edit Resource Adapter Config page contains the following options.

- Resource Adapter Name

  The name of the resource adapter you are configuring. The name is a read-only field. You can only specify a resource adapter name when you create a new resource adapter configuration.

- Thread Pool ID

  The thread pool to be used for the resource adapter configuration.

- Deployment Order

  The deployment order of the resource adapter.

- Additional Properties

  Additional properties for the resource adapter. Available properties are the names of setter methods of the class referenced by the `resourceadapter-class` element in the `ra.xml` file, which defines the class name of the resource adapter JavaBeans component. Any modifications to properties defined here override the default values present in `ra.xml`.

4.6.6 Concurrent Resources

Use the Concurrent Resources pages to configure concurrent resources of the following types:

- Context Services
- Managed Thread Factories
- Managed Executor Services
- Managed Scheduled Executor Services

Concurrent resources are managed objects that provide concurrency capabilities to Java EE applications as defined in the JSR 236 specification, Concurrency Utilities for Java EE.

In Apusic Application Server, you configure concurrent resources and make them available for use by application components such as servlets and EJBs. Concurrent resources are accessed through JNDI lookup or resource injection.

4.6.6.1 Context Services

Context services are used to create dynamic proxy objects that capture the context of a container and enable applications to run within that context at a later time. The context of the container is propagated to the thread executing the task.

Use the Context Services page to configure context service resources.

For each context service resource, the following information is provided.

- JNDI Name

  A unique name that identifies the resource.

- Logical JNDI Name

  The logical JNDI name for the resource. See Logical JNDI Names for more information.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Context Information

  The container contexts propagated to threads.

- Description

  A description of the resource.

The Context Services table also contains the following options.

- New

  Button to create a new resource.

- Delete

  Button to delete one or more selected resources.

- Enable

  Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

4.6.6.1.1 NEW CONTEXT SERVICE

Use the New Context Service page to create a context service resource.

The New Context Service page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Description

  A description of this resource.

- Status

  The status of the resource. The resource can be enabled or disabled. The default value is Enabled.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for the resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.6.1.2 EDIT CONTEXT SERVICE

Use the Edit Context Service page to edit a context service resource.

The Edit Context Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource

- Logical JNDI Name

  The logical JNDI name for the resource. It is a Java EE standard name by which the resource can be accessed. The logical JNDI name is a read-only field.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Deployment Order

  The deployment order of this resource. Resources with lower numbers are loaded first at server startup. The default value is 100.

- Description

  A description of this resource.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

4.6.6.2 Managed Executor Services

Managed executor services are used by applications to execute submitted tasks asynchronously. Tasks are executed on threads that are started and managed by the container. The context of the container is propagated to the thread executing the task.

Use the Managed Executor Services page to configure managed executor service resources.

For each managed executor service resource, the following information is provided.

- JNDI Name

  A unique name that identifies the resource.

- Logical JNDI Name

  The logical JNDI name for the resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Context Information

  The container contexts propagated to threads.

- Thread Priority

  The priority assigned to threads.

- Description

  A description of the resource.

The Managed Executor Services table also contains the following options.

- New

  Button to create a new resource.

- Delete

  Button to delete one or more selected resources.

- Enable

  Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

4.6.6.2.1 NEW MANAGED EXECUTOR SERVICE

Use the New Managed Executor Service page to create a managed executor service resource.

The New Managed Executor Service page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Context Information

The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Status

  The status of the resource. The resource can be enabled or disabled. The default value is Enabled.

- Thread Priority

  The priority to assign to created threads. The default value is 5.

- Long-Running Tasks

  Whether the resource should be used for long-running tasks. If enabled, long-running tasks are not reported as stuck. This option is disabled by default.

- Hung After

  The number of seconds tasks can execute before they are considered unresponsive. The default value is 0.

- Description

  A description of this resource.

- Core Size

  The number of threads to keep in a thread pool. The default value is 0.

- Maximum Pool Size

  The maximum number of threads a thread pool can contain. The default value is 2147483647.

- Keep Alive

  The number of seconds threads can remain idle when the number of threads is greater than core size. The default value is 60.

- Thread Lifetime

  The number of seconds threads can remain in a thread pool before being purged, regardless of whether the number of threads is greater than core size or whether the threads are idle. The default value is 0.

- Task Queue Capacity

  The number of submitted tasks that can be stored in the task queue awaiting execution. The default value is 2147483647.

- URI Prefix

  Set the VIP thread URL prefix.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for the resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.6.2.2 EDIT MANAGED EXECUTOR SERVICE

Use the Edit Managed Executor Service page to edit a managed executor service resource.

The Edit Managed Executor Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource

- Logical JNDI Name

  The logical JNDI name for the resource. It is a Java EE standard name by which the resource can be accessed. The logical JNDI name is a read-only field.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Thread Priority

  The priority to assign to created threads. The default value is 5.

- Long-Running Tasks

  Whether the resource should be used for long-running tasks. If enabled, long-running tasks are not reported as stuck. This option is disabled by default.

- Hung After

  The number of seconds tasks can execute before they are considered unresponsive. The default value is 0.

- Deployment Order

  The deployment order of this resource. Resources with lower numbers are loaded first at server startup. The default value is 100.

- Description

  A description of this resource.

- Core Size

  The number of threads to keep in a thread pool. The default value is 0.

- Maximum Pool Size

  The maximum number of threads a thread pool can contain. The default value is 2147483647.

- Keep Alive

  The number of seconds threads can remain idle when the number of threads is greater than core size. The default value is 60.

- Thread Lifetime

  The number of seconds threads can remain in a thread pool before being purged, regardless of whether the number of threads is greater than core size or whether the threads are idle. The default value is 0.

- Task Queue Capacity

  The number of submitted tasks that can be stored in the task queue awaiting execution. The default value is 2147483647.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

4.6.6.3 Managed Scheduled Executor Services

Managed scheduled executor services are used by applications to execute submitted tasks asynchronously at specific times. Tasks are executed on threads that are started and managed by the container. The context of the container is propagated to the thread executing the task.

Use the Managed Scheduled Executor Services page to configure managed scheduled executor service resources.

For each managed scheduled executor service resource, the following information is provided.

- JNDI Name

  A unique name that identifies the resource.

- Logical JNDI Name

  The logical JNDI name for the resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Context Information

  The container contexts propagated to threads.

- Thread Priority

  The priority assigned to threads.

- Description

  A description of the resource.

The Managed Scheduled Executor Services table also contains the following options.

- New

  Button to create a new resource.

- Delete

  Button to delete one or more selected resources.

- Enable

  Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

4.6.6.3.1 NEW MANAGED SCHEDULED EXECUTOR SERVICE

Use the New Managed Scheduled Executor Service page to create a managed scheduled executor service resource.

The New Managed Scheduled Executor Service page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Status

  The status of the resource. The resource can be enabled or disabled. The default value is Enabled.

- Thread Priority

The priority to assign to created threads. The default value is 5.

- Long-Running Tasks

  Whether the resource should be used for long-running tasks. If enabled, long-running tasks are not reported as stuck. This option is disabled by default.

- Hung After

  The number of seconds tasks can execute before they are considered unresponsive. The default value is 0.

- Description

  A description of this resource.

- Core Size

  The number of threads to keep in a thread pool. The default value is 0.

- Keep Alive

  The number of seconds threads can remain idle when the number of threads is greater than core size. The default value is 60.

- Thread Lifetime

  The number of seconds threads can remain in a thread pool before being purged, regardless of whether the number of threads is greater than core size or whether the threads are idle. The default value is 0.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for the resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons

4.6.6.3.2 EDIT MANAGED SCHEDULED EXECUTOR SERVICE

Use the Edit Managed Scheduled Executor Service page to edit a managed scheduled executor service resource.

The Edit Managed Scheduled Executor Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource

- Logical JNDI Name

  The logical JNDI name for the resource. It is a Java EE standard name by which the resource can be accessed. The logical JNDI name is a read-only field.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Thread Priority

  The priority to assign to created threads. The default value is 5.

- Long-Running Tasks

  Whether the resource should be used for long-running tasks. If enabled, long-running tasks are not reported as stuck. This option is disabled by default.

- Hung After

  The number of seconds tasks can execute before they are considered unresponsive. The default value is 0.

- Deployment Order

  The deployment order of this resource. Resources with lower numbers are loaded first at server startup. The default value is 100.

- Description

  A description of this resource.

- Core Size

  The number of threads to keep in a thread pool. The default value is 0.

- Keep Alive

  The number of seconds threads can remain idle when the number of threads is greater than core size. The default value is 60.

- Thread Lifetime

  The number of seconds threads can remain in a thread pool before being purged, regardless of whether the number of threads is greater than core size or whether the threads are idle. The default value is 0.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

4.6.6.4 Managed Thread Factories

Managed thread factories are used by applications to create managed threads on demand. The threads are started and managed by the container. The context of the container is propagated to the thread executing the task.

Use the Managed Thread Factories page to configure managed thread factory resources.

For each managed thread factory resource, the following information is provided.

- JNDI Name

  A unique name that identifies the resource.

- Logical JNDI Name

  The logical JNDI name for the resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Context Information

  The container contexts propagated to threads.

- Thread Priority

  The priority assigned to threads.

- Description

A description of the resource.

The Managed Thread Factories table also contains the following options.

- New

  Button to create a new resource.

- Delete

  Button to delete one or more selected resources.

- Enable

  Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

4.6.6.4.1 NEW MANAGED THREAD FACTORY

Use the New Managed Thread Factory page to create a managed thread factory resource.

The New Managed Thread Factory page contains the following options.

- JNDI Name

  A unique JNDI name that identifies the resource. By convention, the name begins with a resource-type indicator and a slash, such as concurrent/, jdbc/, or jms/.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Status

  The status of the resource. The resource can be enabled or disabled. The default value is Enabled.

- Thread Priority

  The priority to assign to created threads. The default value is 5.

- Description

  A description of this resource.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for the resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.6.4.2 EDIT MANAGED THREAD FACTORY

Use the Edit Managed Thread Factory page to edit a managed thread factory resource.

The Edit Managed Thread Factory page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique name that identifies the resource. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new resource.

- Logical JNDI Name

  The logical JNDI name for the resource. It is a Java EE standard name by which the resource can be accessed. The logical JNDI name is a read-only field.

- Context Information

  The contexts to propagate to threads: Classloader, JNDI, Security, or WorkArea. Context propagation can be enabled or disabled. If enabled, the selected contexts are propagated. If disabled, none of the contexts are propagated, even if they are selected. The default value is Enabled, with all four contexts selected.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Thread Priority

  The priority to assign to created threads. The default value is 5.

- Deployment Order

  The deployment order of this resource. Resources with lower numbers are loaded first at server startup. The default value is 100.

- Description

  A description of this resource.

- Additional Properties

  Additional properties for the resource. Apusic Application Server does not define any additional properties for this resource type. Moreover, this resource type does not currently use any additional properties.

4.6.7 JNDI

4.6.7.1 Custom Resources

A custom resource specifies a custom server-wide resource object factory.

Use the Custom Resources page to configure custom JNDI resources.

For each custom resource, the following information is provided.

- JNDI Name

  A unique name that identifies the custom resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Resource Type

  The type of the custom resource.

- Description

  A description of the custom resource.

The Resources table also contains the following options.

- New

  Button to create a new custom resource.

- Delete

  Button to delete one or more selected custom resources.

- Enable

  Button to enable one or more selected custom resources.

- Disable

  Button to disable one or more selected custom resources.

4.6.7.1.1 NEW CUSTOM RESOURCE

Use the New Custom Resource page to create a custom JNDI resource.

The New Custom Resource page contains the following options.

- JNDI Name

  A unique name that identifies the resource.

- Resource Type

  The fully qualified type definition for the resource.

- Factory Class

  The user-specified name for the factory class. This class implements the `javax.naming.spi.ObjectFactory` interface.

- Description

  A description of this custom resource.

- Status

  The status of the custom resource. The custom resource can be enabled or disabled. The default value is Enabled.

- Additional Properties

  Additional properties for the custom resource. Apusic Application Server does not define any additional properties for custom resources.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.7.1.2 EDIT CUSTOM RESOURCE

Use the Edit Custom Resource page to edit a custom JNDI resource.

The Edit Custom Resource page contains the following options.

- JNDI Name

  A unique name that identifies the resource. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new custom resource.

- Resource Type

  The fully qualified type definition for the resource.

- Factory Class

  The user-specified name for the factory class. This class implements the `javax.naming.spi.ObjectFactory` interface.

- Deployment Order

  The deployment order of the custom resource.

- Description

A description of this custom resource.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

  Additional properties for the custom resource. Apusic Application Server does not define any additional properties for custom resources.

4.6.7.2 External Resources

An external JNDI resource accesses an external JNDI repository, such as an LDAP server.

Use the External Resources page to configure external JNDI resources.

For each external resource, the following information is provided.

- JNDI Name

  A unique name that identifies the external resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Resource Type

  The type of the external resource.

- Description

  A description of the external resource.

The Resources table also contains the following options.

- New

  Button to create a new external resource.

- Delete

  Button to delete one or more selected external resources.

- Enable

  Button to enable one or more selected external resources.

- Disable

  Button to disable one or more selected external resources.

4.6.7.2.1 NEW EXTERNAL RESOURCE

Use the New External Resource page to create an external JNDI resource.

The New External Resource page contains the following options.

- JNDI Name

  A unique name that identifies the resource.

- Resource Type

  The fully qualified type definition for the resource.

- Factory Class

  The user-specified name for the factory class. This class implements the `javax.naming.spi.InitialContextFactory` interface.

- JNDI Lookup

  The JNDI value to look up in the external repository.For example, when creating an external resource to connect to an external repository, to test a bean class, the JNDI Lookup can look like this: `cn=` *testmybean*.

- Description

  A description of this external resource.

- Status

  The status of the external resource. The external resource can be enabled or disabled. The default value is Enabled.

- Additional Properties

  Additional properties for the external resource. Apusic Application Server does not define any additional properties for external resources.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.7.2.2 EDIT EXTERNAL RESOURCE

Use the Edit External Resource page to modify an external JNDI resource.

The Edit External Resource page contains the following options.

- JNDI Name

  A unique name that identifies the resource. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new external resource.

- Resource Type

  The fully qualified type definition for the resource.

- Factory Class

  The user-specified name for the factory class. This class implements the `javax.naming.spi.InitialContextFactory` interface.

- JNDI Lookup

  The JNDI value to look up in the external repository.For example, when creating an external resource to connect to an external repository, to test a bean class, the JNDI Lookup can look like this: `cn=` *testmybean*.

- Deployment Order

  The deployment order of the external resource.

- Description

  A description of this external resource.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

Additional properties for the external resource. Apusic Application Server does not define any additional properties for external resources.

4.6.7.3 Logical JNDI Names

When you create a resource, you specify a unique JNDI name for the resource. Applications use this name to access the resource.

The Java EE standard specifies that certain default resources be made available to applications, and defines specific JNDI names for these default resources. Apusic Application Server makes these names available through the use of logical JNDI names, which map Java EE standard JNDI names to specific Apusic Application Server resources as follows:

- java:comp/DefaultContextService

  This Java EE standard name is mapped to the concurrent/__defaultContextService resource.

- java:comp/DefaultManagedThreadFactory

  This Java EE standard name is mapped to the concurrent/__defaultManagedThreadFactory resource.

- java:comp/DefaultManagedExecutorService

  This Java EE standard name is mapped to the concurrent/__defaultManagedExecutorService resource.

- java:comp/DefaultManagedScheduledExecutorService

  This Java EE standard name is mapped to the concurrent/__defaultManagedScheduledExecutorService resource.

- java:comp/DefaultDataSource

  This Java EE standard name is mapped to the jdbc/__default resource.

- java:comp/DefaultJMSConnectionFactory

  This Java EE standard name is mapped to the jms/__defaultConnectionFactory resource.

4.6.8 Connectors

Use the Connectors page to configure connector resources of the following types:

- Connector Resources
- Connector Connection Pools
- Admin Object Resources
- Work Security Maps

A connector module (also called a resource adapter), is a Java component that enables applications to interact with enterprise information (EIS) software. EIS software includes various types of systems: enterprise resource planning (ERP), mainframe transaction processing, non-relational databases, and messaging systems such as the Java Message Service (JMS), among others.

A connector connection pool is a group of reusable connections for a particular EIS.

A connector resource is a program object that provides an application with a connection to an EIS.

An administered object resource provides specialized functionality for an application.

A work security map maps a principal associated with an incoming work instance to a principal in the security domain of Apusic Application Server.

4.6.8.1 Connector Resources

Use the Connector Resources page to configure connector resources.

A connector resource is a program object that provides an application with a connection to an EIS.

For each resource, the following information is provided.

- JNDI Name

  A unique name that identifies the connector resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Connection Pool

  The connector connection pool associated with the connector resource.

- Description

  A description of the connector resource.

The Resources table also contains the following options.

- New

  Button to create a new connector resource.

- Delete

  Button to delete one or more selected connector resources.

- Enable

  Button to enable one or more selected connector resources.

- Disable

  Button to disable one or more selected connector resources.

4.6.8.1.1 NEW CONNECTOR RESOURCE

Use the New Connector Resource page to create a new connector resource.

The New Connector Resource page contains the following options.

- JNDI Name

  A unique name that identifies the connector resource.

- Pool Name

  The connection pool to which the new connector resource belongs.

- Description

  A description of the connector resource.

- Status

  If this option is selected, the connector resource is enabled. This option is enabled by default.

- Additional Properties

  Additional properties for the connector resource. Available properties depend upon the resource adapter for the connector connection pool. The name-value pairs specified in this table can be used to override the default values for the properties defined by the resource-adapter vendor.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.8.1.2 EDIT CONNECTOR RESOURCE

Use the Edit Connector Resource page to modify the settings for a connector resource.

The Edit Connector Resource page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique name that identifies the connector resource. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new connector resource.

- Pool Name

  The connector connection pool associated with this resource.

- Deployment Order

  The deployment order of the connector resource.

- Description

  A description of the connector resource.

- Status

  Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

  Additional properties for the connector resource. Available properties depend upon the resource adapter for the connector connection pool. The name-value pairs specified in this table can be used to override the default values for the properties defined by the resource-adapter vendor.

4.6.8.2 Connector Connection Pools

Use the Connector Connection Pools page to configure connector connection pools.

A connector connection pool is a group of reusable connections for a particular EIS.

For each resource, the following information is provided.

- Pool Name

  A unique name that identifies the connector connection pool.

- Resource Adapter

  The resource adapter associated with the connector connection pool.

- Connection Definition

  The class name of the connection definition for the connector connection pool.

- Description

  A description of the connector connection pool.

The Pools table also contains the following options.

- New

  Button to create a new connector connection pool.

- Delete

  Button to delete one or more selected connector connection pools.

4.6.8.2.1 NEW CONNECTOR CONNECTION POOL (STEP 1 OF 2)

Use the New Connector Connection Pool (Step 1 of 2) page to create a new connector connection pool.

The New Connector Connection Pool (Step 1 of 2) page contains the following options.

- Name

  A unique name that identifies the connector connection pool.

- Resource Adapter

  The resource adapter associated with the connector connection pool.

- Connection Definition

  The connection definition for the connector connection pool. The choices in the list depend on the resource adapter you selected. The connection definition identifies a resource adapter's `connection-definition` element in the `ra.xml` file.

4.6.8.2.2 NEW CONNECTOR CONNECTION POOL (STEP 2 OF 2)

Use the New Connector Connection Pool (Step 2 of 2) page to create a new connector connection pool.

The New Connector Connection Pool (Step 2 of 2) page contains the following options.

- Ping

  If this option is enabled, the connection pool is pinged during creation or reconfiguration to identify and warn of any erroneous values for its attributes. This option is disabled by default.

- Description

  A description of the connector connection pool.

- Initial and Minimum Pool Size

  The minimum number of connections in the connector connection pool. The default value is 8.

- Maximum Pool Size

  The maximum number of connections in the connector connection pool. The default value is 32.

- Pool Resize Quantity

  The number of connections to be removed when pool idle timeout expires. The default value is 2.

- Idle Timeout

  The maximum time in seconds that a connection can remain idle in the pool. The default value is 300.

- Max Wait Time

  The amount of time the application requesting a connection will wait before getting a connection timeout. The default value is 60000.

- Connection Validation

  If the Required checkbox is selected, connections are validated before being given to the application. If a resource's validation fails, it is destroyed, and a new resource is created and returned. This option is disabled by default.

- On Any Failure

  If the Close All Connections checkbox is selected, Apusic Application Server will close all connections in the pool and reestablish them if a single connection fails. This option is disabled by default.

- Transaction Support

  The chosen transaction support overrides the transaction support attribute in the resource adapter associated with this connection pool in a downward compatible way. In other words, it can support a lower transaction level than that specified in the resource adapter, or the same transaction level as that specified in resource adapter, but it cannot specify a higher level.The value may be any of the following:XATransactionThe resource adapter supports resource manager local and JTA transactions by implementing the LocalTransaction and XAResource interfaces. XA transactions are controlled and coordinated by a transaction manager external to a resource manager.LocalTransactionThe resource adapter supports local transactions by implementing the LocalTransaction interface. Local transactions are managed internal to a resource manager and involve no external transaction managers.NoTransactionThe resource adapter does not support resource manager local or JTA transactions and does not implement the XAResource or LocalTransaction interfaces.

- Additional Properties

Additional properties for the connector connection pool. Available properties depend upon the resource adapter for the connector connection pool. The name-value pairs specified in this table can be used to override the default values for the properties defined by the resource-adapter vendor.If you specify a name but not a value for a property, it is removed from the table when you finish creating the connection pool. If you do not want to set a value for a property now, but you want the property to remain in the table so that the value can be set later, type a set of empty parentheses in the Value field: `()`

4.6.8.2.3 EDIT CONNECTOR CONNECTION POOL

Use the Edit Connector Connection Pool page to modify the settings for a connector connection pool.

The Edit Connector Connection Pool page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Flush

  Button to reinitialize the connections in the connection pool. Clicking this button destroys any existing connections, recreates connections established for the pool, and restores the pool to its initial and minimum pool size.

- Ping

  Button to verify that the connection pool is usable.

- Pool Name

  A unique name that identifies the connector connection pool. This is a read-only field. You can only specify the pool name when you create a new connector connection pool.

- Resource Adapter

  The resource adapter associated with the connector connection pool.

- Connection Definition

  The connection definition for the connector connection pool. The choices in the list depend on the resource adapter you selected. The connection definition identifies a resource adapter's `connection-definition` element in the `ra.xml` file.

- Ping

  If this option is enabled, the connection pool is pinged during creation or reconfiguration to identify and warn of any erroneous values for its attributes. This option is disabled by default.

- Deployment Order

  The deployment order of the connector connection pool.

- Description

  A description of the connector connection pool.

- Initial and Minimum Pool Size

  The minimum number of connections in the connector connection pool. The default value is 8.

- Maximum Pool Size

  The maximum number of connections in the connector connection pool. The default value is 32.

- Pool Resize Quantity

  The number of connections to be removed when pool idle timeout expires. The default value is 2.

- Idle Timeout

  The maximum time in seconds that a connection can remain idle in the pool. The default value is 300.

- Max Wait Time

  The amount of time the application requesting a connection will wait before getting a connection timeout. The default value is 60000.

- Connection Validation

  If the Required checkbox is selected, connections are validated before being given to the application. If a resource's validation fails, it is destroyed, and a new resource is created and returned. This option is disabled by default.

- On Any Failure

  If the Close All Connections checkbox is selected, Apusic Application Server will close all connections in the pool and reestablish them if a single connection fails. This option is disabled by default.

- Transaction Support

  The chosen transaction support overrides the transaction support attribute in the resource adapter associated with this connection pool in a downward compatible way. In other words, it can support a lower transaction level than that specified in the resource adapter, or the same transaction level as that specified in resource adapter, but it cannot specify a higher level.The value may be any of the following:XATransactionThe resource adapter supports resource manager local and JTA transactions by implementing the LocalTransaction and XAResource interfaces. XA transactions are controlled and coordinated by a transaction manager external to a resource manager.LocalTransactionThe resource adapter supports local transactions by implementing the LocalTransaction interface. Local transactions are managed internal to a resource manager and involve no external transaction managers.NoTransactionThe resource adapter does not support resource manager local or JTA transactions and does not implement the XAResource or LocalTransaction interfaces.

4.6.8.2.4 EDIT CONNECTOR CONNECTION POOL ADVANCED ATTRIBUTES

Use the Edit Connector Connection Pool Advanced Attributes page to modify the advanced settings for a connector connection pool.

The Edit Connector Connection Pool Advanced Attributes page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Validate At Most Once

  Specifies the time interval in seconds between successive requests to validate a connection at most once. Default value is 0, which means the attribute is not enabled.

- Leak Timeout

  The amount of time, in seconds, to trace connection leaks in a connection pool. A value of 0 indicates that connection leak tracing is disabled. The default value is 0.

- Leak Reclaim

  If this option is selected, leaked connections will be restored to the pool after leak connection tracing is complete. This option is disabled by default.

- Creation Retry Attempts

  The number of attempts that will be made if there is a failure in creating a new connection. A value of 0 indicates that no attempts will be made to create the connection again. The default value is 0.

- Retry Interval

  The interval, in seconds, between two attempts to create a connection. This attribute is used only if the value of Creation Retry Attempts is greater than 0. The default value is 10 seconds.

- Pooling

  Enables connection pooling for the pool. This option is enabled by default.

- Lazy Association

  If this option is selected, connections will be lazily associated when an operation is performed on them. The connections are disassociated when the transaction is completed and a component method ends, which helps reuse of the physical connections. This option is disabled by default.

- Lazy Connection Enlistment

  If this option is selected, a resource is enlisted in the transaction only when it is actually used in a method. This option is disabled by default.

- Associate with Thread

If this option is selected, a connection is associated with a thread. When the same thread is in need of a connection, it can reuse the connection already associated with it, thereby not incurring the overhead of getting a connection from the pool. This option is disabled by default.

- Match Connections

  If this option is selected, connection matching is enabled for the pool. This option can be disabled if you know that the connections in the pool will always be homogeneous and hence a connection picked from the pool need not be matched by the resource adapter. This option is enabled by default.

- Match Connection Usage

  The number of times a connection should be reused by the pool. Once a connection is reused for the specified number of times, it will be closed. A value of 0 indicates that this attribute is not enabled. The default value is 0.

4.6.8.2.5 EDIT CONNECTOR CONNECTION POOL PROPERTIES

Use the Edit Connector Connection Pool Properties page to modify the properties for a connector connection pool.

The Edit Connector Connection Pool Properties page contains the following options.

- Additional Properties

  Additional properties for the connector connection pool. Available properties depend upon the resource adapter for the connector connection pool. The name-value pairs specified in this table can be used to override the default values for the properties defined by the resource-adapter vendor.If you do not specify a value for a property, it is removed from the table when you finish editing the connection pool. If you do not want to set a value for a property now, but you want the property to remain in the table so that the value can be set later, type a set of empty parentheses in the Value field: `()`

4.6.8.2.6 EDIT CONNECTOR CONNECTION POOL SECURITY MAP

Use the Edit Connector Connection Pool Security Map page to edit a security map for a connector connection pool.

The Edit Connector Connection Pool Security Map page contains the following options.

- Connection Pool Name

  The name of the connector connection pool to which this security map belongs. The pool name is a read-only field.

- Security Map Name

  The name of the security map. The name is a read-only field. You can only specify a security map name when you create a new security map.

- Security Map Name

  The name of the security map.

- User Groups

  If this radio button is selected, the text field must contain the names of the user groups, separated by commas.

- Principals

  If this radio button is selected, the text field must contain the names of the principals, separated by commas.

- Username

  The backend principal username required by the EIS.

- Password

  The backend principal password required by the EIS.

4.6.8.2.7 NEW CONNECTOR CONNECTION POOL SECURITY MAP

Use the New Connector Connection Pool Security Map page to create a new security map for a connector connection pool.

The New Connector Connection Pool Security Map page contains the following options.

- Connection Pool Name

  The name of the connector connection pool to which this security map belongs. The pool name is a read-only field.

- Security Map Name

The name of the security map.

- User Groups

  If this radio button is selected, the text field must contain the names of the user groups, separated by commas.

- Principals

  If this radio button is selected, the text field must contain the names of the principals, separated by commas.

- Username

  The backend principal username required by the EIS.

- Password

  The backend principal password required by the EIS.

4.6.8.3 Admin Object Resources

Use the Admin Object Resources page to configure administered object resources.

An administered object resource provides specialized functionality that is defined by the resource adapter for the deployed connector module.

For each administered object resource, the following information is provided.

- JNDI Name

  A unique name that identifies the administered object resource.

- Enabled

  (This column is displayed if only the default server instance, `server`, exists in the domain.)A check mark if the resource is enabled, or an X if the resource is disabled.

- Status

  (This column is displayed if any clusters or standalone instances have been created in the domain.)The number of clusters and standalone instances to which the resource is targeted and how many of these targets the resource is enabled on. For example, "Enabled on 2 of 4 Target(s)" means that the resource is targeted to four clusters and standalone instances and that it is enabled on two of these four targets.

- Resource Type

  The Java type of the administered object resource.

- Description

  A description of the administered object resource.

The Resources table also contains the following options.

- New

  Button to create a new administered object resource.

- Delete

  Button to delete one or more selected administered object resources.

- Enable

  Button to enable one or more selected administered object resources.

- Disable

  Button to disable one or more selected administered object resources.

4.6.8.3.1 NEW ADMIN OBJECT RESOURCE

Use the New Admin Object Resource page to create a new administered object resource.

The New Admin Object Resource page contains the following options.

- JNDI Name

  A unique name that identifies the administered object resource.

- Resource Adapter

  The resource adapter for the administered object resource.

- Resource Type

  The Java type for the administered object resource (or example, `javax.jms.Topic`).

- Class Name

  The implementation class name associated with the resource type.

- Description

  A description of the administered object resource.

- Status

  If this option is selected, the administered object resource is enabled. This option is enabled by default.

- Additional Properties

  Additional properties for the administered object resource. Available properties depend upon the resource adapter for the connector connection pool. The name-value pairs specified in this table can be used to override the default values for the properties defined by the resource-adapter vendor.

- Targets

  (This option is displayed if any clusters or standalone instances have been created in the domain.)Clusters and standalone instances for resource. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.

4.6.8.3.2 EDIT ADMIN OBJECT RESOURCE

Use the Edit Admin Object Resource page to modify the settings for an administered object resource.

The Edit Admin Object Resource page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- JNDI Name

  A unique name that identifies the administered object resource. The JNDI name is a read-only field. You can only specify the JNDI name when you create a new administered object resource.

- Resource Adapter

  The resource adapter for the administered object resource.

- Resource Type

  The Java type for the administered object resource (or example, `javax.jms.Topic`).

- Class Name

  The implementation class name associated with the resource type.

- Deployment Order

  The deployment order of the administered object resource.

- Description

  A description of the administered object resource.

- Status

Whether the resource is available at runtime.If only the default server instance, `server`, exists in the domain, the resource's runtime availability can be enabled or disabled. Runtime availability is enabled by default.If any clusters or standalone instances have been created in the domain, use the Target tab to control runtime availability of the resource on its targets.

- Additional Properties

  Additional properties for the administered object resource. Available properties depend upon the resource adapter for the connector connection pool. The name-value pairs specified in this table can be used to override the default values for the properties defined by the resource-adapter vendor.

4.6.8.4 Work Security Maps

Use the Work Security Maps page to configure work security maps for resource adapters.

A connector work security map maps the caller identity of the work submitted by a resource adapter EIS principal or EIS user group to a suitable principal or user group in the Apusic Application Server security domain. One or more work security maps may be associated with a resource adapter. The connector work security map configuration supports the use of the wild card asterisk (`*`) to indicate all users or all user groups.

A work security map is different from a security map for a connector connection pool, which maps the principal received during servlet or EJB authentication to the credentials accepted by the EIS. See Connector Connection Pool Security Maps for information about security maps for connector connection pools.

The enterprise information system (EIS) is any system that holds the data of an organization. It can be a mainframe, a messaging system, a database system, or an application.

For each work security map, the following information is provided.

- Name

  The name of the work security map.

The Work Security Maps table also contains the following options.

- New

  Button to create a new work security map.

- Delete

  Button to delete one or more selected work security maps.

4.6.8.4.1 NEW WORK SECURITY MAP

Use the New Work Security Map page to create a new work security map.

The New Work Security Map contains the following options.

- Security Map Name

  The name of the work security map.

- Group

  If the Group Mapping radio button is selected, the text field must specify a map of each back-end EIS user group to the corresponding user group defined in the Apusic Application Server security domain. Use a comma-separated list to specify more than one mapping. For example: `eis-group-1=server-group-1,eis-group-2=server-group-2`

- Principal

  If the Principal Mapping radio button is selected, the text field must specify a map of each back-end EIS principal to the corresponding principal defined in the Apusic Application Server security domain. Use a comma-separated list to specify more than one mapping. For example: `eis-principal-1=server-principal-1,eis-principal-2=server-principal-2`

- Resource Adapter

  The resource adapter for the EIS.

- Description

  A description of the work security map.

- Status

  If this option is selected, the work security map is enabled. This option is enabled by default.

4.6.8.4.2 EDIT WORK SECURITY MAP

Use the Edit Work Security Map page to modify the settings for a work security map.

The Edit Work Security Map page contains the following options.

- Security Map Name

  The name of the work security map. The name is a read-only field. You can only specify a Security Map Name when you create a new work security map.

- Group

  If the Group radio button is selected, the text field must specify a map of each back-end EIS user group to the corresponding user group defined in the Apusic Application Server security domain. Use a comma-separated list to specify more than one mapping. For example: `eis-group-1=server-group-1,eis-group-2=server-group-2`

- Principal

  If the Principal radio button is selected, the text field must specify a map of each back-end EIS principal to the corresponding principal defined in the Apusic Application Server security domain. Use a comma-separated list to specify more than one mapping. For example: `eis-principal-1=server-principal-1,eis-principal-2=server-principal-2`

- Resource Adapter

  The resource adapter for the EIS.

- Deployment Order

  The deployment order of the work security map.

- Description

  A description of the work security map.

- Status

  If this option is selected, the work security map is enabled. This option is enabled by default.

4.6.9 OSGI Bundle Repository

This module is mainly used to manage OSGI Bundle. It is required to deploy module applications using Java EE and OSGi technology. After uploading the OSGI Bundle resource, it is saved to `{APUSIC_HOME}/domains/mydomain/lib/bundle-repository/` by default.

For each resource, the following information is provided.

- Bundle Name

  The name of the bundle.

The OSGI Bundle Repository table also contains the following options.

- New

  Button to create a new OSGI Bundle repository.

- Delete

  Button to delete one or more selected OSGI Bundle repository.If the resource has already been referenced, it needs to be unlinked before it can be deleted.

4.6.9.0.1 NEW OSGI BUNDLE REPOSITORY

Use the New OSGI Bundle Repository page to create a new OSGI Bundle.

The New OSGI Bundle Repository contains the following options.

- Location

The location of the archive for the bundle that you are adding.The type of the file is `.jar` or `.cba`.

The following options specify from where the archive is accessible and whether the archive is a file or a directory.

- Packaged File to Be Uploaded to the Server

The archive is in a file that resides on or is accessible from the client machine.

The client machine is the host on which you are viewing the Administration Console through a browser.

- Local Packaged File or Directory That Is Accessible From the Apusic Application Server

The archive is a file that resides on the server machine, or is an unpackaged application in an exploded directory.

The server machine is the host that is running the Apusic Application Server domain administration server.

After selecting the file, you need to click the "Upload" button.

After adding the Bundle file, you can deploy the related application. You can view the module information in the application's [Module Information]. When there is a new version of the module, you can enable the new version. You can also "enable" or "disable" the extension module. After setting up, click "Save" to take effect.

4.7 Nodes

In the Apusic Application Server, as a system administrator, you can configure the nodes to ensure that the application server can run correctly on the corresponding host.

A node represents a host on which the Apusic Application Server software is installed. A node must exist for every host on which Apusic Application Server instances reside.

In the management console, click the navigation bar to enter the "Nodes" module.

The Nodes page displays a list of nodes in the domain. For each node, the following information is displayed:

- Name

  The name that was assigned to the node when the node was created. Clicking the name opens the Edit Node page for the node.

- Node Host

  The name of the host that the node represents.

- Type

  The type of the node, which is one of the following types:CONFIGThe node does not support remote communication.SSHThe node supports communication over secure shell (SSH).

- Instances

  A list of the instances that reside on the node with an indication of whether each instance is running:If the instance is running, a check mark (✓) is displayed.If the instance is stopped, an exclamation point (`!`) is displayed.Clicking the name of an instance opens the General Information page for the instance.

- Action

  Nodes of type SSH only: A link to the Ping action for the node. Clicking the Ping link tests whether the node is reachable over the communication protocol that the node supports.No actions are possible for CONFIG nodes.

The Nodes table also contains the following options.

- New

  Button to create a node.

- Delete

  Button to delete one or more selected nodes.

- Delete and Uninstall

  Button to delete one or more selected nodes and uninstall Apusic Application Server software from the host that each deleted node represents.

Note:

The Apusic Application Server software is not uninstalled from any host that a node of type CONFIG represents. For nodes of type CONFIG, the action of this button is identical to the action of the Delete button.

4.7.1 New Node

Use the New Node page to create a Apusic Application Server node. A node represents a host on which the Apusic Application Server software is installed. A node must exist for every host on which Apusic Application Server instances reside.

The options that the New Node page contains depend on the type of the node.

4.7.1.1 Options for All Types of Nodes

The following options are available for all types of nodes:

- Name

  The name of the node.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase lettersNumbersHyphenPeriodUnderscoreThe name must be unique in the domain and must not be the name of another node, a cluster, a named configuration, or a Apusic Application Server instance.The name must not be `domain`, `server`, or any other keyword that is reserved by Apusic Application Server.

- Type

  Drop-down list to specify the type of the node, which is one of the following types:CONFIGThe node does not support remote communication. If CONFIG, is selected, no options for SSH configuration is displayed and the Installation Directory field is cleared.SSHThe node supports communication over secure shell (SSH).

- Node Host

  The name of the host that the node represents. Whether this field is mandatory depends on the type of the node:If the type of the node is CONFIG, the name of the host may be omitted. The name of the host can be determined when instances that reside on the node are created.If the type of the node is SSH, the name of the host must be specified. Otherwise, an error occurs.

- Node Directory

  The path to the directory that is to contain Apusic Application Server instances that are created on the node. If a relative path is specified, the path is relative to the *as-install* directory.The default depends on the type of the node:If the type of the node is CONFIG, no directory for instances is specified for the node in the domain administration server (DAS) configuration. The node directory can be determined when instances that reside on the node are created.If the type of the node is SSH, the default is *as-install*`/nodes`, where *as-install* is the base installation directory of the Apusic Application Server software on the host.

- Installation Directory

  The full path to the *parent* of the base installation directory of the Apusic Application Server software on the host, for example, `/export/Apusic Application3`.The default depends on the type of the node:If the type of the node is CONFIG, no installation directory is specified for the node in the DAS configuration. The installation directory can be determined when instances that reside on the node are created.If the type of the node is SSH, the default is the parent of the default base installation directory of the Apusic Application Server software for the DAS. This default is useful only if Apusic Application Server is installed in the same location on all hosts.

4.7.1.2 Additional Options for SSH Nodes

The following options are available only if the Type field is set to SSH:

- Install Apusic Application Server

  If the Enabled option is selected, the DAS will copy Apusic Application Server software from the DAS host to the node host. To copy Apusic Application Server software to the host, the DAS must be able to contact the host through SSH.

- Node Agent Port

  Specify the node proxy port,used to manage agents such as load balancers and cache cluster instances; default 1099, the port needs to be open to the outside world. During the creation of a node of type "SSH", the node agent program nodemanager is installed by default. If the node agent program installation is unsuccessful or there is a conflict with the node agent port, it will not affect the installation and creation of the node.

- Force

  If the Enabled option is selected, the node is created in the DAS configuration even if validation of the node's parameters fails. To validate a node's parameters, the DAS must be able to contact the node's host through SSH.

- SSH Port

  The port to use for SSH connections to this node's host. The default is 22. If the Node Host option is set to `localhost-`*domain*, the SSH port option is ignored.

- SSH User Name

  The user on this node's host that is to run the process for connecting to the host through SSH. The default is the user that is running the DAS process. To ensure that the DAS can read this user's SSH private key file, specify the user that is running the DAS process. If the Node Host option is set to `localhost-`*domain*, the User Name option is ignored.

- SSH User Authentication

  Drop-down list to specify how the SSH user is authenticated when logging in to this node's host. The possible methods for authentication are as follows:Key FileThe SSH user is authenticated through an SSH private key file. If you select this option, specify the absolute path to the file in the Key File field.PasswordThe SSH user is authenticated through a password that is stored in the DAS configuration.**Caution:**This password is stored in clear text in the DAS configuration. For enhanced security, use a key file or a password alias.If you select this option, the SSH User Password field is activated to enable you to provide the password.Password AliasThe SSH user is authenticated through a password alias. If you select this option, select the password alias from the Password Alias drop-down list.

- Setup SSH

  If the Enabled option is selected, Apusic Application Server sets up an SSH key on the node host. Apusic Application Server sets up the SSH key by copying existing key files from the DAS host or, if required, by generating the files before copying the files. This option is available only if the SSH User Authentication option is set to Key File.

- Generate Key

  If the Enabled option is selected, Apusic Application Server generates the SSH key files before copying the files to the node host if the files do not exist. If the SSH key files exist, Apusic Application Server does not generate the files, even if the Enabled option is selected.**Note:**If the SSH key files do not exist, the Enabled option must be selected. Otherwise, an error occurs when you attempt to save your changes.

- SSH User Password

  The password that the SSH user will use when logging in to this node's host.

- Password Alias

  Drop-down list of password aliases that are defined for this domain.

- Public Key File

  The absolute path to the SSH public key file for user that the SSH User Name option specifies. The content of the public key file is appended to the user's `.ssh/authorized_keys` file on the node host. If the `.ssh/authorized_keys` file does not exist on the host, Apusic Application Server creates the file.

- Key File

  The absolute path to the SSH private key file for user that the SSH User Name option specifies. This file is used for authentication to the `sshd` daemon on the node's host.**Note:**Apusic Application Server also supports password authentication through the `AS_ADMIN_SSHPASSWORD` entry in the password file. The password file is specified in the `--paswordfile` option of the `asadmin` utility.If the SSH private key file is protected by a passphrase, the password file must contain the `AS_ADMIN_SSHKEYPASSPHRASE` entry.The path to the key file must be reachable by the DAS and the key file must be readable by the DAS.The default is a key file in the user's `.ssh` directory. If multiple key files are found, Apusic Application Server uses the following order of preference: `id_rsa``id_dsa``identity`

4.7.2 Edit Node

Use the Edit Node page to update the configuration data of a Apusic Application Server node. You can also use this page to change the type of a node from CONFIG to SSH to enable remote communication for the node.

If you change the type of the node from CONFIG to SSH, default values are applied if any of the following fields is left empty:

- SSH Port
- User Name
- Key File

The options that the Edit Node page contains depend on the type of the node.

4.7.2.1 Options for All Types of Nodes

The following options are available for all types of nodes:

- Name

  The name that was assigned to the node when the node was created. This field is read only.

- Type

  Drop-down list to specify the type of the node, which is one of the following types:CONFIGThe node does not support remote communication. If CONFIG, is selected, no options for SSH configuration are displayed. If the current type of the node is SSH and you select CONFIG, SSH configuration data is removed from the node's configuration when you save your changes.

- Node Host

  The name of the host that the node is to represent after the node is updated. Whether this field is mandatory depends on the type of the node:If the type of the node is CONFIG, the name of the host may be omitted. The name of the host can be determined when instances that reside on the node are created.If the type of the node is SSH, the name of the host must be specified. Otherwise, an error occurs.

- Node Directory

  The path to the directory that is to contain Apusic Application Server instances that are created on the node. If a relative path is specified, the path is relative to the *as-install* directory, where *as-install* is the base installation directory of the Apusic Application Server software on the host.If this field is left empty, the configuration of the node is updated as follows:If the type of the node is CONFIG, no directory for instances is specified for the node in the domain administration server (DAS) configuration. The node directory can be determined when instances that reside on the node are created.If the type of the node is SSH, the node directory is *as-install*`/nodes`, where *as-install* is the base installation directory of the Apusic Application Server software on the host.

- Installation Directory

  The full path to the *parent* of the base installation directory of the Apusic Application Server software on the host, for example, `/export/Apusic Application3/`.If this field is left empty, the configuration of the node is updated as follows:If the type of the node is CONFIG, no installation directory is specified for the node in the DAS configuration. The installation directory can be determined when instances that reside on the node are created.If the type of the node is SSH, the installation directory is the parent of the default base installation directory of the Apusic Application Server software for the DAS. This setting is useful only if Apusic Application Server is installed in the same location on all hosts.

- License File

  Display the license of Apusci Application Server. If you need to update, copy the content of the license file to the input box, replace the current license, and click "Update License File" to take effect in real time. If the license of the node needs to be replaced, replace it in the corresponding location in Node page. After enabling email verification, a reminder message will be sent to the email address 7 days before the license expires. The reminder will only stop sending after the license is replaced.

4.7.2.2 Additional Options for SSH Nodes

The following options are available only if the Type field is set to SSH:

- Node Agent Port

  Specify the node proxy port,used to manage agents such as load balancers and cache cluster instances; default 1099, the port needs to be open to the outside world. During the creation of a node of type "SSH", the node agent program nodemanager is installed by default. If the node agent program installation is unsuccessful or there is a conflict with the node agent port, it will not affect the installation and creation of the node.

- Start the Node Agent

  If the node agent status is unknown or stopped, you can click the button to start node agent.

- Stop the Node Agent

  If the node agent status is started, you can click the button to stop node agent.

- Force

  If the Enabled option is selected, the node is updated in the domain administration server (DAS) configuration even if validation of the node's parameters fails. To validate a node's parameters, the DAS must be able to contact the node's host through SSH.

- SSH Port

The port to use for SSH connections to this node's host.If you change the type of the node from CONFIG to SSH and leave this field empty, the SSH Port option is set to 22.If the Node Host option is set to `localhost-`*domain*, the SSH Port option is ignored.

- SSH User Name

  The user on this node's host that is to run the process for connecting to the host through SSH.If you change the type of the node from CONFIG to SSH and leave this field empty, the User Name option is set to the user that is running the DAS process.If the Node Host option is set to `localhost-`*domain*, the SSH User option is ignored.

- SSH User Authentication

  Drop-down list to specify how the SSH user is authenticated when logging in to this node's host. The possible methods for authentication are as follows:Key FileThe SSH user is authenticated through an SSH private key file. If you select this option, specify the absolute path to the file in the Key File field.PasswordThe SSH user is authenticated through a password that is stored in the DAS configuration.**Caution:**This password is stored in clear text in the DAS configuration. For enhanced security, use a key file or a password alias.If you select this option, the SSH User Password field is activated to enable you to provide the password.Password AliasThe SSH user is authenticated through a password alias. If you select this option, select the password alias from the Password Alias drop-down list.

- Setup SSH

  If the Enabled option is selected, Apusic Application Server sets up an SSH key on the node host. Apusic Application Server sets up the SSH key by copying existing key files from the DAS host or, if required, by generating the files before copying the files. This option is available only if the SSH User Authentication option is set to Key File.

- Generate Key

  If the Enabled option is selected, Apusic Application Server generates the SSH key files before copying the files to the node host if the files do not exist. If the SSH key files exist, Apusic Application Server does not generate the files, even if the Enabled option is selected.**Note:**If the SSH key files do not exist, the Enabled option must be selected. Otherwise, an error occurs when you attempt to save your changes.

- SSH User Password

  The password that the SSH user will use when logging in to this node's host.

- Password Alias

  Drop-down list of password aliases that are defined for this domain.

- Public Key File

  The absolute path to the SSH public key file for user that the SSH User Name option specifies. The content of the public key file is appended to the user's `.ssh/authorized_keys` file on the node host. If the `.ssh/authorized_keys` file does not exist on the host, Apusic Application Server creates the file.

- Key File

  The absolute path to the SSH private key file for user that the User Name option specifies. This file is used for authentication to the `sshd` daemon on the node's host.**Note:**Apusic Application Server also supports password authentication through the `AS_ADMIN_SSHPASSWORD` entry in the password file. The password file is specified in the `--paswordfile` option of the `asadmin` utility.If the SSH private key file is protected by a passphrase, the password file must contain the `AS_ADMIN_SSHKEYPASSPHRASE` entry.The path to the key file must be reachable by the DAS and the key file must be readable by the DAS.The default depends on whether you change the type of the node from CONFIG to SSH to enable SSH communication for the node:If you change the type of the node from CONFIG to SSH and leave this field empty, the Key File option is set to a key file in the user's `.ssh` directory. If multiple key files are found, Apusic Application Server uses the following order of preference: `id_rsa``id_dsa``identity`

4.7.2.3 LoadBalancer Instances

Use the LoadBalancer Instances page to display information about all load balancers under the current node, you can perform operations such as start, stop, restart, and delete. Clicking on the name of a load balancer will redirect you to the edit page for that specific load balancer instance.

4.7.2.4 Cache Cluster Instances

Use the Cache Cluster Instances page to display information about all cache managerment under the current node, you can perform operations such as start, stop, and delete.

4.7.3 Delete Node

In the node list, select the nodes that need to be deleted or uninstalled, and you can perform deletion or uninstallation operations.

Delete: A button used to delete one or multiple selected nodes. The deletion operation only removes the node information from the management platform and does not delete the node information from the node installation directory.

Delete and Uninstall: A button used to delete and uninstall one or multiple selected nodes, and uninstall the Apusic application server software from the host represented by each deleted node. After uninstallation, the node agent thread will also be terminated.

Note:

1. When there are instances with a status of "Running" under a node, the deletion or uninstallation operation cannot be performed. You need to manually stop the instances and delete them before you can delete or uninstall the node.
2. The option of Delete, Delete and Uninstalling will affect all information of the node, please be careful.

4.8 Standalone Server Instances

The "Standalone Server Instance" page allows you to create and manage standalone server instances. If you want a server to be accessible through multiple HTTP ports, you can use the Create Independent Instance method. Independent instances include local independent instances and remote independent instances.

Use the Standalone Server Instances page to create and manage standalone Apusic Application Server instances.

The Standalone Server Instances page displays a list of all standalone instances in the domain. For each instance, the following information is provided.

- Name

  The name that was assigned to the instance when the instance was created. Clicking the name opens the General Information page for the instance.

- Configuration

  The configuration that was selected for the instance when the instance was created. Clicking the configuration opens the Configurations page for the instance.

- Node

  The name of the node on which the instance resides. Clicking the name opens the Edit Node page for the node.

- Status

  An indication of whether the instance is running or not running.

The Server Instances table also contains the following options.

- New

  Button to create a standalone instance.

- Delete

  Button to delete one or more selected standalone instances.

- Start

  Button to start one or more selected standalone instances.

- Stop

  Button to stop one or more selected standalone instances.

4.8.1 New Standalone Server Instance

Use the New Standalone Server Instance page to create a standalone Apusic Application Server instance.

The New Standalone Server Instance contains the following options.

- Instance Name

  The name of the instance that is being created.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase lettersNumbersHyphenPeriodUnderscoreThe name must be unique in the domain and must not be the name of another Apusic Application Server instance, a cluster, a named configuration, or a node.The name must not be `domain`, `server`, or any other keyword that is reserved by Apusic Application Server.

- Node

Drop-down list of existing nodes that define the hosts where the instance can reside. One node must be selected from the list. If the instance is to be created on the host where the domain administration server (DAS) is running, select the predefined node `localhost-` *domain*.

- Configuration

  Drop-down list of existing named configurations that the instance can use. One configuration must be selected from the list. The instance will use the selected configuration when the instance is created.

- Make a Copy of the Selected Configuration

  If this option is selected, the selected configuration is copied when the instance is created. The copy of the configuration is assigned the name *instance-name* `-config`, where *instance-name* is the name of the instance as typed in the Instance Name field.If `default-config` is selected from the Configuration drop-down list, you must select this option. The `default-config` configuration can only be copied, not referenced.

- Reference the Selected Configuration

  If this option is selected, the instance will use the specified existing named configuration.If `default-config` is selected from the Configuration drop-down list, you must *not* select this option. The `default-config` configuration cannot be referenced, only copied.

4.8.2 Edit Standalone Server Instance

By clicking on the instance name, you can enter the General Information page of that instance where you can perform operations such as starting, stopping, deploying applications, and more.

4.8.2.1 General Information

The General Information page contains the following options.

- Start

  Button to start the instance. If the instance is running, this button is deactivated.

- Stop

  Button to stop the instance. If the instance is stopped, this button is deactivated.

- Restart

  Button to restart the instance. If the instance is stopped, this button is deactivated.

- View Log Files

  Button to view log files for the instance.

- View Raw log

  Button to view real time log files for the instance.By accessing the View Row Log page, you can view or download the access log files.

- Rotate Log

  Button to rotate the log file for the instance. When the log file is rotated, the file is renamed with a timestamp name in the format `server.log_` *date-and-time*, and an empty log file is created for new log messages. The changes are applied dynamically. Server restart is not required.

- Recover Transactions

  Button to recover transactions for the instance.

- View Access Log

  Button to view access log files for the instance.By accessing the Access Log page, you can view or download the access log files.

- Name

  The name that was assigned to the instance when the instance was created. This field is read only.

- JVM

  If you click JVM Report, a separate window opens and displays reports on the Java Virtual Machine, including a summary report, memory management and garbage collection report, class loading report, and a current thread dump.

- Status

  An indication of whether the instance is running or not running. This field is read only.

- Node

  The name of the node on which the instance resides. Clicking the name opens the Edit Node page for the node.

- Configuration

  The configuration that the instance references. Clicking the configuration opens the Configurations page for the instance.

- Debug

  An indication of whether Java Platform Debugger Architecture (JPDA) (`http://www.oracle.com/technetwork/java/javase/tech/jpda-141715.html/`) debugging is enabled for the instance. This field is read only.If an instance is started from the Administration Console, JPDA debugging is not enabled for the instance. To start an instance with JPDA debugging enabled, use the start-instance subcommand.

- Detection Type

  Set the instance detection method. Detection methods include UR, JDBC, Ping, and Socket URL: Request the corresponding instance through URL, and return the required basic information of the instance. JDBC: Through JDBC. Test whether the Jndi data source of the specified instance correctly executes the specified SQL statement. Ping: Determine whether the physical machine where the instance is located is Ping-able through Ping. Socket: Determine whether the physical machine IP and port of the instance are normal and can be connected through Ping.

- HTTP Ports

  A comma-separated list of the port numbers of the ports on which the instance listens for the following types of requests:Administration requestsHTTP requestsHTTPS requestsThis field is read only.

- IIOP Ports

  A comma-separated list of the port numbers of the ports that the instance uses for the following types of connections:Secure IIOP connectionSecure IIOP connections with client authenticationIIOP connectionsThis field is read only.

4.8.2.2 Applications

Use the Applications page to perform the following tasks on the selected standalone Apusic Application Server instance:

- Viewing and managing the applications that are already deployed to the instance
- Deploying more applications to the instance

The Instance Name field is a read-only field that displays the name of the current instance.

The Applications page displays a list of applications that are deployed to the selected instance. For each application, the following information is provided.

- Name

  The application name.

- Enabled

  An indication of whether the application is enabled:If the application is enabled, a check mark (✓) is displayed.If the application is disabled, a cross (✗) is displayed.

- Engines

  The types of containers that the application uses. Container types can be any of the following types: `web``webservices``ejb``connector``appclient``weld` — the container for Contexts and Dependency Injection for the Java EE Platform applications

The Deployed Applications table also contains the following options.

- Deploy

  Button to deploy an application.

- Undeploy

Button to undeploy one or more selected applications.Undeploying an application removes the application from all targets and the domain. When an application is undeployed, the application is no longer listed on the Applications page that displays all applications that are deployed in Apusic Application Server.To remove applications only from the current target, use the Remove button on this page.

- Remove

  Button to remove one or more selected applications.Removing an application removes the application from only the current target. When an application is removed from a target, the application continues to be listed on the Applications page that displays all applications that are deployed in Apusic Application Server. The application also remains deployed on any other targets on which it was deployed.To remove applications from all targets and the domain, use the Undeploy button on this page.

- More Actions

  Drop-down list that performs additional actions on one or more selected applications.EnableEnables one or more selected applications.DisableDisables one or more selected applications.Load Balancer EnableEnables one or more selected applications for load balancing.Load Balancer DisableDisables one or more selected applications for load balancing.

4.8.2.3 Resources

Use the Resources page to perform the following tasks on the selected standalone Apusic Application Server instance:

- Viewing and managing the instance's existing resources
- Creating additional resources for the instance

The Instance Name field is a read-only field that displays the name of the current instance.

The Resources page displays a list of the instances's existing resources. For each resource, the following information is provided.

- Resource Name

  The name that was assigned to the resource when the resource was created. Clicking the name opens one of the following pages to edit the resource.

- Enabled

  An indication of whether the resource is enabled:If the resource is enabled, a check mark (✓) is displayed.If the resource is disabled, a cross (✗) is displayed.

- Type

  The type of the resource.

The Resources table also contains the following options.

- Enable

  Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

- New

  Drop-down list that creates a resource of the selected type. Selecting a resource type from the list opens one of the following pages to create a resource of the selected type.

- Filter

  Drop-down list that filters the list of displayed resources by type.

4.8.2.4 Instance Properties

The Instance Properties page displays a list of the properties that are set for the current Apusic Application Server instance. These properties add optional configuration information about the instance.

The Instance Name field is a read-only field that displays the name of the current instance.

For each property, the following information is displayed:

- Name

The name of the property.

- Value

  The value that of the property that is set for the selected instance.

- Description

  A textual description that provides more information about the property.

The Additional Properties table also contains the following options.

- Add Property

  Button to add a property. Clicking this button adds a row to the Additional Properties table.

- Delete Properties

  Button to delete one or more selected properties. Any property that is deleted reverts to its default value or, if no default value is set, is undefined.

Apusic Application Server defines the following instance properties:

- `rendezvousOccurred`

  Specifies whether the instance has contacted the domain administration server (DAS). Possible values are as follows: `true` The instance has contacted the DAS. `false` The instance has *not* contacted the DAS.

4.8.2.5 Instance System Properties

The Instance System Properties page displays a list of the Java system properties that are set for the current Apusic Application Server instance. Java system properties are passed to the Java application launcher through the `-D` option of the Java application launcher whenApusic Application Server is started.

These properties override property definitions for port settings in the instance's configuration. Predefined port settings must be overridden if, for example, two clustered instances reside on the same host. In this situation, port settings for one instance must be overridden because both instances share the same configuration.

The Instance Name field is a read-only field that displays the name of the current instance.

For each property, the following information is displayed:

- Instance Variable Name

  The name of the system property.

- Current Value

  The value that is currently set for the property. This field is read only.

- Override Value

  The value to which the property will be set after changes are saved.

The Additional Properties table also contains the following options.

- Add Property

  Button to add a property. Clicking this button adds a row to the Additional Properties table.

- Delete Properties

  Button to delete one or more selected properties. Any property that is deleted reverts to its default value or, if no default value is set, is undefined.

4.8.2.6 Monitoring

Use the Monitoring page to configure monitoring and view monitoring data for the server instance,the details please see Monitoring Data.

4.8.3 Delete Standalone Instance

Select the instance that need to be deleted, and you can perform deletion operations.The option of Delete will affect all information of the instance, please be careful.

4.9 Clusters

4.9.1 Application Server Clusters

Application Server Cluster refers to a larger computer service system composed of a group of several independent computers through a high-speed communication network. Each cluster node, that is, each computer in the cluster, is an independent server that runs its own services. These servers can communicate with each other and collaborate to provide users with applications, system resources, and data, and are managed in a single system model.

Application Server Clusters have multiple benefits:

- High availability: When a server in the cluster fails or goes down, other servers can take over its work to ensure the continuous availability of services.
- Load balancing: The cluster can distribute requests to different servers to balance the load of the servers and avoid overload of a single server.
- Improved system performance: By using multiple servers in a cluster, computing and processing tasks can be distributed across different servers, thereby improving the overall performance and throughput of the system
- Easy scalability: The cluster can be expanded according to demand, by adding new servers to increase the system's capacity and performance.
- Improve system security: By distributing sensitive data and services across multiple servers, the risk of a single server being attacked or data being leaked can be reduced.
- Simplify maintenance and management: Cluster can simplify maintenance and management through centralized management and monitoring tools, reduce the workload of administrators, and improve the reliability and stability of the system.

The Application Server Clusters page displays a list of clusters in the domain. For each cluster, the following information is displayed:

- Name

  The name that was assigned to the cluster when the cluster was created. Clicking the name opens the General Information page for the cluster.

- Configuration

  The configuration that was selected for the cluster when the cluster was created. Clicking the configuration opens the Configurations page for the cluster.

- Instances

  The name and the status of each instance in the cluster. Clicking the name of the instance opens the General Informationpage for the instance.

- Status

  Display the current load status of the server cluster instance, mainly for the "Elastic Scaling" function. After enabling the "Elastic Scaling" function, there will be load information.

The Clusters table also contains the following options.

- New

  Button to create a cluster.

- Delete

  Button to delete a cluster.

- Start Cluster

  Button to start a cluster.

- Stop Cluster

  Button to stop a cluster.

4.9.1.1 New Application Server Cluster

Use the New Apllication Server Cluster page to create a cluster.

The New Application Server Cluster page contains the following options.

- Cluster Name

  The name of the cluster.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase lettersNumbersHyphenPeriodUnderscoreThe name

must be unique in the domain and must not be the name of a another cluster, a named configuration, a Apusic Application Server instance, or a node.The name must not be `domain`, `server`, or any other keyword that is reserved by Apusic Application Server.

- Configuration

  Drop-down list of existing named configurations that the cluster can use. A cluster requires a named configuration that defines the configuration of all instances that are added to the cluster. One configuration must be selected from the list. The cluster will use the selected configuration when the cluster is created.

- Make a Copy of the Selected Configuration

  If this option is selected, the selected configuration is copied when the cluster is created. The copy of the configuration is assigned the name *cluster-name* `-config`, where *cluster-name* is the name of the cluster as typed in the Cluster Name field.If `default-config` is selected from the Configuration drop-down list, you must select this option. The `default-config` configuration can only be copied, not referenced.

- Reference the Selected Configuration

  If this option is selected, the cluster will use the specified existing named configuration.If `default-config` is selected from the Configuration drop-down list, you must *not* select this option. The `default-config` configuration cannot be referenced, only copied.

- Message Queue Cluster Config Type

  Use this option to specify the type of Message Queue broker cluster to use for the new Apusic Application Server cluster. By default, Apusic Application Server uses a conventional Message Queue broker cluster with a master broker, with a broker embedded in each of the instances in the Apusic Application Server cluster.If this option is set to Custom, several additional options are displayed to configure the Message Queue broker cluster to use for the Apusic Application Server cluster. See Options for Custom Message Queue Broker Clusters.

- Server Instances to be Created

  A list of Apusic Application Server instances that are to be created when the cluster is created. For each instance, the following information is provided:Instance NameThe name of the instance.WeightAn integer that represents the load-balancing weight of the instance.The load-balancing weight determines the proportion of all requests to the cluster that the instance should process.For example, in a two-instance cluster, you might require one instance to process one out of four requests, and the other instance to process three out of four requests. In this situation, set the weight of the instance that is process one out of four requests to 1 and set the weight of the other instance to 3. If you prefer to use percentages, set the weights of the instances to 25 and 75 respectively.The default weight is 100.NodeDrop-down list of existing nodes where the instance can reside. One node must be selected from the list. The instance will reside on the selected node when the cluster is created.The Server Instances to be Created table also contains the following options.NewButton to create an instance. Clicking this button adds a row to the Server Instances to be Created table.DeleteButton to delete an instance.

4.9.1.1.1 OPTIONS FOR CUSTOM MESSAGE QUEUE BROKER CLUSTERS

When the Message Queue Cluster Config Type is set to Custom, the following options are displayed to configure the Message Queue broker cluster to use for the Apusic Application Server cluster.

- JMS Service Type

  The type of brokers (called JMS hosts in Apusic Application Server) to use in the broker cluster. The type chosen determines what types of broker clusters are available. The broker types are as follows.EmbeddedBroker configuration and lifecycle are managed by Apusic Application Server. Each Apusic Application clustered instance is serviced by a broker running in the same JVM as the instance.If the JMS Service Type option is set to Embedded, only conventional Message Queue broker clusters are supported. Options to configure the conventional cluster are displayed.LocalBroker configuration and lifecycle are managed by Apusic Application Server. Each Apusic Application clustered instance is serviced by a broker running in a separate JVM on the same host as the instance.If the JMS Service Type option is set to Local, both conventional and enhanced (highly available) Message Queue broker clusters are supported. Options to select the cluster type and to configure the cluster are displayed.RemoteBroker and broker cluster configuration and lifecycle are managed using Message Queue administrative tools.If the JMS Service Type option is set to Remote, no other options are displayed.

- JMS Cluster Type

  The type of Message Queue broker cluster to use for the Apusic Application Server cluster.If the JMS Cluster Type option is set to Conventional, the following options are displayed.JMS Configuration Store TypeThe type of data store to use for the conventional cluster's configuration data. If this option is set to Master Broker, one broker in the cluster is designated as the master broker and the configuration data is stored by it. If this option is set to Shared DB, the configuration data is stored in a JDBC data store accessible to all the brokers. In this case, the database-related options are displayed.JMS Message Store TypeThe type of data store each broker is to use to store its message data. If this option is set to File, each broker stores its message data in a file-based data store. If this option is set to JDBC, each broker stores its message data in a JDBC data store.If the JMS Cluster Type option is set to Enhanced (HA), the database-related options are displayed.

- Database Vendor

- Database URL
- Database User

  The database vendor, access url, and user of the JDBC database to use in any of these situations:When the configuration store type of a conventional cluster is set to Shared DBWhen the message store type of a conventional cluster is set to JDBC .

- Database Authentication
- Database Password
- Password Alias

  The password information for the JDBC database user specified in Database User.If Database Authentication is set to Password, the Database Password option is displayed. If Database Authentication is set to Password Alias, the Password Alias option is displayed.

- JMS Cluster Properties

  A list of one or more Message Queue broker properties to configure the brokers. The list is colon-separated (`:`) and has the form: `*prop1Name*=*prop1Value*:*prop2Name*=*prop2Value*:...` If a broker property name includes dots, preface the dots with two backslashes (`\\`); for example, to include the `imq.system.max_count` property, specify `imq\\.system\\.max_count` in the list.

4.9.1.2 General Information

The General Information page contains the following options.

- Start Cluster

  Button to start the cluster. If all instance in the cluster are running, this button is deactivated.

- Stop Cluster

  Button to stop the cluster. If all instances in the cluster are stopped, this button is deactivated.

- Migrate EJB Timers

  Button to migrate Enterprise JavaBeans (EJB) technology timers to from a stopped Apusic Application Server instance to a running instance in the cluster. Clicking this button opens the Migrate EJB Timers page. If all instances in the cluster are running, this button is deactivated.

- Rotate Log

  Button to rotate the log file for the cluster. When the log file is rotated, the file is renamed with a timestamp name in the format `server.log_` *date-and-time*, and an empty log file is created for new log messages. The changes are applied dynamically. Server restart is not required.

- Cluster Name

  The name that was assigned to the cluster when the cluster was created. This field is read only.

- JVM

  If you click JVM Report, a separate window opens and displays reports on the Java Virtual Machine, including a summary report, memory management and garbage collection report, class loading report, and a current thread dump.

- Configuration

  The configuration that was selected for the cluster when the cluster was created. Clicking the configuration opens the Configurations page for the cluster. This field is read only.

- Cache Cluster

  When the 'Cache Cluster' module has configuration information, you can select the corresponding cache cluster, and the applications in the server cluster will utilize the cache cluster functionality. After configuring the cache cluster, it is necessary to reload the applications in the server cluster.

- Balance

  When there is configuration information for the 'Load Balancer' or 'Load Balancer Cluster' module, this option can be selected. By selecting the corresponding load balancer or load balancer cluster, you can access the application through the load balancer. The configuration of the load balancer will take effect immediately.

- GMS

  If this option is enabled, the Group Management Service (GMS) is started in each Apusic Application Server instance in the cluster and in the domain administration server (DAS). The DAS participates in each cluster for which GMS is enabled. GMS provides cluster monitoring, cluster membership, and group communication services. This option is enabled by default.

- Multicast Port

  The port number of communication port on which GMS listens for group events. This option must specify a valid port number in the range 2048-49151. The default is an automatically generated value in this range.

- Multicast Address

  The address on which GMS listens for group events. This option must specify a multicast address in the range 224.0.0.0 through 239.255.255.255. The default is 228.9.*XX.YY*, where *XX* and *YY* are automatically generated independent values between 0 and 255.

- Bind Interface Address

  The Internet Protocol (IP) address of the network interface to which GMS binds. This option must specify the IP address of a local network interface. The default is all public network interface addresses.On a multihome machine, this option configures the network interface that used for the GMS. A multihome machine possesses two or more network interfaces.To specify an address that is valid for all Apusic Application Server instances in the cluster, use a system property to set the address individually for each instance.For example, use the Cluster System Properties page to create the system property `GMS-BIND-INTERFACE-ADDRESS-` *cluster-name*. Then set the Bind Address option in this page to `${GMS-BIND-INTERFACE-ADDRESS-` *cluster-name* `}` to specify the system property. Finally, for each instance in the cluster, set the `GMS-BIND-INTERFACE-ADDRESS-` *cluster-name* system property to the required network interface address on the instance's machine.

- Detection Type

  Set the instance detection method. Detection methods include UR, JDBC, Ping, and Socket URL: Request the corresponding instance through URL, and return the required basic information of the instance. JDBC: Through JDBC. Test whether the Jndi data source of the specified instance correctly executes the specified SQL statement. Ping: Determine whether the physical machine where the instance is located is Ping-able through Ping. Socket: Determine whether the physical machine IP and port of the instance are normal and can be connected through Ping.

- Status

  An indication of the number of instances in the cluster that are running and are not running. This field is read only.

- Elastic Scaling Switch

  When the status is set to "on", the cluster will enable the elastic scaling function. Please refer to the Elastic Scaling Function for more details.

4.9.1.2.1 ELASTIC SCALING

Enable the elastic scaling function to create or close running instances based on the set policy.

Based on the set "Elastic Scaling Detection Interval Seconds", the CPU or memory usage of the node is detected. When the "New Instance Threshold" is reached, an instance adjustment prompt will be displayed. If the "Elastic Scaling Adjustment Method" is "Automatic", an instance will be automatically added. If the "Elastic Scaling Adjustment Method" is "Manual", a prompt for adjustment will be displayed in the "Status" section of the load balancer cluster list page.

The configuration of Elastic Scaling contains the following options:

- Elastic Scaling Switch

  Set to on will turn on the elastic scaling function, and the default is off.

- Elastic Scaling AdjustMethod

  Automatic: The application server automatically adjusts based on the set elastic scaling policy.

  Manual: When the metrics reach the values set in the elastic scaling policy, an adjustment prompt will be displayed, requiring manual adjustment.

  Default is Automatic.

- Elastic Scaling Type

  To set an elastic scaling type, the options include CPU and memory.

- Elastic Scaling CheckInterval

The interval for elastic scaling detection is 60 seconds by default, which means it checks every 60 seconds.

- Elastic Scaling FrozenTime

  After adding a new cluster instance, the default waiting time for refreshing is 300 seconds, indicating that detection will be performed after 300 seconds.

- Elastic Scaling MinIns

  The minimum number of instances for cluster operation.

- Elastic Scaling MaxIns

  The maximum number of instances for cluster operation

- Elastic Scaling MinCPU

  Minimum threshold for CPU usage .When the average CPU threshold reaches or falls below 20%, close the instance, and then close instances based on the order of the node CPU usage rate.

- Elastic Scaling MaxCPU

  Maximum threshold for CPU usage.When the average CPU threshold reaches or exceeds 80%, create instances, starting with those with the lowest CPU usage across all nodes.

- Elastic Scaling MinMemory

  Minimum threshold for memory usage.When the average memory threshold reaches or falls below 20%, shut down the instance, and close it in order of the CPU usage rate of the node where the instance is located.

- Elastic Scaling MaxMemory

  Maximum threshold for memory usage.When the average memory threshold reaches or exceeds 80%, create instances, starting with those with the lowest CPU usage across all nodes.

Intelligent Routing Functionality: If the average CPU utilization of a server cluster is between the threshold for shutting down instances (%) and the threshold for adding new instances (%), and there are instances whose CPU utilization is below the shutdown threshold (%) while other instances are above the threshold for adding new instances (%), the weights of the server cluster instances are automatically adjusted based on their CPU idleness from highest to lowest.

4.9.1.3 Applications

Use the Applications page to perform the following tasks on the current cluster:

- Viewing and managing the applications that are already deployed to the cluster
- Deploying more applications to the cluster

The Cluster Name field is a read-only field that displays the name of the current cluster.

The Applications page displays a list of applications that are deployed to the current cluster. For each application, the following information is provided.

- Name

  The application name.

- Enabled

  An indication of whether the application is enabled:If the application is enabled, a check mark (✓) is displayed.If the application is disabled, a cross (✗) is displayed.

- Engines

  The types of containers that the application uses. Container types can be any of the following types: `web``webservices``ejb``connector``appclient``weld` — the container for Contexts and Dependency Injection for the Java EE Platform applications

The Deployed Applications table also contains the following options.

- Deploy

  Button to deploy an application.

- Undeploy

  Button to undeploy one or more selected applications.Undeploying an application removes the application from all targets and the domain. When an application is undeployed, the application is no longer listed on the Applications] page that displays all applications that are deployed in Apusic Application Server.To remove applications only from the current target, use the Remove button on this page.

- Remove

  Button to remove one or more selected applications.Removing an application removes the application from only the current target. When an application is removed from a target, the application continues to be listed on the Applications page that displays all applications that are deployed in Apusic Application Server. The application also remains deployed on any other targets on which it was deployed.To remove applications from all targets and the domain, use the Undeploy button on this page.

- More Actions

  Drop-down list that performs additional actions on one or more selected applications.EnableEnables one or more selected applications.DisableDisables one or more selected applications.Load Balancer EnableEnables one or more selected applications for load balancing.Load Balancer DisableDisables one or more selected applications for load balancing.

4.9.1.4 Clustered Server Instances

Use the Clustered Server Instances page to create and manage Apusic Application Server instances in the current cluster.

The Cluster Name field is a read-only field that displays the name of the current cluster.

The Clustered Server Instances page displays a list of instances that are members of the current cluster. For each instance, the following information is provided.

- Name

  The name that was assigned to the instance when the instance was created. Clicking the name opens the General Information page for the instance.

- LB Weight

  An integer that represents the load-balancing weight of the instance.The load-balancing weight determines the proportion of all requests to the cluster that the instance should process.For example, in a two-instance cluster, you might require one instance to process one out of four requests, and the other instance to process three out of four requests. In this situation, set the weight of the instance that is process one out of four requests to 1 and set the weight of the other instance to 3. If you prefer to use percentages, set the weights of the instances to 25 and 75 respectively.The default weight is 100.

- Configuration

  The configuration that the instance references. Clicking the configuration opens the Configurations page for the instance.

- Node

  The name of the node on which the instance resides. Clicking the name opens the Edit Node page for the node.

- Status

  An indication of whether the instance is running, requires restart to be updated with configuration changes, or is not running. This field is read only.

The Server Instances table also contains the following options.

- New

  Button to add an instance to the cluster.

- Delete

  Button to remove one or more selected instances from the cluster.**Caution:**If you are using a Java Message Service (JMS) cluster with a master broker, do not delete the instance that is associated with the master broker.

- Start

  Button to start one or more selected instances in the cluster.

- Stop

  Button to stop one or more selected instances in the cluster.

4.9.1.5 New Clustered Server Instance

Use the New Clustered Server Instance page to add a Apusic Application Server instance to an existing cluster.

The New Clustered Server Instance contains the following options.

- Cluster Name

  The name that was assigned to the cluster when the cluster was created. This field is read only.

- Instance Name

  The name of the instance that is being created.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase lettersNumbersHyphenPeriodUnderscoreThe name must be unique in the domain and must not be the name of another Apusic Application Server instance, a cluster, a named configuration, or a node.The name must not be `domain`, `server`, or any other keyword that is reserved by Apusic Application Server.

- Node

  Drop-down list of existing nodes that define the hosts where the instance can reside. One node must be selected from the list. If the instance is to be created on the host where the domain administration server (DAS) is running, select the predefined node `localhost-`*domain*.

4.9.1.6 Instance Properties

The Instance Properties page displays a list of the properties that are set for the current Apusic Application Server instance. These properties add optional configuration information about the instance.

The Instance Name field is a read-only field that displays the name of the current instance.

For each property, the following information is displayed:

- Name

  The name of the property.

- Value

  The value that of the property that is set for the selected instance.

- Description

  A textual description that provides more information about the property.

The Additional Properties table also contains the following options.

- Add Property

  Button to add a property. Clicking this button adds a row to the Additional Properties table.

- Delete Properties

  Button to delete one or more selected properties. Any property that is deleted reverts to its default value or, if no default value is set, is undefined.

Apusic Application Server defines the following instance properties:

- `rendezvousOccurred`

  Specifies whether the instance has contacted the domain administration server (DAS). Possible values are as follows: `true` The instance has contacted the DAS. `false` The instance has *not* contacted the DAS

4.9.1.7 Instance System Properties

The Instance System Properties page displays a list of the Java system properties that are set for the current Apusic Application Server instance. Java system properties are passed to the Java application launcher through the `-D` option of the Java application launcher whenApusic Application Server is started.

These properties override property definitions for port settings in the instance's configuration. Predefined port settings must be overridden if, for example, two clustered instances reside on the same host. In this situation, port settings for one instance must be overridden because both instances share the same configuration.

The Instance Name field is a read-only field that displays the name of the current instance.

For each property, the following information is displayed:

- Instance Variable Name

  The name of the system property.

- Current Value

  The value that is currently set for the property. This field is read only.

- Override Value

  The value to which the property will be set after changes are saved.

The Additional Properties table also contains the following options.

- Add Property

  Button to add a property. Clicking this button adds a row to the Additional Properties table.

- Delete Properties

  Button to delete one or more selected properties. Any property that is deleted reverts to its default value or, if no default value is set, is undefined.

4.9.1.8 Resources

Use the Resources page to perform the following tasks on the selected cluster:

- Viewing and managing the cluster's existing resources
- Creating additional resources for the cluster

The Cluster Name field is a read-only field that displays the name of the current cluster.

The Resources page displays a list of the clusters's existing resources. For each resource, the following information is provided.

- Resource Name

  The name that was assigned to the resource when the resource was created.

- Enabled

  An indication of whether the resource is enabled:If the resource is enabled, a check mark (✓) is displayed.If the resource is disabled, a cross (✗) is displayed.

- Type

  The type of the resource.

The Resources table also contains the following options.

- Enable

  Button to enable one or more selected resources.

- Disable

  Button to disable one or more selected resources.

- New

  Drop-down list that creates a resource of the selected type.

- Filter

  Drop-down list that filters the list of displayed resources by type.

4.9.2 Node Manager

The node agent is used to monitor the status of the load balancer or session manager and remotely execute commands to start, stop, and restart the load balancer or session manager. It is usually installed on the server of the load balancer and session manager, and attention should be paid to the port permissions.

4.9.2.1 Independently install node agent

Installation steps:

1. Obtain the installation package node-manager.tar.
2. Unzip the installation package and install it.

   Directly extract under Windows.

   Execute unzip node-manager.zip under Linux to decompress.
3. Configure the address and port of the load balancer node agent, edit the "installation path/node-manager/config/config.properties", and fill in the IP address and port of the host load balancer node agent (default is 1099, but it needs to be confirmed that the port is not occupied).
4. Start the load balancer node agent.

   In Windows, switch to the "installation path node-manager\bin" directory and execute the agent.cmd command as an administrator to start.

   In Linux, switch to the directory "installation path../node-manager/bin" and execute the command ./agent.sh to start.

   Or use the background start command: nohup ./agent.sh >/dev/null 2>&1 &.

4.9.2.2 Install node agent through nodes

When installing a remote node from the Kingdee Apusic application server console, the node agent will be automatically installed in the location ${APUSIC_HOME}/aas/nodes/[NODE_NAME]/nodemanager/. The default port is 1099 and the host is the server IP.

4.9.3 Cache Clusters

A caching cluster is a technical architecture designed to improve the availability and scalability of caching services through a set of independent computer nodes connected by a high-speed network. This architecture allows caching data to be shared and synchronized across multiple nodes, ensuring data consistency and reliability.

In a cache cluster, each node can act as an instance of a cache service and is responsible for managing a portion of the cached data. When a client requests data, the cluster selects a node to handle the request based on certain policies such as load balancing. If the required data does not exist on that node, it can obtain the data from other nodes and cache it for future quick access

The Apusic Application Server supports the creation and management of cache clusters.

Use the Cache Clusters page to create and manage cache clusters.

For each cache cluster, the following information is displayed.

- Name

  The name that was assigned to the cache cluster when the cache cluster was created. Clicking the name opens the General Information page for the cache cluster.

- Type

  The type of cache cluster.

- Use Status

  The number of clusters and standalone instances to which the cache cluster is targeted and how many of these targets the cache cluster is enabled on. .

- Status

  An indication of whether the cache instance is running or not running.

The Cache Clusters table also contains the following options.

- New

  Button to create a cache cluster.

- Delete

  Button to delete one or more selected cache cluster.You cannot delete an instance that is currently running; if the cache cluster is referenced by the server cluster, you need to cancel the reference before deleting it.

- Start

Button to start one or more selected cache clusters.

- Stop

  Button to stop one or more selected cache clusters.

4.9.3.1 New Cache Clusters

Use the New Cache Clusters page to create a cluster.

The New Cache Cluster page contains the following options.

- Cluster Name

  The name of the cluster.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase letters Numbers Hyphen Period Underscore.

- Cluster Type

  The type of cache cluster.

- Cluster Mode

  This option is displayed when the "Cache Cluster Type" is "amdc". The modes are cluster and master-slave.

  - cluster: Cluster mode. When the amdc cluster is set to cluster mode, at least three amdc node instances need to be set. To ensure high availability, the "follower quantity" needs to be set
  - master-slave: master-slave mode. When the amdc cluster is set to master-slave mode, at least the amdc master node instance needs to be specified

- Cluster Replicas

  This option is displayed when the "Cache Cluster Mode" is "cluster". To ensure high availability, it is necessary to specify the number of primary nodes to follow (the number of slave nodes). The cluster requires at least 3 node instances. When the "Follow Number" is set to 0, it means that no slave node instances are attached. When an instance failure occurs in the amdc cluster, all nodes under the cluster cannot be used. When the "Follow Number" is set to 1, it means that at least one slave node instance needs to be attached under the primary node instance. At this time, at least 6 amdc instances need to be created. And so on. The master-slave status of node instances is automatically assigned by amdc

- External Reference

  Whether to use an external cache instance is disabled by default. The cache instance is installed on the node at the location ${APUSIC_HOME}/nodes/[node_name]/[cacheins_name].

- Cache Instances to be Created

A list of cache instances that are to be created when the cluster is created. It is necessary to determine whether it is required based on the specific cache cluster type and cache cluster mode.

If the External Refrence is disabled,you need to first create a node with a Type of SSL in the Node module, and also need to install and run the node agent. For each instance, the following information is provided:

- Instance Name:The name of the instance.Automatically generated, the generation rule is [Cache Cluster Name] -> [Cache Cluster Type] -> [Time], which can be modified and can only contain letters, numbers, horizontal lines, or underscores. It must be unique.
- Master:Set as the master instance of the cluster instance,is disabled when the External Refrence is disabled.
- Node:The cache instance is installed on a node, and you can select any node under the current console's "Node" module. The selected node needs to have the node agent started.
- Port:The port of the cache instance can be set. The port needs to have the permission to be open to the outside world. When it is empty, the port will be automatically assigned after confirmation. The default range is 5701 - 5800.

If the External Refrence is enabled, you need to ensure that you have installed a cache manager, such as Hazelcast or ADMC, and you need to install and start the node agent on the server of the session manager. For each instance, the following information is provided:

- Instance Name:The name of the instance.Automatically generated, the generation rule is [Cache Cluster Name] -> [Cache Cluster Type] -> [Time], which can be modified and can only contain letters, numbers, horizontal lines, or underscores. It must be unique.
- Master:Set as the master instance of the cluster instance.When the "Cluster Mode" of the amdc cluster is set to master-slave, it needs to be set.

- Agent IP:Node Agent IP refers to the IP where the NodeManager (for the node agent) is installed.
- Agent Port:Node Agent prot .
- Excute Path:The cache session management execution file path, such as the execution file path of Hazelcast is ${HAZELCAST_INSTALL_DIR}/bin/start.sh; The execution file path of amdc is ${AMDC_INSTALL_DIR}/amdc-server
- Config Path:The configuration file path of the session manager, such as the configuration file path of Hazelcast, is ${HAZELCAST_INSTALL_DIR}/bin/hazelcast.xml; The execution file path of amdc is ${AMDC_INSTALL_DIR}/conf.yaml
- Port:The port of the cache instance can be set. The port needs to have the permission to be open to the outside world. When it is empty, the port will be automatically assigned after confirmation. The default range is 5701 - 5800.

The Instances to be Created table also contains the following options.NewButton to create an instance. Clicking this button adds a row to the cache Instances to be Created table.DeleteButton to delete an instance.

After confirmation, the cache cluster is successfully created, and upon returning to the list page, you need to click "Start" to run the cluster instance.

Note:

1. External cache instances and node cache instances cannot be in the same cluster at the same time.
2. When setting up external cache instances, it is important to note that the same version of Hazelcast needs to be used in the same cluster. For example, if you use Hazelcast4.0.3, all other instances need to be of the Hazelcast4.0 series, not 4.1/4.2, etc. AAS currently only supports versions of Hazelcast4.0 and above.
3. On the New Cache Cluster page, it is allowed not to create a cache cluster instance, but it is necessary to add an instance for the cluster on the Edit page, otherwise it cannot be started.

4.9.3.2 Edit Cache Clusters

By clicking on the cache cluster name, you can enter the General Information page of that cluster.

4.9.3.2.1 GENERAL INFORMATION

The General Information page contains the following options.

- Start Cluster

  Button to start the cluster. If the cluster instance is running, this button is deactivated.

- Stop Cluster

  Button to stop the cluster. If the cluster instance is stopped, this button is deactivated.

- Cluster Name: The name of the current cache cluster, which cannot be edited.
- Cluster Type: The type of the current cache cluster, which cannot be edited.
- Reference Cluster: The current cache cluster reference type. node indicates the node cache cluster instance, and external indicates the reference external cluster instance. It cannot be edited.
- Status: The current status of the cluster instance.

Note:

When adding or modifying the cache cluster option in the server cluster during the application's runtime, the application needs to be reloaded for the changes to take effect

4.9.3.2.2 INSTANCES

Use the Clustered Server Instances page to create new instances, delete, start, stop, or edit them. If the cluster type is "amdc", to ensure high availability of the cache cluster, you need to stop the cluster before adding or removing instances. If the cluster mode is "cluster", it is necessary to properly configure the number of instances when creating a new one.

The Cluster Name field is a read-only field that displays the name of the current cluster.

The Cluster Type field is a read-only filed that displays the type of the current cluster.

The Instances page displays a list of instances that are members of the current cluster. For each instance, the following information is provided.

- Instance Name

The name that was assigned to the instance when the instance was created. By clicking on the instance name, you can enter the instance editing page where editing is possible. For Hazelcast, the editable attribute is the port, and when modifying it, you need to ensure that the port is not in use and has access permissions. If the instance is running, the port cannot be modified, and you need to stop the running instance to change the port..

- Master

  Set as the master instance of the cluster instance.

- Port

  The port of the cache instance.

  Version

  The version of the session manager.

- Node

  The name of the node on which the instance resides.

- Status

  An indication of whether the instance is running, requires restart to be updated with configuration changes, or is not running. This field is read only.

The Clustered Cache Instances table also contains the following options.

- New

  Button to add an instance to the cluster.Display properties based on whether the created cluster references a type

- Delete

  Button to remove one or more selected instances from the cluster.

- Start

  Button to start one or more selected instances in the cluster.

- Stop

  Button to stop one or more selected instances in the cluster.

4.9.4 Load Balancers

Load Balancer is a network device or software that can distribute workload such as network traffic, data requests, or computing tasks to multiple computing resources such as servers, virtual machines, containers, etc., to optimize performance, improve reliability, and increase scalability

The Apusic Application Server console provides a graphical interface for load balancer management, which enables unified configuration management of load balancers, including apache, nginx, and ALB.

Use the Load Balancers page to create and manage load balancers.

For each cache load balancer, the following information is displayed.

- Name

  The name that was assigned to the load balancer when the load balacer was created. Clicking the name opens the [eneral Information page for the load balancer.

- Type

  The type of load balancer.

- LoadBalancer Version

  The version of load balancer.

- Status

  The number of clusters and standalone instances to which the cache cluster is targeted and how many of these targets the load balancer is enabled on.

- Running Status

An indication of whether the load balancer is running or not running.

The Load Balancers table also contains the following options.

- New

  Button to create a load balancer.

- Delete

  Button to delete one or more selected load balancer.Before delete load balancer ,ensure the load balancer have stopped running.

- Restart

  Button to restart load balancer.

- Stop

  Button to stop load balancer.

4.9.4.1 New Load Balancers

For the centralized management and configuration of load balancers and server clusters, Apusic Application Servers provide Apache, Nginx, and ALB as front-end processors. If Apache is used, the Apusic Application Server is proxied through mod_jk, requiring the creation or opening of a JK listener program.

Apusic Application Service supports node installation load balancer and external load balancer reference. Note that no matter what kind of reference mode is used, it is necessary to ensure that the node manager has been installed and started first.

Use the New Load Balancer page to create a cluster.

4.9.4.1.1 INSTALL LOAD BALANCERS THROUGH NODES

The New Load Balancer page contains the following options.

- Load Balancer Name

  The name of the load balancer.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase letters Numbers Hyphen Period Underscore.

- Type

  The type of load balancer. Such as Apache、NGINX or ALB.

- Failover

  When the "Type" is set to "Apache", this option will be displayed. Enabling the failover function will automatically transfer sessions to another node instance when one node instance becomes unavailable.

- Is the Session Sticky

  When the "Type" is set to "Apache", this option will be displayed. By enabling the session stickiness feature, users will continue to access the current session until the session expires.

- Static Cache

  If this option is selected, the static cache function is enabled. This option is selected by default.

- SSL

  If this option is selected,the SSL protocol is enable. Confirm whether to enable it on the new page. It is not supported to edit on the editing page. It is not enabled by default. After enabling, it supports one-way authentication and two-way authentication.

- SSL Authentication Method

  Enable SSL to display this item. Supports one-way authentication or two-way authentication

- Crt File Path

  Enable SSL to display this item. Upload the certificate file, such as server.crt, and save the file to the load balancer installation directory conf/.

- Key File Path

Enable SSL to display this item. Upload the private key file, such as server.key, and save the file in the load balancer installation directory conf/ after uploading.

- Chain File Path

  Enable SSL, and this option will be displayed when the "SSL authentication method" is "two-way authentication". Upload the root certificate, such as ca.crt, and the file will be saved in the load balancer installation directory conf/ after uploading.

- Domain Name

  The load balancer type is "alb", and SSL mutual authentication is enabled. This item is displayed and needs to be filled in based on the actual domain name of the certificate

- GM

  When the load balancer type is 'alb', enabling SSL will display this item. Checking the option to enable GM (Guomi) configuration requires uploading corresponding files in the .pem format. To enable GM functionality on ALB, you need to upload two certificate files and a private key file.

- Balancer SSL Port

  The load balancer type is "alb", and the support for national cryptography is displayed when it is enabled. Set the port for the load balancing SSL service; access this port when accessing.

- Algorithm

  The algorithm of load balancer,the option is deference between each load balancer.

  Apache's load balancing algorithm includes:

  - By weight: Access will be based on the weight of the server instance, from high to low.
  - By traffic: Access will be based on the server traffic of the server instance, from high to low.
  - By Busy Level: Access will be based on the server CPU usage rate of the server instance, from low to high

  The load balancing algorithms of NGINX or ALB include:

  - By weight: Access will be based on the weight of the server instance, from high to low.
  - Press the hash of access IP(session sticky): It maps the IP requested by the user to a hash value and then assigns it to a specific server. This algorithm is required to achieve session stickiness.
  - By response time: Access will be based on the response time of the server instance, from fast to slow.
  - Access by HASH of the URL: Access will be based on the HASH of the URL.
  - Sticky: The sticky module will be added. This item will only be displayed when the node is installed with nginx

- Whether to Refer to External Load Balancer

  If the option is selected, the load balancer have installed independently. This option is selected by default,but here you should cancel selection.

- Balancer IP

  The IP of load balancer.

- Balancer Port

  Customize the port of the load balancer, such as 8085, and confirm whether the port you entered is free on the load balancer machine.

- Node Proxy

  Select the management node of the load balancer, which is the node where the load balancer needs to be installed. If you select node74, the load balancer will be installed on node74 at ${NODE_HOME}/node74/[load_balancer_name] Note: When installing the load balancer apache or nginx on the node, a load balancer execution directory is required. The default is /opt/AAS/alb/ahs or /opt/AAS/alb/nhs, which cannot be modified. If files are later created in this directory, the corresponding load balancer cannot be created. It is necessary to ensure that there are no files in the /opt/AAS/alb/ahs or /opt/AAS/alb/nhs directory and that there is execution permission.

- Agent IP
- Agent Port

- Excution File Path
- Config File Path
- Connect Timeout

  Set the timeout for connecting to the backend server, with a limit of 1-3600 seconds, and the default value is 5 seconds.
- Socket Timeout

  Set the timeout for request processing by the backend, with a limit of 1-3600 seconds, and the default value is 300 seconds.
- Client Header Timeout

  Set the timeout for reading the request header, with a limit of 1-3600 seconds, and the default value is 30 seconds.
- Client Bodytimeout

  Set the timeout for reading the request body, with a limit of 1-3600 seconds, and the default value is 30 seconds.
- Max Processes

  Set the maximum number of processes, with a limit of 1-50, and the default value is 5.
- Max Connections

  Set the maximum concurrency, with a limit range of 1-50000, and the default value is 5000.

4.9.4.1.2 INSTALL LOAD BALANCERS INDEPENDENTLY

The New Load Balancer page contains the following options.

- Load Balancer Name

  The name of the load balancer.The name must meet the following requirements:The name may contain only ASCII characters.The name must start with a letter, a number, or an underscore.The name may contain only the following characters:Lowercase lettersUppercase letters Numbers Hyphen Period Underscore.
- Type

  The type of load balancer. Such as Apache、NGINX or ALB.
- Failover

  When the "Type" is set to "Apache", this option will be displayed. Enabling the failover function will automatically transfer sessions to another node instance when one node instance becomes unavailable.
- Is the Session Sticky

  When the "Type" is set to "Apache", this option will be displayed. By enabling the session stickiness feature, users will continue to access the current session until the session expires.
- Static Cache

  If this option is selected, the static cache function is enabled. This option is selected by default.
- SSL

  If this option is selected,the SSL protocol is enable. Confirm whether to enable it on the new page. It is not supported to edit on the editing page. It is not enabled by default. After enabling, it supports one-way authentication and two-way authentication.
- SSL Authentication Method

  Enable SSL to display this item. Supports one-way authentication or two-way authentication
- Crt File Path

  Enable SSL to display this item. Upload the certificate file, such as server.crt, and save the file to the load balancer installation directory conf/.
- Key File Path

  Enable SSL to display this item. Upload the private key file, such as server.key, and save the file in the load balancer installation directory conf/ after uploading.

- Chain File Path

  Enable SSL, and this option will be displayed when the "SSL authentication method" is "two-way authentication". Upload the root certificate, such as ca.crt, and the file will be saved in the load balancer installation directory conf/ after uploading.

- Domain Name

  The load balancer type is "alb", and SSL mutual authentication is enabled. This item is displayed and needs to be filled in based on the actual domain name of the certificate

- GM

  When the load balancer type is 'alb', enabling SSL will display this item. Checking the option to enable GM (Guomi) configuration requires uploading corresponding files in the .pem format. To enable GM functionality on ALB, you need to upload two certificate files and a private key file.

- Balancer SSL Port

  The load balancer type is "alb", and the support for national cryptography is displayed when it is enabled. Set the port for the load balancing SSL service; access this port when accessing.

- Algorithm

  The algorithm of load balancer,the option is deference between each load balancer.

  Apache's load balancing algorithm includes:

  - By weight: Access will be based on the weight of the server instance, from high to low.
  - By traffic: Access will be based on the server traffic of the server instance, from high to low.
  - By Busy Level: Access will be based on the server CPU usage rate of the server instance, from low to high

  The load balancing algorithms of NGINX or ALB include:

  - By weight: Access will be based on the weight of the server instance, from high to low.
  - Press the hash of access IP(session sticky): It maps the IP requested by the user to a hash value and then assigns it to a specific server. This algorithm is required to achieve session stickiness.
  - By response time: Access will be based on the response time of the server instance, from fast to slow.
  - Access by HASH of the URL: Access will be based on the HASH of the URL.
  - Sticky: The sticky module will be added. This item will only be displayed when the node is installed with nginx

- Whether to Refer to External Load Balancer

  If the option is selected, the load balancer have installed independently. This option is selected by default.

- Balancer IP

  The IP of load balancer.

- Balancer Port

  Customize the port of the load balancer, such as 8085, and confirm whether the port you entered is free on the load balancer machine.

- Agent IP

  Monitor the IP of the node agent manager of the load balancer.

- Agent Port

  Monitor the port of the node agent manager of the load balancer.

- Excution File Path

  Load balancer run command path. Example: nginx - /usr/sbin/nginx; apache - /usr/sbin/apachectl; alb - ${installation location}/bin/alb.

- Config File Path

Load balancer configuration file path. Example: nginx - /usr/conf/nginx.conf; apache - /usr/conf/httpd.conf; alb - ${installation location}/conf/conf.yaml.

- Connect Timeout

  Set the timeout for connecting to the backend server, with a limit of 1-3600 seconds, and the default value is 5 seconds.

- Socket Timeout

  Set the timeout for request processing by the backend, with a limit of 1-3600 seconds, and the default value is 300 seconds.

- Client Header Timeout

  Set the timeout for reading the request header, with a limit of 1-3600 seconds, and the default value is 30 seconds.

- Client Bodytimeout

  Set the timeout for reading the request body, with a limit of 1-3600 seconds, and the default value is 30 seconds.

- Max Processes

  Set the maximum number of processes, with a limit of 1-50, and the default value is 5.

- Max Connections

  Set the maximum concurrency, with a limit range of 1-50000, and the default value is 5000.

4.9.4.2 Edit Load Balancers

By clicking on the load balancer name, you can enter the General Information page of that load balancer. After modifying the properties, it is usually necessary to restart the load balancer before the properties take effect.

4.9.4.2.1 GENERAL INFORMATION

The General Information page contains the following options.

- Start

  Button to start the load balancer. If the load balancer is running, this button is deactivated.

- Stop

  Button to stop the load balancer. If the load balancer is stopped, this button is deactivated.

- Restart

  Button to restart the load balancer.

- Load Balancer Name

  The name that was assigned to the load balancer when the load balancer was created. This field is read only.

- Status

  The status of the load balancer.

- Type

  The type that was assigned to the load balancer when the load balancer was created. This field is read only.

- Oprator

  The oprator of the load balancer where installed.

- Failover

  When the "Type" is set to "Apache", this option will be displayed. Enabling the failover function will automatically transfer sessions to another node instance when one node instance becomes unavailable.

- Is the Session Sticky

  When the "Type" is set to "Apache", this option will be displayed. By enabling the session stickiness feature, users will continue to access the current session until the session expires.

- Static Cache

  If this option is selected, the static cache function is enabled. This option is selected by default.

- SSL

  If this option is selected,the SSL protocol is enable. Confirm whether to enable it on the new page. It is not supported to edit on the editing page. It is not enabled by default. After enabling, it supports one-way authentication and two-way authentication.This field is read only.

- SSL Authentication Method

  Enable SSL to display this item. Supports one-way authentication or two-way authentication

- Crt File Path

  Enable SSL to display this item. Upload the certificate file, such as server.crt, and save the file to the load balancer installation directory conf/.

- Key File Path

  Enable SSL to display this item. Upload the private key file, such as server.key, and save the file in the load balancer installation directory conf/ after uploading.

- Chain File Path

  Enable SSL, and this option will be displayed when the "SSL authentication method" is "two-way authentication". Upload the root certificate, such as ca.crt, and the file will be saved in the load balancer installation directory conf/ after uploading.

- Domain Name

  The load balancer type is "alb", and SSL mutual authentication is enabled. This item is displayed and needs to be filled in based on the actual domain name of the certificate

- GM

  When the load balancer type is 'alb', enabling SSL will display this item. Checking the option to enable GM (Guomi) configuration requires uploading corresponding files in the .pem format. To enable GM functionality on ALB, you need to upload two certificate files and a private key file.

- Balancer SSL Port

  The load balancer type is "alb", and the support for national cryptography is displayed when it is enabled. Set the port for the load balancing SSL service; access this port when accessing.

- Algorithm

  The algorithm of load balancer,the option is deference between each load balancer.

  Apache's load balancing algorithm includes:

  - By weight: Access will be based on the weight of the server instance, from high to low.
  - By traffic: Access will be based on the server traffic of the server instance, from high to low.
  - By Busy Level: Access will be based on the server CPU usage rate of the server instance, from low to high

  The load balancing algorithms of NGINX or ALB include:

  - By weight: Access will be based on the weight of the server instance, from high to low.
  - Press the hash of access IP(session sticky): It maps the IP requested by the user to a hash value and then assigns it to a specific server. This algorithm is required to achieve session stickiness.
  - By response time: Access will be based on the response time of the server instance, from fast to slow.
  - Access by HASH of the URL: Access will be based on the HASH of the URL.
  - Sticky: The sticky module will be added. This item will only be displayed when the node is installed with nginx

- Whether to Refer to External Load Balancer

  If the option is selected, the load balancer have installed independently. This option is selected by default.

- Node Proxy

  View the node where the load balancer is installed.When the option of Whether to Refer to External Load Balancer cancel select,and the field is read only.

- Balancer IP

  The IP of load balancer.

- Balancer Port

  Customize the port of the load balancer, such as 8085, and confirm whether the port you entered is free on the load balancer machine.

- Agent IP

  Monitor the IP of the node agent manager of the load balancer.

- Agent Port

  Monitor the port of the node agent manager of the load balancer.

- Excution File Path

  Load balancer run command path. Example: nginx - /usr/sbin/nginx; apache - /usr/sbin/apachectl; alb - ${installation location}/bin/alb.

- Config File Path

  Load balancer configuration file path. Example: nginx - /usr/conf/nginx.conf; apache - /usr/conf/httpd.conf; alb - ${installation location}/conf/conf.yaml.

- Connect Timeout

  Set the timeout for connecting to the backend server, with a limit of 1-3600 seconds, and the default value is 5 seconds.

- Socket Timeout

  Set the timeout for request processing by the backend, with a limit of 1-3600 seconds, and the default value is 300 seconds.

- Client Header Timeout

  Set the timeout for reading the request header, with a limit of 1-3600 seconds, and the default value is 30 seconds.

- Client Bodytimeout

  Set the timeout for reading the request body, with a limit of 1-3600 seconds, and the default value is 30 seconds.

- Max Processes

  Set the maximum number of processes, with a limit of 1-50, and the default value is 5.

- Max Connections

  Set the maximum concurrency, with a limit range of 1-50000, and the default value is 5000.

4.9.4.2.2 TAGETS

Use the Manage Targets page to change the target clusters and standalone server instances on which an load balancer can be enabled. The load balancer can be enabled only on targets in the Selected Targets column.

The Manage Targets page contains the following information.

- Available Targets

  The clusters and instances on which the load balancer is not deployed.

- Selected Targets

  The clusters and instances on which the load balancer is deployed.

- Add

  Button to move the selected target from Available Targets to Selected Targets.

- Add All

Button to move all Available Targets to Selected Targets.

- Remove

  Button to move the selected target from Selected Targets to Available Targets.

- Remove All

  Button to move all Selected Targets to Available Targets.

4.9.5 Load Balancer Clsuters

The Apusic Application Server supports the load balancer cluster function. The load balancer cluster in the LVS+keepalived+LB mode is referenced in the AAS server cluster to improve the anti-load capacity of the server cluster and to manage the load balancer cluster.

4.9.5.1 Environmental preparation

Before creating a load balancer cluster, you need to configure the load balancer in AAS. Refer to Load Balancer Configuration.

Before using the load balancer cluster, you need to prepare a load scheduling server with keepalived and ipvsadm installed.

The load scheduling server is configured with a virtual IP, as detailed in the virtual IP configuration section.

Load scheduling server installation keepalived reference keepalived installation.

Load scheduling server installation ipvsadm Reference ipvsadm installation.

4.9.5.2 Load Balancer Cluster

The Apusic Application supports the configuration and management of load balancer clusters, which can be configured through the control platform.

The Load Balancer Cluster page displays the following attributes:

- Name

  The name of the load balancer cluster. Click to enter the General page.

- Type

  The type of load balancer cluster.

- Usage status

  The load balancer cluster is referenced by the server cluster.

- Running status

  The status of the keepalived instance of the load balancer cluster. Click the keepalived instance name to enter the detailed information page of the keepalived instance.

- LoadBalancer Cluster User Status

  The status of the load balancer instances referenced by this load balancer cluster. Click on the number to enter the Load Balancer Instance page.

The Load Balancer Clusters table also contains the following options.

- New

  Button to create a load balancer cluster.

- Delete

  Button to delete one or more selected load balancer cluster.Before delete load balancer clsuter ,ensure the load balancer have stopped running.

- Start Cluster

  Button to start load balancer cluster.

- Stop Clsuter

  Button to stop load balancer cluster.

4.9.5.3 New Load Balancer Clusters

Use the New Load Balancer Clusters page to create a cluster.

The New Load Balancer Cluster page contains the following options.

- Cluster Name

  The name of the cluster. The name can only contain letters, numbers, horizontal lines, or underscores and must be unique

- Cluster Type

  The type of the cluster.Including Active/standby mode and Dual master mode,the default value is Active/standby mode.

- Virtual IP

  The virtual IP of the load balancers.The IP not assigned to the real host needs to be configured in the load scheduling server, that is, the load balancer instance host. In the Active/standby mode, only one Virtual IP needs to be configured; in the Dual master mode, two Virtual IPs need to be configured. Access the cluster through the Virtual IP.

- Domain name

  Defines the domain name of the Virtual IP, which is only displayed in the Dual master mode; the domain name needs to be configured in the client host, and the cluster can be accessed by accessing the Domain name.

- Load balancing forwarding rule

  The forwarding rule of the load balancer.The value is DR direct routing.

- Load balancing strategy

  Select load balancing scheduling strategies, including Polling(RR),Weighted polling(WRR),Source address hash(SH),IMininum connection(LC),Weighted least connection(WLC).

- Time interval between active and standby synchronization check

  The heartbeat packet sending cycle of the keepalived instance, the value between 1and 60 seconds,default is 3 seconds.

- Health check interval

  The time interval for checking the backend real load balancer status, the value between 1and 60 seconds,default is 6 seconds.

- Session retention timeout

  Keep client requests sent to the same real load balancer within a specified time period, the value between 1 and 3600 seconds,default is 0.

- Health detection method

  The health detection method of the cluster,default is TCP_CHECK.

- Connection timeout

  The timeout for connection to the load balancer instance, the value between 1 and 60 seconds, default to 6 seconds.

- Number of reconnections after detection failure

  The number of retries after the load balancer heartbeat detection fails. If the number of retries is reached and the failure still occurs, it will be removed from the server pool.The value between 1 and 10 ,default is 3 times.

- Failed reconnection interval

  The interval between reconnecting to the load balancer,the value between 1 and 60 , default is 3 seconds.

- Load Balancer Instance:

Select the target of the load balancer instance to be associated, choose an instance of the same type and port number with a virtual IP configured, and click Shift or Ctrl to select multiple instances. Once a load balancer is created in the Load Balancer section, it will be displayed synchronously there. This is a required field, and at least one load balancer instance is required to create a load balancer cluster

- Keepalived instance to be created: You can click "New..." to create a keepalived instance, or leave it blank and create it later.

The table of Keepalived instance to be created contains the following attributes:

- Instance name

  Set the name of the keepalived instance. The name can only contain letters, numbers, horizontal lines, or underscores and must be unique.

- Select active and standby

  Set the keepalived instance type. When the cluster mode is "master-backup mode", you need to create a master instance master and a backup instance backup; when the cluster mode is "dual master mode", you need to create two master instances master.

- Node agent IP

  The IP where the node agent is located.

- Node agent port

  The port set by the node manager, which requires access permission and is mandatory.

- Execute file path

  The execute file of keepalived, such as ${keepalived_root}/keepalived/sbin/keepalived. When empty, the default is /usr/sbin/keepalived.

- Profile path

  The profile path of keepalived, such as ${keepalived_root}/keepalived/etc/keepalived/keepalived.conf. When empty, the default is /etc/keepalived/keepalived.conf.

4.9.5.4 Edit Load Balancer Clusters

Use the Edit Load Balancer Clusters page to modify a cluster.

The Edit Load Balacer Cluster page contains the following options.

- Start Cluster

  Button to start the cluster.

- Stop Cluster

  Button to stop the cluster.

- Cluster Name

  The name of the cluster. This field is read only.

- Cluster Type

  The type of the cluster. This field is read only.

- Virtual IP

  The virtual IP of the load balancers. This field is read only.

- Domain name

  Defines the domain name of the Virtual IP, which is only displayed in the Dual master mode; the domain name needs to be configured in the client host, and the cluster can be accessed by accessing the Domain name.

- Load balancing forwarding rule

  The forwarding rule of the load balancer.The value is DR direct routing.

- Load balancing strategy

  Select load balancing scheduling strategies, including Polling(RR),Weighted polling(WRR),Source address hash(SH),IMininum connection(LC),Weighted least connection(WLC).

- Time interval between active and standby synchronization check

  The heartbeat packet sending cycle of the keepalived instance, the value between 1and 60 seconds,default is 3 seconds.

- Health check interval

The time interval for checking the backend real load balancer status, the value between 1and 60 seconds,default is 6 seconds.

- Session retention timeout

  Keep client requests sent to the same real load balancer within a specified time period, the value between 1 and 3600 seconds,default is 0.

- Health detection method

  The health detection method of the cluster,default is TCP_CHECK.

- Connection timeout

  The timeout for connection to the load balancer instance, the value between 1 and 60 seconds, default to 6 seconds.

- Number of reconnections after detection failure

  The number of retries after the load balancer heartbeat detection fails. If the number of retries is reached and the failure still occurs, it will be removed from the server pool.The value between 1 and 10 ,default is 3 times.

- Failed reconnection interval

  The interval between reconnecting to the load balancer,the value between 1 and 60 , default is 3 seconds.

- Status

  The status of the keepalived instance.

Note:

After modifying the properties, the load balancer cluster properties need to be restarted for them to take effect.

4.9.5.5 Edit Keepalived Instances

Use the Edit Keepalived Instances page to modify a cluster.

The Edit Keepalived Instances page contains the following options.

- Instance name

  The name of the keepalived instance under the current cluster. Click to view the instance information.

- Instance type

  Display the type of the keepalived instance.

- Version

  Displays the keepalived version of the keepalived instance.

- Status

  Displays the status of the keepalived instance. By default, the status is obtained every 6 seconds. If the displayed status is "unknown", refreshing will result in the status being obtained.

The Keepalived Instance table also contains the following options.

- New

  Button to create a keepalived instance.

- Delete

  Button to delete one or more selected keepalived instance.Before deleting, you need to stop the instance. After deleting, the instance information is deleted from the cluster information, and the domain.xml synchronizes the deleted configuration information. The keepalived file is not deleted, and the keepalived configuration file keepalived.conf is updated each time it is called.

- Start

  Button to start keepalived instance.

- Stop

Button to stop keepalived instance.

4.9.5.6 New Keepalived Instance

Use the New Keepalived Instance page to create a cluster.

The New Keepalived Instance page contains the following options.

- Cluster Name

  The name of the cluster,this field is read only.

- Cluster Type

  The type of the cluster. This field is read only.

- Instance name

  Set the name of the keepalived instance. The name can only contain letters, numbers, horizontal lines, or underscores and must be unique.

- Select active and standby

  Set the keepalived instance type. When the cluster mode is "master-backup mode", you need to create a master instance master and a backup instance backup; when the cluster mode is "dual master mode", you need to create two master instances master.

- Node agent IP

  The IP where the node agent is located.

- Node agent port

  The port set by the node manager, which requires access permission and is mandatory.

- Execute file path

  The execute file of keepalived, such as ${keepalived_root}/keepalived/sbin/keepalived. When empty, the default is /usr/sbin/keepalived.

- Profile path

  The profile path of keepalived, such as ${keepalived_root}/keepalived/etc/keepalived/keepalived.conf. When empty, the default is /etc/keepalived/keepalived.conf.

Note:

- Keepalived needs to be installed correctly, and the configuration file permissions need to be 644.
- The node agent needs to be in a running state.
- It is recommended to create only two keepalived instances, with the cluster mode set to "active-standby". When the cluster mode is "active-active", two master instances are required, and no standby instances can be created. If more than two keepalived instances are created, it is recommended to only start two.
- The Keepalived instance can be on the same server as the load balancer instance.

4.9.5.7 Edit Load Balancer

You can manage the load balancer instances of the cluster on the load balancer instance information page. The load balancer instances that join the load balancer cluster need to be configured before referencing the virtual IP load balancer instance list page, which displays the following attributes:

- Cluster Name

  The name of the cluster to which the load balancer instance belongs, read-only mode.

- Name

  The name of the load balancer instance. Click to enter the detailed information page of the load balancer.

- Load Balancer Type

  Displays the current load balancer type.

- Version

  Displays the version of the current load balancer.

- Weight

The proportion of this load balancer instance in the load balancer cluster. You need to enter a positive integer. After clicking "Save", you need to restart the load balancer cluster for the changes to take effect.

- Running Status

  The running status of the load balancer instance.

- Operation

  You can start, stop, or restart the load balancer.

- Update: Click to enter the page of selecting "Load Balancer Instance", and display the load balancers of type nginx, Apache, and ALB in the "Load Balancer" module. After selecting the load balancer, click "OK" to select the load balancer instance for the load balancer cluster.

Note:

- The server where the load balancer is located needs to be configured with a virtual IP, which should be consistent with the "virtual IP" filled in the general information of the load balancer cluster. When there is a "dual master mode", two virtual IPs need to be configured.
- The same type of load balancer with the same port configuration and the same virtual IP must be selected for the same load balancer cluster.
- After selecting OK, you need to restart the load balancer cluster for it to take effect.
- After being selected by the load balancer cluster, the load balancer does not affect the individual reference by the server cluster.
- The same load balancer can only be referenced by one load balancer cluster.

4.9.5.8 Attachment

4.9.5.8.1 VIRTUAL IP CONFIGURATION

Create lvsdr.sh and add the following content. VIP is the virtual IP that needs to be set. If multiple virtual IPs need to be configured at the same time, modify lo:0 to another number, such as lo:1.

```
#!/bin/bash
VIP=172.20.140.245
case "$1" in
start)
ifconfig lo:0 $VIP broadcast $VIP netmask 255.255.255.255 up
echo "0" > /proc/sys/net/ipv4/ip_forward
echo "1" > /proc/sys/net/ipv4/conf/lo/arp_ignore
echo "2" > /proc/sys/net/ipv4/conf/lo/arp_announce
echo "1" > /proc/sys/net/ipv4/conf/all/arp_ignore
echo "2" > /proc/sys/net/ipv4/conf/all/arp_announce
sysctl -p >/dev/null 2>&1
echo "lo:0 port starting"
;;
stop)
ifconfig lo:0 down
route del $SNS_VIP >/dev/null 2>&1
echo "0" > /proc/sys/net/ipv4/conf/lo/arp_ignore
echo "0" > /proc/sys/net/ipv4/conf/lo/arp_announce
echo "0" > /proc/sys/net/ipv4/conf/all/arp_ignore
echo "0" > /proc/sys/net/ipv4/conf/all/arp_announce
echo "lo:0 port closing"
;;
*)
echo "Usage: $0 {start | stop}"
exit 1
esac
exit 0
```

Execute `lvsdr.sh start` to start.

Execute `lvsdr.sh stop` to stop.

4.9.5.8.2 INSTALL KEEPALIVED

Execute `yum install -y keepalived` to install keepalived.

4.9.5.8.3 INSTALL IPVSADM

Execute `yum install -y ipvsadm` to install ipvsadm.

4.10 SessionCache

This module is mainly used to manage the session manager. It can configure centralized storage for each application and supports storage in memcached and redis.

This module can only create or edit cache manager information, and cannot start or stop the session manager; therefore, it is necessary to start the session manager externally to connect with it.

The name of the created session manager will be displayed on the web application deployment or editing page of [Application Management], and you can associate the application with the session manager according to your actual needs.

The Load Balancer Cluster page displays the following attributes:

- Name

  The name of the session cache. Click to enter the General page.

- Type

  The type of the session cache.

- Usage status

  Displays the number of targets currently referenced by the session manager, supporting simultaneous reference by multiple applications for the same session management.

The SessionCache table also contains the following options.

- New

  Button to create a session cache manager.

- Delete

  Button to delete one or more selected session cache manager.If the session manager has been referenced, it needs to be unreferenced before it can be deleted. The deletion here refers to deleting the configuration of the session manager from AAS, and does not delete the actual installed session manager.

4.10.0.1 New Seesion Cache Manager

Use the New Seesion Cache Manager page to create a session cache manager.The session cache type includes Redis and Memcached. Choosing different types requires configuring different properties.

If the session cache type is redis,the New Session Cache Manager page contains the following options.

- sessionCacheName

  The name of the session cache manager.

- sessionCacheType

  The type of the session cache manager.

- Immediately store attribute modifications

  Whether to store the session.setAttribute method call immediately in the centralized storage. Not checked by default.

- BlockWhenExhausted:

  If this option is selected,enabled the connection blocking function when the connection is exhausted.Enabled by default.

- EvictionPolicyClassName

  Set the class name of eviction policy.

- Fairness

  If this option is selected,enabled faireness.By default, it is disabled.

- JmxEnabled

  If this option is selected,enabled JMX manager.Enabled by default.

- JMX base name

The base name of JMX.

- JmxNamePrefix

The prefix name of JMX.

- Lifo

  If this option is selected,enabled the last in, first out strategy.

- MaxIdle

  Set the maximum number of idle connections,the default value is 8.

- MaxTotal

Set the maximum number of connections,the default value is 8.

- MaxWaitMillis

Set maximum waiting milliseconds, -1 means not enabled,the default value is -1.

- MinEvictableIdleTimeMillis

  Set the minimum idle time for eliminating connections,the default value is 1800000 milliseconds.

- MinIdle

  Set the minimum number of idle connections,the default value is 0.

- NumTestsPerEvictionRun

  Set the maximum number of elimination during inspection,the default value is 3 times.

- SoftMinEvictableIdleTimeMillis

  Set the minimum idle time for soft elimination,the default value is -1.

- TestOnBorrow

  If this option is selected ,enable checking when calling the border object method.Disabled by default.

- TestOnCreate

  If this option is selected ,enable checking when calling the Create Object method.Disabled by default.

- TestOnReturn

  If this option is selected ,enable checking when calling the return Object method.Disabled by default.

- TestWhileIdle

  If this option is selected ,enable check when idle.Disabled by default.

- ConnectionMode

  Set the connection type for Redis, with options including single/sequential/hard/cluster.

- Host

  Displayed when the ConnectionMode is single; Set Redis server address, required field.

- Port

  Displayed when the ConnectionMode is single; Set Redis server port, required field.

- Timeout

  Displayed when the ConnectionMode is single or sentinel; Set connection timeout, optional fields.

- Password

  Displayed when the ConnectionMode is single/continuous; Set Redis server password, optional fields.

- SentinelList

  Displayed when the ConnectionMode is sentinel; Redis server list, multiple servers separated by commas, for example: 192.168.101.44:26379192.168.101.45:26379192.168.101.46:26379.

- MasterName

  Displayed when the ConnectionMode is sentinel; MasterName is required field.

- ShardList

  Displayed when the ConnectionMode is sharp; shard list, multiple separated by commas. For example: `instance01 (host=xxx; port=xxx;password=xxx;connectionTimeout=xxx;soTimeout=xxx;weight=xxx)` .

- ClusterList

  Displayed when the ConnectionMode is cluster; Redis cluster list, multiple separated by commas. For example: `host: port, host: port` .

If the session cache type is memcached,the New Session Cache Manager page contains the following options.

- sessionCacheName

  The name of the session cache manager.

- sessionCacheType

  The type of the session cache manager.

- Immediately store attribute modifications

  Whether to store the session.setAttribute method call immediately in the centralized storage. Not checked by default.

- ServerAddress

  Set the address of memecached, in the format of "IP:port number", multiple addresses are separated by commas. This field is mandatory.

- BackupServerAddress

  Set the backup server address list in the format IP: Port Number with multiple addresses separated by commas. The number of addresses in this list can not exceed the number of addresses in the serverAddress.

- Weights

  The weight of memcached.

- Fairness

  If this option is selected,enabled faireness.By default, it is disabled.

- ConnectionPoolSize

  Set the connection pool size,default value is 1.

- BinaryProtocol

  If this option is selected ,enable the binary protocol.Disabled by default.

- TokyoTyrant

If this option is selected ,enable cTokyoTyrant.enabled by default.

- ConsistentHash

  If this option is selected ,enable consistentHash.Disabled by default.

- Connection timeout

  Set the connection timeout,the default value is 10 seconds.

- OpTimeout

  Set the optimeout,the default value is 3 seconds.

- HealSessionInterval

  Set the time interval for attempting to restore sessions,the default value is 2 seconds

- SessionIdleTimeout

  Set session idle time,the default value is 5 seconds.

- CheckSessionTimeoutInterval

  Set the timeout interval for detecting sessions,the default value is 1 second.

- ReadThreadCount

  Set the total number of read threads,the default value is 1.

- WriteThreadCount

  Set the total number of write threads,the default value is 0.

- EnableHeartBeat

  If this option is selected,enable heartbeat activation function.Enabled by default.

- SelectorPoolSize

  Set selector pool size,the default value is 1.

4.10.0.2 Edit Session Cache Manager

Use the Edit Session Cache Manager page to modify the session cache manager informations.

After modifying the properties, if the session manager has already been referenced, the application properties need to be reloaded for them to take effect.

4.10.0.3 Reference Session Cache Manager

The configured session manager can be referenced by applications in Application Management. Enter the deployment properties or editing page of the web application, which displays the sessionCache option. Select the name of the session manager that needs to be associated to be applied by the application.

The same session manager can be referenced by multiple applications.

After configuring session management on the application editing page, the application needs to be reloaded for it to take effect.

4.11 Monitoring Data

Use the Monitoring page to configure monitoring and view monitoring data for clustered and non-clustered server instances.

For each server instance, the following information is provided.

- Instance Name

  The name of the server instance. Click on the instance name to go to the General Information page for the instance.

- Cluster Name

  The name of the cluster to which the server instance belongs, if it is a clustered instance. Click on the cluster name to go to the General Information page for the cluster.

- Action

  Links that allow you to perform the following operations on the server instances.

  - Configure Monitoring

    Displays the Monitoring Service page for the cluster or stand-alone instance configuration.

- View Monitoring Data

  Links that allow you to view the following types of monitoring data for the server instances.

  - Application

    Displays the Application Monitoring page for the server instance.

  - Server

    Displays the Server Monitoring page for the server instance.

  - Resources

    Displays the Resource Monitoring page for the server instance

  - Graphic Monitoring

    Displays the Graphic Monitoring page for the server instance.

  - History Graphic Monitoring

    Displays the History Graphic Monitoring page for the server instance.

  - Class Loading Tree

    Displays the Class Loading Tree Monitoring page for the server instance.

  - JNDI Tree

    Displays the JNDI Tree Monitoring page for the server instance.

  - Sql Tracing

    Displays the Sql Tracing Monitoring page for the server instance.

  - Long Thread

    Displays the Long Thread Monitoring page for the server instance.

  - Snapshots

    Displays the snapshots Monitoring page for the server instance.

  - Full-GC

    Displays the Full-GC Monitoring page for the server instance.

4.11.1 To View Application Monitoring Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Service for details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the Application link that corresponds to the server instance you are monitoring.

   The Application Monitoring page opens.

3. From the Application drop-down list, select the application for which you want to view monitoring data.

4. From the Component drop-down list, select the application component for which you want to view monitoring data.

5. Click the Refresh button to update displayed monitoring data.
6. Click the Suspended button to stop refresh displayed monitoring data.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.
2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.
3. From the Application drop-down list, select the application for which you want to view monitoring data.
4. From the Component drop-down list, select the application component for which you want to view monitoring data.
5. Click the Refresh button to update displayed monitoring data.
6. Click the Suspended button to stop refresh displayed monitoring data.

4.11.2 To View Server Monitoring Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.
2. In the View Monitoring Data column, click the Server link that corresponds to the server instance you are monitoring.

   The Server Monitoring page opens.
3. From the View drop-down list, select the component for which you want to view monitoring data.
4. Click the Refresh button to update displayed monitoring data.
5. Click the Suspended button to stop refresh displayed monitoring data.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.
2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.
3. Click the Server tab.

   The Server Monitoring page opens.
4. From the View drop-down list, select the component for which you want to view monitoring data.
5. Click the Refresh button to update displayed monitoring data.
6. Click the Suspended button to stop refresh displayed monitoring data.

4.11.3 To View Resource Monitoring Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.
2. In the View Monitoring Data column, click the Resources link that corresponds to the server instance you are monitoring.

   The Resource Monitoring page opens.

3. From the Connection Pool drop-down list, select the connection pool for which you want to view monitoring data.
4. Click the Refresh button to update displayed monitoring data.
5. Click the Suspended button to stop refresh displayed monitoring data.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the Resources tab.

   The Resource Monitoring page opens.

4. From the Connection Pool drop-down list, select the connection pool for which you want to view monitoring data.
5. Click the Refresh button to update displayed monitoring data.
6. Click the Suspended button to stop refresh displayed monitoring data.

4.11.4 To View Graphic Monitoring Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the Graphic Monitoring link that corresponds to the server instance you are monitoring.

   The Graphic Monitoring page opens.

3. From the View list, select the option for which you want to view monitoring data.For example,jvm/cpu/thread/jdbc.
4. Click the Refresh button to update displayed monitoring data.
5. Click the Suspended button to stop refresh monitoring data.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the Graphic Monitoring tab.

   The Graphic Monitoring page opens.

4. From the View list, select the option for which you want to view monitoring data.For example,jvm/cpu/thread/jdbc.
5. Click the Refresh button to update displayed monitoring data.
6. Click the Suspended button to stop refresh displayed monitoring data.

4.11.5 To View History Graphic Monitoring Data

Before you can view monitoring data, you must configure monitoring. Enable the monitoring playback function in the monitoring service configuration interface, set the total amount of monitoring data, which is 8000 by default; set the monitoring sampling time in seconds, which is 30 seconds by default, indicating that monitoring is saved every 30 seconds.See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the History Graphic Monitoring link that corresponds to the server instance you are monitoring.

   The History Graphic Monitoring page opens.

3. From the View list, select the option for which you want to view monitoring data.

4. Select a time period and click Search button to search history monitoring data.

5. Click Slow Down button to slow down the playback speed; click Speed Up button to speed up the playback speed.

6. Click the Suspended button to stop refresh monitoring data,the page will no longer refresh, but the data will still be recorded.Click Play button to continue refreshing the data

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the History Graphic Monitoring tab.

   The History Graphic Monitoring page opens.

4. From the View list, select the option for which you want to view monitoring data.

5. Select a time period and click Search button to search history monitoring data.

6. Click Slow Down button to slow down the playback speed; click Speed Up button to speed up the playback speed.

7. Click the Suspended button to stop refresh monitoring data,the page will no longer refresh, but the data will still be recorded.Click Play button to continue refreshing the data.

4.11.6 To View Class Loading Tree Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the Class Loading Tree link that corresponds to the server instance you are monitoring.

   The Class Loading Tree Monitoring page opens.

3. The left tree view displays all the class names of AAS, and the right tree view displays the detailed information of the currently selected class loader.

4. Select the node in the tree diagram on the left, and the corresponding node class loading details will be displayed on the right.

5. Enter the full name of the class name in the condition query, such as "com.opensymphony.xwork2.util.ValueStack", and click "Search". All application class loading information using this class name will be displayed. If the class name does not exist, an error prompt will be displayed. The left node can be combined with the query condition input box for combined querying.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the Class Loading Tree tab.

   The Class Loading Tree Monitoring page opens.

4. The left tree view displays all the class names of AAS, and the right tree view displays the detailed information of the currently selected class loader.

5. Select the node in the tree diagram on the left, and the corresponding node class loading details will be displayed on the right.

6. Enter the full name of the class name in the condition query, such as "com.opensymphony.xwork2.util.ValueStack", and click "Search". All application class loading information using this class name will be displayed. If the class name does not exist, an error prompt will be displayed. The left node can be combined with the query condition input box for combined querying.

4.11.7 To View JNDI Tree Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the JNDI Tree link that corresponds to the server instance you are monitoring.

   The JNDI Tree Monitoring page opens.

3. The left tree view displays all the JNDI names of AAS, and the right tree view displays the detailed information of the currently selected JNDI name.

4. Select the node in the tree diagram on the left, and the corresponding node JNDI loading details will be displayed on the right.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the JNDI Tree tab.

   The JNDI Tree Monitoring page opens.

4. The left tree view displays all the JNDI names of AAS, and the right tree view displays the detailed information of the currently selected JNDI name.

5. Select the node in the tree diagram on the left, and the corresponding node JNDI loading details will be displayed on the right.

4.11.8 To View SQL Tracing Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the SQL Tracing link that corresponds to the server instance you are monitoring.

   The SQL Tracing Monitoring page opens.The trace list displays events, connection pool names, SQL statements, execution times, execution durations, maximum durations, execution times, and stack information.

3. Conditions for enabling the SQL tracing function:

   - Ensure that the "Security" in the instance management service-JMX connector is not enabled.
   - Ensure that the jdbc connection pool property "Wrap JDBC Objects" is checked.

   If it is a standalone instance or cluster instance:

   - Configure cluster or standalone instance management services to ensure that security is not enabled.
   - On the cluster or standalone instance management service page, modify the "address" to ${JMX_SYSTEM_CONNECTOR_URL}.
   - Enter the cluster or standalone instance system settings, add the system property JMX_SYSTEM_CONNECTOR_URL, and save it.

- Configure the instance value of the property JMX_SYSTEM_CONNECTOR_URL, and overwrite the value with the IP of the instance.

Start or restart the cluster or standalone instance.

4. Click "Start Trace" and "Trace Stack", select the data source you want to view, and the SQL tracking information for that data source under the current instance will be displayed.
5. Click "Show Configuration" to display the information that can be configured. The data source needs to be enabled to track; you can set tracking events by click Events button and click Confirm Configuration buttonto complete the setting.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the SQL Tracing tab.

   The SQL Tracing Monitoring page opens.

4. Click "Start Trace" and "Trace Stack", select the data source you want to view, and the SQL tracking information for that data source under the current instance will be displayed.
5. Click "Show Configuration" to display the information that can be configured. The data source needs to be enabled to track; you can set tracking events by click Events button and click Confirm Configuration buttonto complete the setting.

4.11.9 To View Thread Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the Thread link that corresponds to the server instance you are monitoring.

   The Thread Monitoring page opens.Monitoring busy threads, blocked threads, and long thread information during operation.

   - Instance name:

     The instance currently monitoring the long thread.

   - Refresh frequency:

     The frequency of refreshing the front-end page, default 10 seconds.

   - Pause

     Clicking will pause the refresh of the page without affecting the monitoring task. If monitoring has already been enabled, data monitoring will still occur in the background. After clicking, the button will change to "Auto Refresh"; clicking the "Auto Refresh" button again will change to "Pause".

   - Refresh

     Manually refresh the page data.

   - Save

     After modifying the configuration, click "Save" to take effect.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

The Application Monitoring page opens.

3. Click the Thread tab.

   The Thread Monitoring page opens.

4.11.9.1 Busy Threads

You can collect the CPU usage and print the stack traces of the top 10 busiest threads, and allow downloading the stack information.

- Start Busy Thread Monitoring

  Whether to enable busy thread monitoring. It is not enabled by default.

- Busy Thread Detection Cycle

  Perform busy thread detection at regular intervals, with a default of 200ms (minimum 200ms).

- Download

  Click to retrieve the information for downloading again; the downloaded file is a .zip file, and the stack information is in a .txt format.

- Stop

  Select a thread and click "Stop" to manually stop the busy thread.

4.11.9.2 Blocking thread

You can monitor the blocking thread and allow the download of stack information.

- Start Blocking Thread Monitoring

  Whether to enable blocking thread monitoring. The default is not enabled.

- Download

  Click to obtain information again and download the information obtained again. The download file is .zip and the stack information is .txt.

- Stop

  Select a thread and click "Stop" to manually stop the busy thread.

4.11.9.3 Long Threads

You can view the current information of long-running threads. For threads with excessively long request times, you can stop them as needed to prevent them from being occupied ineffectively for extended periods. Additionally, you can view the current thread call stack.

- Start Long Thread Monitoring

  Whether to enable monitoring of long-running threads. It is not enabled by default.

- Long Thread Threshold

  Monitors long-running threads exceeding a specified threshold, with a default of 60 seconds.

- Long Thread Detection Cycle

  Performs long thread detection at regular intervals, with a default of 3 seconds.

- Log Thread Log File Path

  Specifies the path for the log file where the stack information of long-running threads is output; the default is ${com.apusic.aas.instanceRoot}/logs/longthread.

- Download

  Click to retrieve the information for downloading again; the downloaded file is a .zip file, and the stack information is in a .txt format.

- Stop

  Select a thread and click "Stop" to manually stop the thread.

4.11.10 To View Snapshots Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the Snapshots link that corresponds to the server instance you are monitoring.

   The Snapshots Monitoring page opens.Generate thread snapshots, heap memory snapshots, JVM snapshots, process snapshot core files, GC snapshots, server log snapshots, access log snapshots, and configuration file snapshots. Snapshots can be automatically generated based on CPU and memory settings.

   The snapshot file is stored in ${instanceRoot}/snapshots/. It is generally named as [snapshot type]_[date].

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the Snapshots tab.

   The Snapshots Monitoring page opens.

4.11.10.1 Configure automatic snapshot generation

You can configure automatic snapshots generation by following options:

- CPU to Snap

  Generate snapshot CPU threshold, -1 indicates not to enable, and the default is -1. Generate JVM snapshot, with snapshot type being process snapshot.

- Memory to Snap

  Generate snapshot memory threshold, -1 indicates not to enable, the default is -1. Generate JVM snapshot, snapshot type is snapshot.

- Snap Interval

  The interval for detecting the threshold value, which is 60 minutes by default.

The modification will take effect directly after saving.

4.11.10.2 Thread Snapshot

Thread snapshots are generated as .txt files; click "Download" in the operation to download and view the file.

- Interval: Set the interval for thread collection, in seconds. The default value is 1 second.
- Check Times: Set the number of times for this sampling, which is 3 by default.

4.11.10.3 Heap Snapshot

The Heap Snapshot is generated as a .hprof file; click "Download" in the operation to download and view the file.

4.11.10.4 JVM Snapshot

The JVM Snapshot is generated as a .txt file; click "Download" in the operation to download and view the file.

4.11.10.5 Process Snapshot By Gcore

The Process Snapshot By Gcore button is generated as a .zip file; click "Download" in the operation to download and view the file. This feature requires gdb to be installed on the server.

4.11.10.6 GC Snapshot

The GC Snapshot is generated as a .txt file; click "Download" in the operation to download and view the file.

4.11.10.7 Server Log Snapshot

The Server Log Snapshot is generated as a .txt file; click "Download" in the operation to download and view the file.

4.11.10.8 Access Log Snapshot

Access Log Snapshots are generated as .zip files; click "Download" in the operation to download and view the file. This feature requires enabling "Access Log Record" in [Configuration] -> [Instance]-config -> [HTTP Service]. After enabling, the instance needs to be restarted.

4.11.10.9 Config File Snapshot

The Config File Snapshot is generated as a .txt file; click "Download" in the operation to download and view the file.

4.11.10.10 Performance Collection

It can monitor the CPU usage, collect data, and generate flame charts.

- Event

  Sampling based on CPU usage

- Duration

  The duration of the sampling, in seconds, with a default value of 60 seconds

- Type

  The format of the generated report, single choice

  - Flamegraph: Generate flame graph, document format instance name+timestamp.html
  - Jfr: Generate a jfr document in the format of instance name + timestamp.jfr

4.11.10.11 Delete Snapshot File

Select the snapshot file, click "Delete", and confirm the deletion before deleting the snapshot file.

4.11.11 To View Full-GC Data

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

1. In the navigation tree, select the Monitoring Data node.

   The Monitoring page opens.

2. In the View Monitoring Data column, click the Full-GC link that corresponds to the server instance you are monitoring.

   The Full-GC Monitoring page opens.The FGC monitoring page can obtain garbage collection information in real time, displaying the FullGC start time, end time, duration, details before collection, reason, details after collection, and cumulative collection times.

3. Click Start Trace button to start monitoring garbage collection information.

4. Click Stop Trace button will turn off monitoring of garbage collection information.

Alternative steps:

1. In the navigation tree, select the server (Admin Server) node.

   The General Information page opens.

2. On the General Information page, click the Monitor tab.

   The Application Monitoring page opens.

3. Click the Full-GC tab.

   The Full-GC Monitoring page opens.

4. Click Start Trace button to start monitoring garbage collection information.

5. Click Stop Trace button will turn off monitoring of garbage collection information.

4.12 Configurations

Use the Configurations page to manage configurations and view the standalone server instances and clusters using the configurations.

The `default-config` configuration is predefined by Apusic Application Server. The `server-config` configuration is automatically created for the domain administration server (DAS) when the domain is created.

For each configuration, the following information is provided.

- Name

  The name of the configuration.

- Instance

  The list of standalone instances and clusters using the configuration.

The Configurations table also contains the following options.

- New

  Button to create a new configuration.

- Delete

  Button to delete one or more selected configurations.

4.12.1 Admin Service

4.12.1.1 Edit JMX Connector

Use the Edit JMX Connector page to edit the JSR 360-compliant JMX connector.

The Edit JMX Connector page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the JMX connector. This is a read-only field.

- JMX Protocol

  The name of the protocol that this JMX connector supports. This is a read-only field.

- JMXServiceURL

  The JMX Service URL. This is a read-only field.

- Security

  If the Enabled checkbox is selected, JMX communication is encrypted. This option is disabled by default.

- Address

  The IP address of the naming service where the JMX connector server stub is registered. This is not the port of the server socket that does the actual JMX communication. This is the address of the network interface where the RMI registry is started. If your system has multiple network interfaces, modify this value so that only a particular interface is selected. The default value is `0.0.0.0`.

- Port

  The port number on which the naming service (RMI registry) listens for RMI client connections. The only use of this naming service is to download the RMI stubs. If the default port is occupied, a free port is used. Legal values are 1 - 65535. On UNIX systems, creating sockets that listen on ports 1 - 1024 requires superuser privileges. The default value is 8686.

- Realm Name

  The name that represents the special administrative realm. All authentication is handled by this realm. The default value is `admin-realm`.

- Additional Properties

Additional properties for the JMX Connector. Several properties are specified by default.

4.12.1.2 SSL

Use the SSL page to modify the SSL settings for the JMX connector.

The settings on this page are meaningful only if security is enabled for the JMX connector on the Edit JMX Connector page.

The SSL page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- SSL3

  If this checkbox is selected, the SSL3 protocol is enabled for the JMX connector. This option is enabled by default.

- TLS

  If this checkbox is selected, the TLS protocol is enabled for the JMX connector. This option is enabled by default.

- Client Authentication

  If this checkbox is selected, clients must identify themselves to the server on every request. This option is disabled by default.

- Certificate Nickname

  The nickname of the server certificate in the certificate database or the PKCS#11 token. In the certificate, the name format is *tokenname* `:` *nickname*. Including the *tokenname* `:` part of the name in this attribute is optional.

- Key Store

  The name of the keystore file (for example, `keystore.jks`).

- Trust Algorithm

  The name of the trust management algorithm (for example, PKIX) to use for certification path validation.

- Max Certificate Length

  The maximum number of non-self-issued intermediate certificates that can exist in a certification path. This field is used only if the Trust Algorithm field is set to PKIX. A value of 0 implies that the path can only contain a single certificate. A value of -1 implies that the path length is unconstrained (there is no maximum). Setting a value less than -1 causes an exception to be thrown.

- Trust Store

  The name of the truststore file (for example, `cacerts.jks`).

- Cipher Suites

  An area where you can add or remove cipher suites. If you do not add any cipher suites, all cipher suites will be used.

4.12.2 JVM Settings

4.12.2.1 JVM General Settings

The Java Virtual Machine (JVM) enables Java byte codes to run on a physical machine. Tuning the JVM settings improves performance and avoids memory allocation errors.

The Java Virtual Machine is included in the Java Platform, Standard Edition (Java SE platform), which is required by the Apusic Application Server. Incorrect JVM settings will prevent the server from running, so use care when changing these settings.

Note:

The terms "Java Virtual Machine" and "JVM" mean a Virtual Machine for the Java platform.

The JVM General Settings page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Java Home

  The name of the installation directory of the Java SE software, which is required for Apusic Application Server to run. If you enter a nonexistent directory name or the directory for an unsupported version of Java SE software, Apusic Application Server will not start.

- Javac Options

  The command-line options for the Java programming language compiler. The Apusic Application Server runs the compiler when Enterprise JavaBeans (EJB) components are deployed.

- Debug

  Enables debugging with the Java Platform Debugger Architecture (JPDA).JPDA is used by application developers.

- Debug Options

  Specifies the JPDA options passed to the JVM when debugging is enabled.

- RMI Compile Options

  The command-line options for the `rmic` compiler. The Apusic Application Server runs the `rmic` compiler when EJB components are deployed.

- Bytecode Preprocessor

  Names of classes for bytecode preprocessing. Use commas to separate class names. Each class must implement the com.sun.appserv.BytecodePreprocessor interface. The classes are called in the order specified.

- Additional Properties

  Additional properties for the JVM. The Apusic Application Server does not define any additional properties for the JVM.

4.12.2.2 JVM Options

Use the JVM Options page to specify the options of the Java application launcher (`java`) that runs the Apusic Application Server. The `-D` options designate properties that are specific to the Apusic Application Server.

- To modify an option, edit that option's Value field.
- To add an option, click the Add JVM Option button. In the blank row that appears, type the option information in the Value field.

  If the option information contains one or more spaces, enclose the entire string in double quotes.

- To delete an option, select the checkbox to the left of the Value field of the option to be deleted, then click Delete.

4.12.2.3 JVM Path Settings

Use the JVM Path settings page to configure the native library path. The native library path is a concatenation of the following paths:

- The native library path prefix
- The server's path for its native libraries (that is, *as-install* `/lib`)
- The 64-bit native library directories, if applicable
- The value of the `java.library.path` variable for the JVM used to start the Apusic Application Server (this differs between operating systems)
- Any paths that are specified on the JVM Profiler Settings page, if the profiler is enabled
- The native library path suffix

Other settings are shown as read-only values for domains that have been upgraded from an earlier release of Apusic Application Server. They are not supported in the v3 release.

The JVM Path Settings page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Environment Classpath

If the Ignore checkbox is selected, the `CLASSPATH` environment variable is ignored. The `CLASSPATH` environment variable is convenient for basic tutorials in programming, but is not recommended for enterprise environments. This is a read-only field.

- System Classpath

  The path to the classes on your system. This is a read-only field.

- Classpath Prefix

  Full paths to the JAR files to be prepended to the server's classpath, one per line. This is a read-only field.

- Classpath Suffix

  Full paths to the JAR files to be appended to the server's classpath, one per line. This is a read-only field.

- Native Library Path Prefix

  Full paths to the JAR files to be prepended to the native library path, one per line.

- Native Library Path Suffix

  Full paths to the JAR files to be appended to the native library path, one per line.

4.12.2.4 JVM Profiler Settings

A profiler tool generates data that is used to analyze performance and identify potential bottlenecks.

Specify profiler settings. The information that you specify depends on which profiler product you are using.

4.12.3 Thread Pools

Use the Thread Pools page to configure thread pools.

Use thread pools to limit a service to a specific number of concurrent threads.

When you first install the Apusic Application Server, two thread pools exist by default. The thread pool named `http-thread-pool` is configured for use by network listeners, while `thread-pool-1` is configured for use by the ORB for RMI/IIOP requests. You may need to create a thread pool with different settings for use by particular kinds of applications and for particular hardware systems, or to tune the default thread pools for your needs.

The Java Virtual Machine (JVM) can support many threads of execution at once. To help performance, theApusic Application Server maintains one or more thread pools. It is possible to assign specific thread pools to specific services.

One thread pool can serve multiple services. Request threads handle user requests. When the server receives a request, it assigns the request to a free thread from the thread pool. The thread executes the client's requests and returns results. For example, if the request needs to use a system resource that is currently busy, the thread waits until that resource is free before allowing the request to use that resource.

Specify the minimum and maximum number of threads that are reserved for requests. The thread pool is dynamically adjusted between these two values. The minimum thread pool size that is specified signals the server to allocate at least that many threads in reserve for requests. That number is increased up to the maximum thread pool size that is specified.

Increasing the number of threads available to a process allows the process to respond to more requests simultaneously.

For each thread pool, the following information is provided.

- Thread Pool ID

  The name of the thread pool.

- Class Name

  The class name of the thread pool.

- Max Thread Pool Size

  The maximum number of threads in the thread pool.

- Min Thread Pool Size

  The minimum number of threads in the thread pool. These threads are created when the thread pool is instantiated.

- Max Queue Size

The maximum number of messages that can be queued until threads are available to process them.

- Idle Thread Timeout

  The amount of time in seconds after which idle threads are removed from the pool.

- Thread Priority

  Set the priority level of the thread. The smaller the number, the higher the priority of loading.

The Thread Pools table also contains the following options.

- New

  Button to create a new thread pool.

- Delete

  Button to delete one or more selected thread pools.

4.12.3.1 New Thread Pool

Use the New Thread Pool page to create a thread pool.

The New Thread Pool page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the thread pool.

- Class Name

  The fully qualified name of the class that implements the thread pool. The default value is `com.sun.grizzly.http.StatsThreadPool`.

- Max Queue Size

  The maximum number of messages that can be queued until threads are available to process them. A value of -1 indicates that there is no limit to the queue size. The default value is 4096.

- Max Thread Pool Size

  The maximum number of threads in the thread pool. The default value is 5. This is the recommended value for network listener thread pools.

- Min Thread Pool Size

  The minimum number of threads in the thread pool. The default value is 2.

- Thread Priority

  Set the priority level of the thread. The smaller the number, the higher the priority of loading.

- Idle Thread Timeout

  The maximum amount of time, in seconds, that a thread can remain idle in the pool. After this time expires, the thread is removed from the pool. The default value is 900.

4.12.3.2 Edit Thread Pool

Use the Edit Thread Pool page to modify an existing thread pool.

The Edit Thread Pool page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the thread pool. The Name field is a read-only field. You can only specify the name when you create a new thread pool.

- Class Name

  The fully qualified name of the class that implements the thread pool. The default value is `com.sun.grizzly.http.StatsThreadPool`.

- Max Queue Size

  The maximum number of messages that can be queued until threads are available to process them. A value of -1 indicates that there is no limit to the queue size. The default value is 4096.

- Max Thread Pool Size

  The maximum number of threads in the thread pool. The default value is 5.

- Min Thread Pool Size

  The minimum number of threads in the thread pool. The default value is 2.

- Idle Thread Timeout

  The maximum amount of time, in seconds, that a thread can remain idle in the pool. After this time expires, the thread is removed from the pool. The default value is 900.

- Thread Priority

  Set the priority level of the thread. The smaller the number, the higher the priority of loading.

4.12.4 HTTP Service

Use the HTTP Service page to specify the general access log policy for web applications.

The HTTP service provides the facilities for deploying web applications and for making deployed web applications accessible by HTTP clients, along with virtual servers, thread pools, and the Network Configuration capability.

The Apusic Application Server uses the HTTP Service Access Log to enable and configure rotation for the access logs for the virtual servers. These logs are located in the *domain-dir*`/logs/access` directory and are named as follows: *virtual_server_name*`_access_log.` *yyyyMMdd*`-` *HH*`h` *mm*`m` *ss*`s.txt`

The HTTP Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- SSO

  Identifies whether single sign-on is enabled by default for all web applications on all virtual servers on this server that are configured for the same realm. If this option is not enabled, single sign-on is disabled by default for all virtual servers, and users must authenticate separately to every application on each virtual server. You can override this setting for an individual virtual server. SSO can be can be enabled or disabled. This option is disabled by default.

- Access Logging

  If this option is selected, access logging is enabled for all virtual server sub-elements that do not specify this property. If not selected, this option disables access logging for all virtual-server sub-elements that do not specify this property. This option is disabled by default.

- Rotation

  If this option is selected, log file rotation is enabled. This option is enabled by default.

- Rotation Policy

  The only available rotation policy is `time`.

- Rotation Interval

  The number of minutes between rotations of the access log. This field is valid only if the Rotation Policy is `time` . The default value is 1,440 minutes (24 hours).

- Rotation Suffix

  String value to be added to the end of the log file name after rotation. The default value is `yyyy-MM-dd` .Ensure that the rotation suffix contains enough values to create a unique suffix for every log rotation. For example, if the access log is rotated more frequently than once every day, include the time of day in the rotation suffix.A unique suffix is required to ensure that the access log file is rotated every time that the rotation interval has elapsed. If the rotation suffix of the new access log file is the same as the rotation suffix of the existing access log file, the file is not rotated.

- Max File Count

  The maximum number of rotated access log files that are to be kept. A negative value indicates that there is no limit to the file count. The default value is -1.

- Buffer Size

  The size, in bytes, of the access log buffer, or a value less than or equal to 0 for unbuffered access logs. The default value is 32768.

- Write Interval

  The number of seconds before the log will be written to the disk. The access log is written when the buffer is full or when the interval expires. If the value is less than or equal to 0, then the buffer is always written even if it is not full. This means that each time the server is accessed, the log message is stored directly to the file. The default value is 300.

- Log Compress Time

  The time point for compressing the access log every day. The format is 00:00.The default value is 00:03.

- Max Compress Files

  To reserve the maximum number of access log compressed files to be retained, enter 0 to reserve all compressed log files.The default value is 10.

- Log Compress Path

  Specify the path of the compressed file generated by daily log compression. When modifying this path, ensure that there is authorization.The default value is ${com.apusic.aas.instanceRoot}/logs/access.

- Format

  String value that specifies what information is to be captured in the access log and the order in which it is captured. The default value is as follows: `%client.name%` `%auth-user-name%` `%datetime%` `%request%` `%status%` `%response.length%` The following values are available for inclusion in the format string: `%attribute.` *name* `%` Logs the value (by calling its `toString` method) of the request attribute with the given *name* `%auth-user-name%` Name of authorized user `%client.name%` Client host name `%client.dns%` Client DNS `%cookie.` *name* `%` The value of the request cookie with the given *name* `%cookie.value%` The value of the first cookie found in the request `%datetime%` System date `%request%` Full HTTP request line `%status%` Status `%response.length%` Response content length `%header.referer%` Referer header `%header.user-agent%` User agent header `%http-method%` HTTP method `%http-uri%` HTTP URI `%query-str%` HTTP query string `%http-version%` HTTP protocol version `%header.accept%` Accept header `%header.date%` Date header `%header.if-mod-since%` If-Modified-Since header `%header.auth%` Authorization header `%header.` *any* `%` Any valid HTTP header value defined in RFC 2616 (`any` is also a valid header value; it is specified as a variable here) `%session.` *name* `%` The value of the session attribute with the name *name*, or `NULL-SESSION-ATTRIBUTE-` *name* if the named attribute does not exist in the session, or `NULL-SESSION` if no session exists `%time-taken%` Time (in milliseconds) it took to service each request `%vs.id%` Virtual server ID

- Additional Properties

Additional properties for the HTTP Service.The following additional properties are available for the HTTP service. If you specify these properties, they apply to all Network Listeners.

- `ssl-cache-entriese`

  The number of SSL sessions to be cached. The default value is 10,000.

- `ssl3-session-timeout`

  The interval at which SSL3 sessions are cached. The default value is 86,400.

- `ssl-session-timeout`

The interval at which SSL2 sessions are cached. The default value is 100.

- `proxyHandler`

  The fully qualified class name of a custom implementation of the `com.sun.appserv.ProxyHandler` abstract class, which allows a back-end Apusic Application Server instance to retrieve information about the original client request that was intercepted by an SSL-terminating proxy server (for example, a load balancer). An implementation of this abstract class inspects a given request for the custom request headers through which the proxy server communicates the information about the original client request to the Enterprise Server instance, and returns that information to its caller.The default implementation reads the client IP address from an HTTP request header named `Proxy-ip`, the SSL `keysize` from an HTTP request header named `Proxy-keysize`, and the SSL client certificate chain from an HTTP request header named `Proxy-auth-cert`. The `Proxy-auth-cert` value must contain the BASE-64 encoded client certificate chain without the `BEGIN CERTIFICATE` and `END CERTIFICATE` boundaries and with `\n` replaced with `% d% a`.This setting is used only if the Auth Pass Through checkbox for the network listener protocol's HTTP page is set to true.

- `connectionTimeout`

  Specifies the number of seconds that network listeners wait, after accepting a connection, for the request URI line to be presented. The default value is 30.

4.12.4.1 HTTP Listeners

Use the HTTP Listeners page to configure HTTP listeners.

An HTTP listener is a type of network listener. To access advanced features, you can edit an HTTP listener using the Network Listener pages.

When you first install the Apusic Application Server, three HTTP listeners exist by default, with the names `admin-listener`, `http-listener1`, and `http-listener2`. The `http-listener2`, which is configured with SSL, is disabled by default.

For each HTTP listener, the following information is provided.

- Name

  The name of the HTTP listener.Click the name,you can go to General(Edit Network Listener) page of the HTTP listener.

- Port

  The port number on which the HTTP listener is listening.

- Address

  The IP address used by the HTTP listener.

- Enabled

  True if the HTTP listener is enabled, or false if the HTTP listener is not enabled.

The HTTP Listeners table also contains the following options.

- New

  Button to create a new HTTP listener.

- Delete

  Button to delete one or more selected HTTP listeners.

4.12.4.1.1 NEW HTTP LISTENER

Use the New HTTP Listener page to create a new HTTP listener.

When you create an HTTP listener, the SSL tab is available for you to edit as well.

The New HTTP Listener page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique listener name. An HTTP listener name cannot begin with a number.

- Port

  The port number on which the listener will listen. Legal values are 1 through 65535. On a UNIX system, creating sockets that listen on ports 1 through 1024 requires superuser privileges.

- Status

  The status of the HTTP listener. The listener can be enabled or disabled. This option is enabled by default. If the listener is disabled, any attempts to connect to the listener result in a socket exception (`java.net.ConnectException`).In Apusic Application Server versions prior to 9.1, a listener whose enabled attribute was set to false returned a 404 response code for any requests sent to it. To achieve this behavior in the current Apusic Application Server version, set the listener's status to enabled, and set every associated virtual server's state to OFF. A virtual server lists its associated listeners in its Network Listeners drop-down list.

- Security

  If this option is selected, security is enabled for the HTTP listener. This option is disabled by default.If you are configuring the `admin-listener` for the domain administration server (DAS), which uses the `server-config` configuration, this option is read-only. To enable security for this listener, you must enable secure administration.

- JK Listener

  If this option is selected, the listener supports Apache's `mod-jk` protocol. The `mod-jk` protocol is a proprietary communication/network protocol between Apache's `httpd` (web server) and an Apusic Application Server back-end instance. If the listener supports this protocol, `httpd` is supposed to handle all static content, whereas any requests for dynamic resources (i.e., Servlets and JSPs) are routed to the `mod-jk` enabled HTTP listener of the Apusic Application Server, using the `mod-jk` protocol. This option is disabled by default.

- Address

  The IP address on which the network listener will listen. The address can be in dotted-pair or IPv6 notation. It can be `any` (for `INADDR_ANY`) to listen on all IP addresses. It can be a hostname.

- Default Virtual Server

  The virtual server to be associated with this HTTP listener. Use the Virtual Servers page to define virtual servers.

- Thread Pool

  The thread pool associated with the HTTP listener. Normally, you select one of the two thread pools that are configured when you install the Apusic Application Server.

- Server Name

  The host name to be used in the URLs the server sends to the client. This name is the alias name if your server uses an alias. If your server does not use an alias, leave this field blank.This value affects URLs the server automatically generates; it does not affect the URLs for directories and files stored in the server. If your server uses an alias, the server-name should be the alias name. If a colon and port number are appended, that port is used in URLs the server sends to the client.

4.12.4.1.2 EDIT HTTP LISTENER

Use the Edit HTTP Listener page to modify an existing HTTP listener.

When you edit an HTTP listener, the SSL tab is available for you to edit as well.

The Edit HTTP Listener page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The unique listener name. The Name field is a read-only field. You can only specify the name when you create a new HTTP listener.

- Port

The port number on which the listener will listen. Legal values are 1 through 65535. On a UNIX system, creating sockets that listen on ports 1 through 1024 requires superuser privileges.

- Status

  The status of the HTTP listener. The listener can be enabled or disabled. This option is enabled by default. If the listener is disabled, any attempts to connect to the listener result in a socket exception (`java.net.ConnectException`).In Apusic Application Server versions prior to 9.1, a listener whose enabled attribute was set to false returned a 404 response code for any requests sent to it. To achieve this behavior in the current Apusic Application Server version, set the listener's status to enabled, and set every associated virtual server's state to OFF. A virtual server lists its associated listeners in its Network Listeners drop-down list.

- Security

  If this option is selected, security is enabled for the HTTP listener. This option is disabled by default.If you are configuring the `admin-listener` for the domain administration server (DAS), which uses the `server-config` configuration, this option is read-only. To enable security for this listener, you must enable secure administration.

- JK Listener

  If this option is selected, the listener supports Apache's `mod-jk` protocol. The `mod-jk` protocol is a proprietary communication/network protocol between Apache's `httpd` (web server) and an Apusic Application Server back-end instance. If the listener supports this protocol, `httpd` is supposed to handle all static content, whereas any requests for dynamic resources (i.e., Servlets and JSPs) are routed to the `mod-jk` enabled HTTP listener of the Apusic Application Server, using the `mod-jk` protocol. This option is disabled by default.

- Address

  The IP address on which the network listener will listen. The address can be in dotted-pair or IPv6 notation. It can be `any` (for `INADDR_ANY`) to listen on all IP addresses. It can be a hostname.

- Default Virtual Server

  The virtual server to be associated with this HTTP listener. Use the Virtual Servers page to define virtual servers.

- Thread Pool

  The thread pool associated with the HTTP listener.

- Server Name

  The host name to be used in the URLs the server sends to the client. This name is the alias name if your server uses an alias. If your server does not use an alias, leave this field blank.This value affects URLs the server automatically generates; it does not affect the URLs for directories and files stored in the server. If your server uses an alias, the server-name should be the alias name. If a colon and port number are appended, that port is used in URLs the server sends to the client.

4.12.5 Network Config

Use the Network Config page to configure objects of the following types:

- Network Listeners
- Protocols
- Transports

4.12.5.1 Network Listeners

Use the Network Listeners page to configure network listeners.

When you first install the Apusic Application Server, three network listeners exist by default, with the names `admin-listener` , `http-listener1` , and `http-listener2` . The `http-listener2` , which is configured with SSL, is disabled by default.

For each network listener, the following information is provided.

- Name

  The name of the network listener.Click the name,you can go to General(Edit Network Listener) page of the Network listener.

- Port

  The port number on which the network listener is listening.

- Protocol

The protocol used by the network listener.Click the name,you can go to Protocol(Edit Network Listener) page of the network listener.

- Thread Pool

  The thread pool used by the network listener.Click the name,you can go to Edit Thread Pool page of the network listener.

- Enabled

  True if the network listener is enabled, or false if the network listener is not enabled.

The Network Listeners table also contains the following options.

- New

  Button to create a new network listener.

- Delete

  Button to delete one or more selected network listeners.

4.12.5.1.1 NEW NETWORK LISTENER

Use the New Network Listener page to create a new network listener.

When you create a network listener, you can simultaneously create a protocol that the network listener will use, or use a protocol that you create first. Typically, there is a one-to-one relationship between a network listener and its associated protocol, but this is not required.

The New Network Listener page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique listener name. A network listener name cannot begin with a number.

- Protocol

  The protocol associated with the network listener. The following options specify whether to create a protocol or use an existing one.Create a New ProtocolA new protocol with default settings is created for the network listener. The default name of the protocol is the name of the network listener with `-protocol` appended, but you can edit this value.Virtual ServerFor a new protocol, the virtual server to be associated with the protocol.Use an Existing ProtocolThe selected protocol is used. The protocol determines whether security is enabled for the network listener.

- Status

  The status of the network listener. The listener can be enabled or disabled. This option is enabled by default. If the listener is disabled, any attempts to connect to the listener result in a socket exception (`java.net.ConnectException`).In Apusic Application Server versions prior to 9.1, a listener whose enabled attribute was set to false returned a 404 response code for any requests sent to it. To achieve this behavior in the current Apusic Application Server version, set the listener's status to enabled, and set every associated virtual server's state to OFF. A virtual server lists its associated listeners in its Network Listeners drop-down list.

- Security

  If this option is selected, security is enabled for the protocol used by the network listener. The Security field is a read-only field.If you are configuring the `admin-listener` for the domain administration server (DAS), which uses the `server-config` configuration, this option is read-only. To enable security for this listener, you must enable secure administration.

- JK Listener

  If this option is selected, the listener supports Apache's `mod-jk` protocol. The `mod-jk` protocol is a proprietary communication/network protocol between Apache's `httpd` (web server) and an Apusic Application Server back-end instance. If the listener supports this protocol, `httpd` is supposed to handle all static content, whereas any requests for dynamic resources (i.e., Servlets and JSPs) are routed to the `mod-jk` enabled network listener of the Apusic Application Server, using the `mod-jk` protocol. This option is disabled by default.

- Port

The port number on which the listener will listen. Legal values are 1 through 65535. On a UNIX system, creating sockets that listen on ports 1 through 1024 requires superuser privileges.

- Address

  The IP address on which the network listener will listen. The address can be in dotted-pair or IPv6 notation. It can be `any` (for `INADDR_ANY`) to listen on all IP addresses. It can be a hostname.

- Thread Pool

  The thread pool associated with the network listener. Normally, you select one of the two thread pools that are configured when you install the Apusic Application Server.

- Transport

  The type of transport for the network listener. By default, the only choice is `tcp`.

4.12.5.1.2 EDIT NETWORK LISTENER

Use the Edit Network Listener page to modify an existing network listener.

When you edit a network listener, the SSL, HTTP, HTTP2,and File Cache tabs for the associated protocol are available for you to edit as well.

The Edit Network Listener Genneral page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The unique listener name. The Name field is a read-only field. You can only specify the name when you create a new network listener.

- Protocol

  The protocol associated with the network listener. The Protocol field is a read-only field. You can create or specify a protocol only when you create a new network listener.Select the Protocol used by the network listener to change its settings. For example, you can enable security.

- Status

  The status of the network listener. The listener can be enabled or disabled. This option is enabled by default. If the listener is disabled, any attempts to connect to the listener result in a socket exception (`java.net.ConnectException`).In Apusic Application Server versions prior to Application Server 9.1, a listener whose enabled attribute was set to false returned a 404 response code for any requests sent to it. To achieve this behavior in the current Apusic Application Server version, set the listener's status to enabled, and set every associated virtual server's state to OFF. A virtual server lists its associated listeners in its Network Listeners drop-down list.

- Security

  If this option is selected, security is enabled for the protocol used by the network listener. The Security field is a read-only field.If you are configuring the `admin-listener` for the domain administration server (DAS), which uses the `server-config` configuration, this option is read-only. To enable security for this listener, you must enable secure administration.

- JK Listener

  If this option is selected, the listener supports Apache's `mod-jk` protocol. The `mod-jk` protocol is a proprietary communication/network protocol between Apache's `httpd` (web server) and an Apusic Application Server back-end instance. If the listener supports this protocol, `httpd` is supposed to handle all static content, whereas any requests for dynamic resources (i.e., Servlets and JSPs) are routed to the `mod-jk` enabled network listener of the Apusic Application Server, using the `mod-jk` protocol. This option is disabled by default.

- Port

  The port number on which the listener will listen. Legal values are 1 through 65535. On a UNIX system, creating sockets that listen on ports 1 through 1024 requires superuser privileges.

- Address

The IP address on which the network listener will listen. The address can be in dotted-pair or IPv6 notation. It can be `any` (for `INADDR_ANY`) to listen on all IP addresses. It can be a hostname.

- Thread Pool

  The thread pool associated with the network listener.

- Transport

  The type of transport for the network listener. By default, the only transport is `tcp`.

4.12.5.2 Protocols

Use the Protocols page to configure network protocols.

For each protocol, the following information is provided.

- Name

  The name of the protocol.

- Security Enabled

  True if security is enabled for the protocol, or false if security is not enabled for the protocol.

The Protocols table also contains the following options.

- New

  Button to create a new protocol.

- Delete

  Button to delete one or more selected protocols.

4.12.5.2.1 NEW PROTOCOL

Use the New Protocol page to create a new network protocol.

After you create a protocol, you normally create a network listener with the same name, and you associate the network listener with the protocol. Alternatively, you can create a network listener and have a protocol with default settings created for you at the same time.

The New Protocol page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The unique name of the protocol. A protocol name cannot begin with a number.

- Security

  If this option is selected, security is enabled for the protocol. This option is disabled by default. If you enable security, you can use the SSL tab of the Edit Protocol page to modify the security settings.

- Status

  If this option is selected, the file cache is enabled. This option is enabled by default.

- Max Age

  The maximum age, in seconds, for a valid cache entry.This parameter controls how long cached information is used after a file has been cached. An entry older than the maximum age is replaced by a new entry for the same file. If your content changes infrequently, increase this value for improved performance.The optimum maximum age depends on whether existing files are modified regularly. For example, if files are modified four times a day at regular intervals, consider setting the maximum age to 21600 seconds (6 hours). Otherwise, consider setting the maximum age to the longest time for which you are willing to serve the previous version of a file after the file has been modified.The default value is 30.

- Max Cache Size

The maximum total size (in bytes) of the file cache on disk. The default value is 10485760.

- Max File Count

  The maximum number of files that can be stored in the file cache.If the value is too big, the server caches little-needed files, which wastes memory. If the value is too small, the benefit of caching is lost. Try different values of this attribute to find the optimal solution for specific applications.The default value is 1024.

- Server Name

  The host name to be used in the URLs the server sends to the client. This name is the alias name if your server uses an alias. If your server does not use an alias, leave this field blank.This value affects URLs the server automatically generates; it does not affect the URLs for directories and files stored in the server. If your server uses an alias, the server-name should be the alias name. If a colon and port number are appended, that port is used in URLs the server sends to the client.

- Default Virtual Server

  The virtual server to be associated with this protocol. Use the Virtual Servers page to define virtual servers.

- Redirect Port

  Port value that redirects a request to another port. The Apusic Application Server automatically redirects the request if these two conditions exist:The network listener that references this protocol is supporting non-SSL requests.A request is received for which a matching security constraint requires SSL transport.If a redirect port is not specified, the Apusic Application Server uses the port number specified in the original request. By default, no port is specified.

- Max Connections

  Specifies the maximum number of requests that can be pipelined until the connection is closed by the server. The Keep-Alive subsystem periodically polls idle connections. The default value is 256.Set this property to 1 to disable HTTP/1.0 keep-alive, as well as HTTP/1.1 keep-alive and pipelining. A value of 0 means requests are always rejected. A value of -1 sets no limit to the number of keep-alive connections.

- Timeout

  The maximum time in seconds that a connection can be deemed as idle and kept in the keep-alive state. A value of 0 or less means that keep-alive connections are kept open indefinitely. The default value is 30 seconds.

- Upload Timeout

  If this option is selected, the connection for a servlet that reads bytes slowly is closed after the Connection Upload Timeout limit is reached. If this option is disabled, servlet connections do not time out. This option is disabled by default.

- Connection Upload Timeout

  The timeout for uploads, in milliseconds. This field is applicable only if the Upload Timeout Enabled checkbox is selected. The default value is 300000 milliseconds.

- Request Timeout

  The number of seconds before a request times out. If the request is not processed before the timeout value is reached, the request is ignored. The default value is 30 seconds.

- Send Buffer Size

  The size in bytes of the send buffer. The default value is 8192 bytes.

- Header Buffer Length

  The size in bytes of the buffer used by the request processing threads to read the request data. The default value is 8192 bytes.

- Max Post Size

  The maximum size in bytes of POST actions. The default value is 2097152 bytes.

- Max Form Post Size Bytes

  The Maximum size of Form POST actions.-1 means unlimited,the default is -1.

- Max Save Post Size

  The maximum size of a POST which will be saved by the container during authentication.The default is 4096 bytes.

- URI Encoding

The name of the character set used to decode the request URIs received. The value must be a valid IANA character set name. The default value is UTF-8.

- Version

  The version of the HTTP protocol used. The default value is HTTP/1.1.

- Compression

  Specifies the use of HTTP/1.1 GZIP compression to save server bandwidth. Available choices are: `on` Compresses data. `off` Disables compression. `force` Forces data compression in all cases.The default value is `off` .

- Compressible Mime Types

  A comma-separated list of MIME types for which HTTP compression is used. The default value is `text/html,text/xml,text/plain` .

- Compression Minimum Size

  The minimum size of a file when compression is applied. This value must be set if Compression is set to `on` or `force` . The default value is 2048 bytes.

- No-Compression User Agents

  A comma-separated list of regular expressions matching user-agents of HTTP clients for which compression should not be used. By default, this value is an empty string.

- Restricted User Agent

  A list of restricted user agents on which HTTP compression is applied. If no user agents are specified, HTTP compression is applied to all user agents. By default, no user agents are specified.

- Default Response Type

  A string that specifies the default response type. The format is a semicolon-delimited string consisting of the content-type, encoding, language, and charset. The default value is `text/html; charset=iso-8859-1` .

- Forced Response Type

  A string that specifies the request type used if no MIME mapping is available that matches the file extension. The format is a semicolon-delimited string consisting of the content-type, encoding, language, and charset. The default value is `text/html; charset=iso-8859-1` .

- Adapter

  The class name of the static resources adapter. The default value is `com.sun.grizzly.tcp.StaticResourcesAdapter` .

- Comet Support

  If this option is selected, Comet support is enabled for the network listener that references this protocol. By default, this option is disabled. If your servlet or JSP page uses Comet technology, make sure it is initialized when the Apusic Application Server starts up by adding the `load-on-startup` element to your `web.xml` file. For example: `<servlet> <servlet-name>CheckIn</servlet-name> <servlet-class>CheckInServlet</servlet-class> <load-on-startup>0</load-on-startup> </servlet>`

- DNS Lookup

  If this option is selected, DNS lookup is enabled. This option is disabled by default.

- RCM Support

  If this option is selected, Resource Configuration Management (RCM) is enabled. This option is disabled by default.

- Allow PayLoad for Get Head Delete

  If this option is selected,allow payload for Get Head Delete. This option is enabled by default.

- Disabled Http Methods

  Which Http Methods are disabled,TARCE,OPTIONS,HEAD,PUT,and DELETE.The default value is TRACE.

- Allow Custom Methods

  If this option is selected,all the custom methods are disabled.This option is enabled by default.

- Custom Methods

  Costomize HTTP method what you want to enabled.Multiple separated by colons: .

- Relaxed Query Chars

  Set the special characters allowed to access the URL. Check the "Special Query Chars" to complete some special characters.

- Auth Pass Through

  If this option is selected, it indicates that the network listener that uses this protocol receives traffic from an SSL-terminating proxy server. This option is disabled by default.

- Chunking

  If this option is selected, HTTP response chunking is enabled. This option is enabled by default.

- XPowered By

  If this option is selected, `X-Powered-By` headers are used according to the Java Servlet 3.0 and Java Server Pages 2.0 specifications. This option is enabled by default.

- Encoded Slash

  If this option is selected, URIs are permitted to contain encoded slashes. This option is disabled by default.

- Websockets Support

  If this option is selected, the WebSockets protocol is supported. This option is disabled by default.

- WebsocketsTimeoutSeconds

  Set the maximum idle time for websocket connections.

- Max Request Parameters

  Set maximum request parameter size.

4.12.5.2.2 EDIT PROTOCOL

Use the Edit Protocol page to modify the settings of an existing network protocol.

The Edit Protocol page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the protocol. The Name field is a read-only field. You can only specify the Name when you create a new protocol.

- Security

  If this option is selected, security is enabled for the protocol. If security is enabled, you can use the SSL tab to modify the security settings.If you are configuring the `admin-listener` for the domain administration server (DAS), which uses the `server-config` configuration, this option is read-only. To enable security for this protocol, you must enable secure administration.

4.12.5.2.3 SSL

Use the SSL page to modify SSL settings for a network protocol.

The settings on this page are meaningful only if security is enabled on the Edit Protocol page.

The SSL page contains the following options.

- Configuration Name

The name of the configuration to which the settings on this page apply. This field is read only.

- SM11Enabled

  Whether to enable the configuration of SM(ShangMi). Before enabling, it is necessary to complete the configuration of SM certificate information and restart AAS to take effect.

- SSL3

  If this checkbox is selected, the SSL3 protocol is enabled for the protocol. This option is disabled by default.

- TLS

  If this checkbox is selected, the TLS protocol is enabled for the protocol. This option is disabled by default.

- TLS1.1

  If this checkbox is selected, the TLS1.1 protocol is enabled for the protocol. This option is disabled by default.

- TLS1.2

  If this checkbox is selected, the TLS1.2 protocol is enabled for the protocol. This option is enabled by default.

- TLS1.3

  If this checkbox is selected, the TLS1.3 protocol is enabled for the protocol. This option is enabled by default.

- Client Authentication

  If this checkbox is selected, clients must identify themselves to the server on every request. This option is disabled by default.

- Certificate Nickname

  The nickname of the server certificate in the certificate database or the PKCS#11 token. In the certificate, the name format is *tokenname* `:` *nickname*. Including the *tokenname* `:` part of the name in this attribute is optional.

- Key Store

  The name of the keystore file (for example, `keystore.jks`).

- Trust Algorithm

  The name of the trust management algorithm (for example, PKIX) to use for certification path validation.

- Max Certificate Length

  The maximum number of non-self-issued intermediate certificates that can exist in a certification path. This field is used only if the Trust Algorithm field is set to PKIX. A value of 0 implies that the path can only contain a single certificate. A value of -1 implies that the path length is unconstrained (there is no maximum). Setting a value less than -1 causes an exception to be thrown.

- Trust Store

  The name of the truststore file (for example, `cacerts.jks`).

- Cipher Suites

  An area where you can add or remove cipher suites. If you do not add any cipher suites, all cipher suites will be used.

4.12.5.2.4 HTTP

Use the HTTP page to modify HTTP settings for a network protocol.

The HTTP page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Protocol Name

  The name of the protocol. The Protocol Name field is a read-only field. You can only specify the name when you create a new protocol.

- Server Name

  The host name to be used in the URLs the server sends to the client. This name is the alias name if your server uses an alias. If your server does not use an alias, leave this field blank.

- Default Virtual Server

  The virtual server to be associated with the protocol. Use the Virtual Servers page to define virtual servers.

- Redirect Port

  Port value that redirects a request to another port. The Apusic Application Server automatically redirects the request if these two conditions exist:This protocol is supporting non-SSL requests.A request is received for which a matching security constraint requires SSL transport.If a redirect port is not specified, the Apusic Application Server uses the port number specified in the original request.

- Max Connections

  The maximum number of persistent connections to be maintained in Keep-Alive mode. The Keep-Alive subsystem periodically polls idle connections. The default value is 256.

- Timeout

  The maximum time in seconds that a connection can be considered idle and kept in the keep-alive state. A value of 0 or less means that keep-alive connections are kept open indefinitely. The default value is 30 seconds.A value of -1 disables this timeout.

- Upload Timeout

  If this option is selected, the connection for a servlet that reads bytes slowly is closed after the Connection Upload Timeout limit is reached. If this option is disabled, servlet connections do not time out. This option is disabled by default.

- Connection Upload Timeout

  The timeout for uploads, in milliseconds. This field is applicable only if the Upload Timeout Enabled checkbox is selected. The default value is 300000 milliseconds.A value of -1 disables this timeout.

- Request Timeout

  The number of seconds before a request times out. If the request is not processed before the timeout value is reached, the request is ignored. The default value is 30 seconds.A value of -1 disables this timeout.

- Send Buffer Size

  The size in bytes of the send buffer. The default value is 8192 bytes.

- Header Buffer Length

  The size in bytes of the buffer used by the request processing threads to read the request data. The default value is 8192 bytes.

- Max Post Size

  The maximum size in bytes of POST actions using application/x-www-form-urlencoded. The default value is 2097152 bytes.

- Max Form Post Size Bytes

  The Maximum size of Form POST actions.-1 means unlimited,the default is -1.

- Max Save Post Size

  The maximum size of a POST which will be saved by the container during authentication.The default is 4096 bytes.

- URI Encoding

  The name of the character set used to decode the request URIs received. The value must be a valid IANA character set name. The default value is UTF-8.

- Version

The version of the HTTP protocol used. The default value is HTTP/1.1.

- Compression

  Specifies the use of HTTP/1.1 GZIP compression to save server bandwidth. Available choices are: `on` Compresses data. `off` Disables compression. `force` Forces data compression in all cases.The default value is `off`.

- Compressible Mime Type

  A comma-separated list of MIME types for which HTTP compression is used. The default value is `text/html,text/xml,text/plain`.

- Compression Minimum Size

  The minimum size of a file when compression is applied. This value must be set if Compression is set to `on` or `force`. The default value is 2048 bytes.

- No-Compression User Agents

  A comma-separated list of regular expressions matching user agents of HTTP clients for which compression should not be used. By default, this value is an empty string.

- Restricted User Agent

  A list of restricted user agents on which HTTP compression is applied. If no user agents are specified, HTTP compression is applied to all user agents. By default, no user agents are specified.

- Default Response Type

  A string that specifies the default response type. The format is a semicolon-delimited string consisting of the content-type, encoding, language, and charset. The default value is `text/html; charset=iso-8859-1`.

- Forced Response Type

  A string that specifies the request type used if no MIME mapping is available that matches the file extension. The format is a semicolon-delimited string consisting of the content-type, encoding, language, and charset. The default value is `text/html; charset=iso-8859-1`.

- Adapter

  The class name of the static resources adapter. The default value is `com.sun.grizzly.tcp.StaticResourcesAdapter`.

- Comet Support

  If this option is selected, Comet support is enabled for the protocol. This option is disabled by default.

- DNS Lookup

  If this option is selected, Domain Name System (DNS) lookup is enabled. This option is disabled by default.

- RCM Support

  If this option is selected, Resource Configuration Management (RCM) is enabled. This option is disabled by default.

- Allow PayLoad for Get Head Delete

  If this option is selected,allow payload for Get Head Delete. This option is enabled by default.

- Disabled Http Methods

  Which Http Methods are disabled,TARCE,OPTIONS,HEAD,PUT,and DELETE.The default value is TRACE.

- Allow Custom Methods

  If this option is selected,all the custom methods are disabled.This option is enabled by default.

- Custom Methods

  Costomize HTTP method what you want to enabled.Multiple separated by colons: .

- Relaxed Query Chars

  Set the special characters allowed to access the URL. Check the "Special Query Chars" to complete some special characters.

- Auth Pass Through

  If this option is selected, it indicates that the network listener that uses this protocol receives traffic from an SSL-terminating proxy server. This option is disabled by default.

- Chunking

  If this option is selected, HTTP response chunking is enabled. This option is enabled by default.

- XPowered By

  If this option is selected, `X-Powered-By` headers are used according to the Java Servlet 3.0 and Java Server Pages 2.0 specifications. This option is enabled by default.

- Encoded Slash

  If this option is selected, URIs are permitted to contain encoded slashes. This option is disabled by default.

- Websockets Support

  If this option is selected, the WebSockets protocol is supported. This option is disabled by default.

- WebsocketsTimeoutSeconds

  Set the maximum idle time for websocket connections.

- Max Request Parameters

  Set maximum request parameter size.

- Header Options

  Set the header options for the protocol needs.

4.12.5.2.5 HTTP2

Use the HTTP2 page to modify HTTP2 settings for a network protocol.

The HTTP2 page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- HTTP2 Enabled

  If this option is selected,the HTTP/2 protocol is supported.Disabled by default.

- Http2MaxConcurrentStreams

  Control the maximum number of streams allowed for any connection.The default is 100.

- http2InitialWindowSizeInBytes

  The initial size of the flow control window for the stream sent by the control server to the client. The default value is 64*1024 - 1,65535 bytes.

- http2MaxFramePayloadSizeInBytes

  The size of data that can be carried in each frame.The default is 16777215 bytes.

- http2MaxHeaderListSizeInBytes

  The number of request headers that can be carried in each request.The default is 4096 bytes.

- http2DisableCipherCheck

  If this option is selected,enable cipher suites check.Disabled by default.

- Header Options

  Set the header options for the protocol needs.

After enabled HTTP/2,the HTTP/2 protocol will be used preferentially.Through the network of browser, you can view the URL protocol as `h2`.

4.12.5.2.6 FILE CACHE

Use the File Cache page to modify file cache settings for a network protocol.

The File Cache page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Protocol Name

  The name of the protocol. The Protocol Name field is a read-only field. You can only specify the name when you create a new protocol.

- Status

  If this option is selected, the file cache status is enabled. This option is selected by default.

- Max Age

  The maximum age, in seconds, for a valid cache entry.This parameter controls how long cached information is used after a file has been cached. An entry older than the maximum age is replaced by a new entry for the same file. If your content changes infrequently, increase this value for improved performance.The optimum maximum age depends on whether existing files are modified regularly. For example, if files are modified four times a day at regular intervals, consider setting the maximum age to 21600 seconds (6 hours). Otherwise, consider setting the maximum age to the longest time for which you are willing to serve the previous version of a file after the file has been modified.The default value is 30.

- Max Cache Size

  The maximum total size (in bytes) of the file cache on disk. The default value is 10485760.

- Max File Count

  The maximum number of files that can be stored in the file cache.If the value is too big, the server caches little-needed files, which wastes memory. If the value is too small, the benefit of caching is lost. Try different values of this attribute to find the optimal solution for specific applications.The default value is 1024.

4.12.5.3 Transports

Use the Transports page to configure transports. Each network listener is associated with a transport.

When you first install the Apusic Application Server, one TCP transport exists by default, with the name `tcp`. You can modify or create a TCP or UDP transport or provide a custom transport implementation.

For each transport, the following information is provided.

- Transport Name

  The name of the transport.

- Classname

  The class name of the transport implementation.

- Byte Buffer Type

  The type of `ByteBuffer` to be used with the transport.

The Transports table also contains the following options.

- New

Button to create a new transport.

- Delete

  Button to delete one or more selected transports.

4.12.5.3.1 NEW TRANSPORT

Use the New Transport page to create a new transport.

It is not common to need another transport in addition to the default `tcp` transport.

The New Transport page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Transport Name

  A unique name for the transport.

- Classname

  The class name of the transport implementation. The default value is `com.sun.grizzly.TCPSelectorHandler`.

- Selection Key Handler

  The class name of the selection key handler. A selection key handler is an abstract class that manages the lifecycle of selection keys. If you do not specify a value, a default implementation is used.

- Byte Buffer Type

  The type of `ByteBuffer` to be used. Available choices are `HEAP` and `DIRECT`. The default value is `HEAP`.

- Acceptor Threads

  The number of processors in the machine. A value of -1 indicates that Grizzly will calculate the number of acceptor threads itself, based on the number of processors in the machine. The default value is 1.To set the number of request processing threads, set the Max Thread Pool Size value of the thread pool used by a network listener that uses this transport.

- Max Connections Count

  The maximum number of pending connections on a network listener that uses this transport. The default value is 4096.

- Buffer Size

  The size, in bytes, of the buffer to be provided for input streams created by a network listener that uses this transport. The default value is 8192.

- Idle Key Timeout

  The number of seconds after which an idle key will be cancelled and the channel closed. The default value is 30.

- Read Timeout

  The number of milliseconds the Apusic Application Server waits during the header and body parsing phase of a read operation. The default value is 30,000.

- Selector Poll Timeout

  The number of milliseconds an NIO selector will block waiting for events (user requests). The default value is 1000.

- Write Timeout

  The number of milliseconds the Apusic Application Server waits before considering the remote client disconnected when writing the response. The default value is 30,000.

- IO Strategy

  Set the IO strategy, `org.Apusic Application.grizzly.strategies.WorkerThreadIOStrategy`, `org.Apusic Application.grizzly.strategies.SameThreadIOStrategy`, `org.Apusic Application.grizzly.strategies.LeaderFollowerNIOStrategy`.The default value is `org.Apusic Application.grizzly.strategies.WorkerThreadIOStrategy`.

- Display Configuration

  If this option is selected, Grizzly's internal configuration is flushed to the server logs. This option may provide useful information for debugging. This option is disabled by default.

- Snoop

  If this option is selected, the request/response information is dumped to the server log. This option may provide useful information for debugging, but will significantly reduce performance. This option is disabled by default.

- Slow HTTP Detect

  If this option is selected,enable slow HTTP attack detection, and automatically close the timeout path when the connection exceeds the set timeout.Disabled by default.

- Slow HTTP Detect Debug

  If this option is selected,enable slow HTTP detect debug ,output slow HTTP attack logs.Disabled by default.

- Slow HTTP Write Count

  Connecting to write N times, and if the data written is less than the configured minimum data, it is determined to be a slow HTTP attack.The default value is 5 times.

- Slow HTTP Write Min Bytes

  Connecting to write N times, and if the data written is less than the value, it is judged as a slow HTTP attack.The default value is 128 bytes.

4.12.5.3.2 EDIT TRANSPORT

Use the Edit Transport page to modify the settings of an existing transport.After modifying the information, you need to restart the instance.

The Edit Transport page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the transport. The name is a read-only field. You can only specify a name when you create a new transport.

- Classname

  The class name of the transport implementation. The default value is `com.sun.grizzly.TCPSelectorHandler`.

- Selection Key Handler

  The class name of the selection key handler. A selection key handler is an abstract class that manages the lifecycle of selection keys. If you do not specify a value, a default implementation is used.

- Byte Buffer Type

  The type of `ByteBuffer` to be used. Available choices are `HEAP` and `DIRECT`. The default value is `HEAP`.

- Acceptor Threads

  The number of processors in the machine. A value of -1 indicates that Grizzly will calculate the number of acceptor threads itself, based on the number of processors in the machine. The default value is 1.To set the number of request processing threads, set the Max Thread Pool Size value of the thread pool used by a network listener that uses this transport.

- Max Connections Count

  The maximum number of pending connections on a network listener that uses this transport. The default value is 4096.

- Buffer Size

The size, in bytes, of the buffer to be provided for input streams created by a network listener that uses this transport. The default value is 8192.

- Idle Key Timeout

  The number of seconds after which an idle key will be cancelled and the channel closed. The default value is 30.

- Read Timeout

  The number of milliseconds the Apusic Application Serverwaits during the header and body parsing phase of a read operation. The default value is 30,000.

- Selector Poll Timeout

  The number of milliseconds an NIO Selector will block waiting for events (users' requests). The default value is 1000.

- Write Timeout

  The number of milliseconds the Apusic Application Serverwaits before considering the remote client disconnected when writing the response. The default value is 30,000.

- Display Configuration

  If this option is selected, Grizzly's internal configuration is flushed to the server logs. This option may provide useful information for debugging. This option is disabled by default.

- Snoop

  If this option is selected, the requests/response information is dumped to the server log. This option may provide useful information for debugging, but will significantly reduce performance. This option is disabled by default.

- IO Strategy

  Set the IO strategy, `org.Apusic Application.grizzly.strategies.WorkerThreadIOStrategy`, `org.Apusic Application.grizzly.strategies.SameThreadIOStrategy`, `org.Apusic Application.grizzly.strategies.LeaderFollowerNIOStrategy`.The default value is `org.Apusic Application.grizzly.strategies.WorkerThreadIOStrategy`.

- Display Configuration

  If this option is selected, Grizzly's internal configuration is flushed to the server logs. This option may provide useful information for debugging. This option is disabled by default.

- Snoop

  If this option is selected, the request/response information is dumped to the server log. This option may provide useful information for debugging, but will significantly reduce performance. This option is disabled by default.

- Slow HTTP Detect

  If this option is selected,enable slow HTTP attack detection, and automatically close the timeout path when the connection exceeds the set timeout.Disabled by default.

- Slow HTTP Detect Debug

  If this option is selected,enable slow HTTP detect debug ,output slow HTTP attack logs.Disabled by default.

- Slow HTTP Write Count

  Connecting to write N times, and if the data written is less than the configured minimum data, it is determined to be a slow HTTP attack.The default value is 5 times.

- Slow HTTP Write Min Bytes

  Connecting to write N times, and if the data written is less than the value, it is judged as a slow HTTP attack.The default value is 128 bytes.

4.12.6 Logger Settings

4.12.6.1 General

Use the Logger General Settings page to configure logging for the selected Apusic Application Server instance or cluster.

The Logger Settings page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Write to System Log

  When enabled, logging output is sent to the `syslog` facility in addition to the server log. This feature is available on Solaris and Linux systems only. This option is disabled by default.

- Log to Console

  When enabled, logging output is sent to the system console. This option is disabled by default.

- Rotation on Date Change

  When enabled, Apusic Application Server rotates log files every day at midnight. This option is disabled by default.

- Multiline Mode

  When enabled, the body of a log message starts on a new line after the message header. This option is enabled by default.

- Console Logging Format

  The format used for logging to the console. Available choices are ULF (UniformLogFormatter) or ODL (Oracle Diagnostics Logging). The default format is ULF.

- Log File Logging Format

  The format used for logging to the server log file. Available choices are ULF or ODL. The default format is ODL.

- Format Exclude Fields

  The fields to exclude in log records. All fields are included by default.

- File Rotation Limit

  The size in bytes that the server log file must reach before it is rotated. When the server log reaches the specified size, Apusic Application Server creates a new, empty file named `server.log` and renames the old file `server.log_`*date*, where *date* is the date and time when the file was rotated. The default value is 2 megabytes. The minimum value for the limit is 500 kilobytes; if you specify a lower value, the file rotates when it reaches 500 KB. To turn off log file rotation, set the value to 0.

- File Rotation Time Limit

  The number of minutes after which the server log must be rotated. The default value is 0, which means that the file is rotated when it reaches the size specified in the File Rotation Limit field. If you specify one or more minutes, the time limit takes precedence over the size limit specified by the File Rotation Limit field.

- Flush Frequency

  The maximum number of messages to be written from the queue to the server log at a time. The default value is 100.

- Maximum History Files

  The maximum number of log files that Apusic Application Server keeps before deleting the oldest file. If you set this value to zero (0), Apusic Application Server does not delete any old log files.

- Log Compress Time

  The time point for compressing the log every day. The format is 00:00.The default value is 00:03.

- Max Compress Files

  To reserve the maximum number of log compressed files to be retained, enter 0 to reserve all compressed log files.The default value is 10.

- Log Compress Path

  Specify the path of the compressed file generated by daily log compression. When modifying this path, ensure that there is authorization.The default value is ${com.apusic.aas.instanceRoot}/logs.

- Log File

  An alternative name or location for the server log file. The default location is *domain-dir*`/logs/server.log`.

- Log Handler

  The absolute class name of a custom log handler. A custom log handler enables you to send logs to a destination other than `server.log` or `syslog`. The custom handler must extend the class `java.util.logging.Handler` (a JSR 047 compliant API). Put the handler class in the Apusic Application Server classpath so that the handler is installed during server startup.

- Push To Kafka

  If this option is selected,the log information is pushed to Kafka.

  The Push to Kafka function contains the following options:

  - Kafka Log Format

    Kafka log format,ODL or ULF.The default value is ODL.

  - Compressed Data Type

    Producer Indicates the compression type used to compress data, including none, gzip, snappy, lz4, and zstd.The default value is none.

  - Services Confirme Number

    The number of service confirmations required for message push completion, with values of all, 0, and 1.The default value is 0.

  - Kafka Cluster Service

    Configure multiple Kafka cluster addresses in the format of host1:port1 and host2:port2.

  - Topic

    Kafka subscribes to topics.The default value is log.

  - Client Id

    The string id used to trace the source of the request.The default value is aas.

  - Block Time

    The blocking duration when the buffer space is insufficient or the metadata is lost, with a value ranging from 0 to 3600000.The default value is 10000ms.

  - Retry Times

    The number of times to retry the push failure, with a value ranging from 0 to 65535.The default value is 1time.

  - Message Send Delay Time

    The maximum idle time for batch data. Batches exceeding this time will also be sent to the broker, with a value ranging from 0 to 3600000.The default value is zero(o)ms.

  - Buffer Memory

    Memory size for cache push data. The value ranges from 0 to 33554432, that is, a maximum of 32 MB.The default value is 33554432 bytes.

  - Batch Size

    Batch capacity to be sent to broker in batches. The value ranges from 0 to 33554432, that is, a maximum of 32 MB.The default value is 16384

  bytes.

4.12.6.2 Module Log Levels

Use the Module Log Levels page to configure the logging levels for individual modules.

For each module, the following information is provided.

- Logger Name

The name of the logger for the module (for example, `javax.enterprise.system.tools.admin`).

- Log Level

  The current logging level for the module.

The Logger Settings table also contains the following options.

- Add Logger

  Button to add a logger module. Clicking this button adds a row to the Logger Settings table. Enter information about the logger module in the new row.

- Delete Logger

  Button to delete one or more selected modules.

- Level

  Drop-down list from which you can select a logging level to be applied to all selected modules.

- Change Level

  Button to change the logging level for one or more selected modules.

The following log levels are available. They are listed in order from highest to lowest.

- EMERGENCY

  The server is in an unusable state. A severe system failure or panic has occurred.

- ALERT

  A particular service is in an unusable state. Automatic recovery is not possible.

- SEVERE

  Events that interfere with normal program execution.

- WARNING

  Warnings, including exceptions.

- INFO

  Messages related to server configuration or server status, excluding errors.

- CONFIG

  Messages related to server configuration.

- FINE

  Minimal verbosity.

- FINER

  Moderate verbosity.

- FINEST

  Maximum verbosity.

- OFF

  No logging messages.

4.12.7 Monitoring Configuration

4.12.7.1 Monitoring Service

Use the Monitoring Service page to configure monitoring options for individual server components or services.

You must deploy an application, component, or service before you configure monitoring for it.

The Monitoring Service page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Monitoring Service

  If the Monitoring Service Enabled checkbox is selected, monitoring is enabled for the Apusic Application Server. This option is enabled by default.

- Monitoring MBeans

  If the Monitoring MBeans Enabled checkbox is selected, all MBeans needed for monitoring are deployed on the Apusic Application Server. This option is enabled by default.If the checkbox is not selected, you cannot view monitoring data in the Administration Console, even if the Monitoring Service checkbox is selected. Instead, you must use the command line to view monitoring data. If you deselect the checkbox, the memory footprint of the Apusic Application Server is reduced.

- Monitoring and acquisition cycle

  Set the collection cycle of monitoring data,60 second by default.

- Monitoring history

  If this option is selected,enabled history monitoring.Disabled by default.

- Exporter Enabled

  If this option is selected,deploy the exporter for AMP.

- Snaps total

  Monitor the total number of data stored in playback.The default value is 8000.

- Monitor playback sampling time

  Configure the time for monitoring playback sampling, which is 30 seconds by default, indicating that monitoring is saved every 30 seconds.

For each available component, the following information is provided.

- Component

  The name of the Apusic Application Server component.

- Monitoring Level

  The level of monitoring set for the component. For each component, available levels are LOW, HIGH, and OFF. The default level is OFF. There is no difference between levels of monitoring offered by the LOW and HIGH options in this release.

The Component Level Settings table also contains the following options.

- Change Level

  Button to change the monitoring level for one or more selected components.

- Level

  Drop-down list from which you can select a monitoring level to be applied to all selected components.

4.12.7.2 Monitoring Alarm

Enable the monitoring alarm function to provide monitoring functions for running threads, memory and CPU,ect.. To use this function, please tick the specified monitoring item and set the threshold and alarm mode.

The Monitoring Alarm page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Alarm

  If this option is selected ,enabled monitoring alarm.Disabled by default.

- Alarm interval (minutes)

  Set the alarm interval,default value is 30 minutes.

- Item

  Select the alarm items, including Thread, CPU, Virtual-memory, Physical-memory, Full-GC, Database Connection Pool, and multiple choices are allowed. The alarm threshold can be set in the drop-down box. When the detection value of the alarm item exceeds the set threshold, an alarm will be triggered.

  - Thread

    Set an alarm based on the HTTP thread pool threshold, which is currently calculated as the number of HTTP threads in use/the total number of all HTTP thread pools; the threshold unit is percentage.

  - CPU

    Set the alarm according to the CPU usage, the current CPU usage value of the machine/total CPU usage value of the machine; the threshold unit is percentage.

  - Virtual-memory

    Set an alert based on the virtual memory usage, AAS current instance heap memory usage/AASXmx; the threshold unit is percentage.

  - Physical-memory

    Set an alert based on the usage of physical memory, with the current physical machine memory usage/total physical machine memory; the threshold unit is percentage.

  - Full-gc

    Set an alarm based on the garbage collection status; it can be based on Single duration: Set the duration of each FullGCC event. If the duration exceeds this time, an alarm event will be generated.Frequency limit: Set the frequency of FullGCC events. For example, 10 times per hour means that ten FullGCC events will generate an alarm within one hour.

  - Database Connection Pool

    Set the alarm according to the connection status of the database connection pool; it can be set:

    1. Validation failed: set an alarm based on the number of "validation failed" times for the database connection pool, with the unit of "time" and a default value of 10 times.
    2. Connection timeout: Set an alarm based on the number of "connection timeouts" in the database connection pool, with the unit of "times". The default value is 10 times.
    3. Idle Timeout: Set the alarm based on the number of "idle timeout" times of the database connection pool, with the unit of "time", and the default value is 10 times.

- Alarm Strategy

  Set the alarm stragegy,send message or snapshots.

  If set Send Message,there are Mail,JMS,SNMP,and Apusic Alarm for choosing.

  If you want to send message by mail,set the following options:

  - mail

    Select to send alert information via email. A security administrator needs to configure the email information in System Settings to configure receiving emails.

  - Receive email

    mailSendToAddress, set the email address to receive alert information. Multiple addresses should be separated by semicolons. The email address must be a valid email address.

  If you want to send message by JMS,set the following options:

- JMS

  After selecting, receive alarm messages through JMS. You need to set up the connection factory and destination resource in the "JMS Resources" section.

- JMS Connection Factory

  To configure the JMS Connection Factory, it needs to be set up in the "Connection Factory" section within the "JMS Resources".

- JMS Destination

  To configure the JMS Destination, it needs to be set up in the "Destination Resource" section within the "JMS Resources".

If you want to send message by SNMP,set the following options:

- SNMP

  After selecting, the alarm message will be sent through SNMP. SNMP v1/v2c/v3 is supported, and only one SNMP version can be set.

- SNMP version

  Select SNMP version, only one can be set.

  - V1:SNMP version is v1:

    SNMP community name: Configure the SNMP community name. The default value is public.

    SNMP address: configure the SNMP address in the format of IP/port, for example, 127.0.0.1/171.

  - V2c:SNMP version is v2c:

    SNMP community name: Configure the SNMP community name. The default value is public.

    SNMP address: configure the SNMP address in the format of IP/port, for example, 127.0.0.1/171.

  - V3:SNMP version is v3:

    SNMP security level: respectively non-security, authentication without encryption, and authentication with encryption. Only one can be set.

    SNMP address: configure the SNMP address in the format of IP/port, for example, 127.0.0.1/171.

    SNMP User Name: Configure the SNMP community name, which needs to be consistent with the SNMP client settings.

    SNMP authentication method: Configure the SNMP authentication method, which can be MD5 or SHA. Only one method can be set, and it needs to be consistent with the SNMP client settings.

    SNMP authentication password: Enter the SNMP authentication password, which needs to be consistent with the SNMP client settings.

    SNMP encryption algorithm: Set the SNMP encryption algorithm, which includes DES, 3DES, AES128, AES192, and AES256. Only one can be set, which needs to be consistent with the SNMP client settings.

    SNMP encryption password: Set the SNMP encryption password, which can only be set once and needs to be consistent with the SNMP client settings.

If you want to send message by Apusic Alarm or other alarm system,set the following options:

- Alarm platform

  Alarm messages can be sent through the alarm platform, requiring configuration of "alarm level" and "alarm platform address".

- Alarm level

  Set the alarm level, indicating that the alarm information generated by AAS belongs to this level in the alarm platform. It can be set to prompt, warning, serious, and disaster

- Alarm platform address

  configure the address of the alarm platform, format: protocol:ip/port, for example: http://127.0.0.1:8888

4.12.7.3 SNMP Listening

By configuring the SNMP agent service, you can obtain the monitoring information of the application server through SNMP.

The SNMP Listening page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- SNMP Proxy Service

  If this option is selected,enabled SNMP Proxy Service.Disabled by default.

- SNMP Address

  The adress of SNMP,the default value is 0.0.0.0.

- SNMP Port

  The port of SNMP,the default value is 161.

- SNMP Transport Type

  Set the transport type of SNMP,TCP or UDP.The default value is UDP.

- SNMP Protocol Version

  The protocol version can be V2C or V3.The default is V2C.

- SNMP Security User Name

  Usm security user name,for example, `public` .

- SNMP Engine Id

  Unique identifier that distinguishes the current SNMP agent device from other devices.For example, `80:00:13:70:01:c0:a8:65` .

- SNMP Validation Protocol

  Authentication algorithm, support MD5,SHA.The default value is MD5.

- SNMP Validation Secret Key

  Ensure that only authorized users can request or receive the used secret key, which length ranges from 8 to 16 bits.The default value is `nmsAuthKey` .

- SNMP Private Protocol

  Encryption algorithms, support DES, 3DES, either AES128.The default value is DES.

- SNMP Private Secret Key

  The secret key used to encrypt and decrypt messages,which length ranges from 8 to 16 bits.The default value is `nmsPrivKey` .

4.12.8 Virtual Servers

A virtual server, sometimes called a virtual host, is an object that allows the same physical server to host multiple Internet domain names. All virtual servers hosted on the same physical server share the Internet Protocol (IP) address of that physical server. A virtual server associates a domain name for a server (such as `www.sun.com`) with the particular server on which the Apusic Application Server is running.

Use the Virtual Servers page to configure virtual servers.

For each virtual server, the following information is provided.

- Name

  The name of the virtual server.

- State

  Either `on` , `off` , or `disabled` .

- Default Web Module

The deployed web module (if any) that is to respond to all requests that cannot be mapped to other web modules deployed to the virtual server.

The Virtual Servers table also contains the following options.

- New

  Button to create a new virtual server.

- Delete

  Button to delete one or more selected virtual servers.

4.12.8.1 New Virtual Server

Use the New Virtual Server page to create a virtual server.

The New Virtual Server page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Id

  Internally visible virtual server identifier. It is not exposed to HTTP clients. The host names that are exposed to HTTP clients must be specified in the Hosts field.

- Hosts

  The host name or names for the machine on which the server is running. Use either actual or virtual host names that are registered with the DNS server for your network (and, on a UNIX system, in your `/etc/hosts` file). The default value is the system property value `${com.sun.aas.hostName}`.

- State

  The desired state for the virtual server. The value may be any of the following:OnThe virtual server is active. This value is the default.OffThe virtual server is inactive. Attempts to access the server will return the error code 404 (resource not available).DisabledThe virtual server is disabled. Attempts to access the server will return the error code 403 (refused to fulfill the request).

- SSO

  Specifies whether single sign-on behavior is inherited from the HTTP Service, enabled, or disabled. By default, it is inherited from the HTTP Service. If this option is enabled, single sign-on is enabled for web applications on this virtual server that are configured for the same realm. If this option is disabled, single sign-on is disabled for this virtual server, and users must authenticate separately to every application on the virtual server.

- Network Listeners

  The network listener or listeners associated with this server, if any.

- Default Web Module

  The deployed web module (if any) that is to respond to all requests that cannot be mapped to other web modules deployed to the virtual server. If a Default Web Module is not specified, the web module that has an empty context root is used. If there is no web module with an empty context root, a system default web module is created and used.

- Log File

  The path name of the file where logging messages from this virtual server will appear. By default, logging messages will appear in *domain-dir*`/logs/server.log`.

- Docroot

  The absolute path to the root document directory for the server. The default value is *domain-dir*`/docroot`.

- Access Logging

  Specifies whether access logging is inherited from the HTTP Service, enabled, or disabled. By default, it is inherited from the HTTP Service.

- Directory

  The absolute directory path to the server access logs. The default value is *domain-dir*`/logs/access`.

- Access address white list

  IP addresses allowed to access. Multiple IP addresses are separated by commas. Wildcards can be used in regular expressions.

- Access address blacklist

  Forbidden IP addresses. Multiple IP addresses are separated by commas. Wildcards can be used in regular expressions.

- Access the domain name white list

  Allowed domain name addresses. Multiple domain name addresses are separated by commas. Wildcards can be used in regular expressions.Enabled 'DNS Lookup' before setting the domain name.

- Access domain name blacklist

  Forbidden domain name addresses. Multiple domain name addresses are separated by commas. Wildcards can be used in regular expressions.Enabled 'DNS Lookup' before setting the domain name.

- Access Control Times

  Setting that can be access between begin and end, format HH:mm:ss(24h).

- Additional Properties

  Additional properties for the virtual server.For information on available properties, see Properties Specific to Virtual Servers.

4.12.8.2 Properties Specific to Virtual Servers

The following additional properties are available for a virtual server.

- `sso-max-inactive-seconds`

  Specifies the number of seconds after which a user's single sign-on record becomes eligible for purging if no client activity is received. Since single sign-on applies across several applications on the same virtual server, access to any of the applications keeps the single sign-on record active.The default value is 300 seconds (5 minutes). Higher values provide longer single sign-on persistence for users at the expense of more memory use on the server.

- `sso-reap-interval-seconds`

  Specifies the number of seconds between purges of expired single sign-on records.The default value is 60.

- `ssoCookieSecure`

  Sets the `Secure` attribute of any `JSESSIONIDSSO` cookies associated with the web applications deployed to this virtual server. This property is applicable only if single sign-on is enabled. Allowed values are as follows: `true` Sets `Secure` to true. `false` Sets `Secure` to false. `dynamic` The `JSESSIONIDSSO` cookie inherits the `Secure` setting of the first session participating in SSO.To set the Secure attribute of a `JSESSIONID` cookie, use the `cookieSecure` `cookie-properties` property in the `sun-web.xml` file.

- `setCacheControl`

  Specifies a comma-separated list of `Cache-Control` response directives.

- `accessLogBufferSize`

  Specifies the size, in bytes, of the buffer where access log calls are stored. If the value is less than 5120, a warning message is issued, and the value is set to 5120. The default value is 32768.

- `accessLogWriterInterval`

  Specifies the number of seconds before the log is written to the disk. The access log is written when the buffer is full or when the interval expires. If the value is 0, the buffer is always written even if it is not full. This means that each time the server is accessed, the log message is stored directly to the file. The default value is 300.

- `allowRemoteAddress`

  Specifies a comma-separated list of regular expression patterns that the remote client's IP address is compared to. If this property is specified, the remote address *must* match for this request to be accepted. If this property is not specified, all requests are accepted *unless* the remote address matches a `denyRemoteAddress` pattern.

- `denyRemoteAddress`

  Specifies a comma-separated list of regular expression patterns that the remote client's IP address is compared to. If this property is specified, the remote address must *not* match for this request to be accepted. If this property is not specified, request acceptance is governed solely by the `allowRemoteAddress` property.

- `allowRemoteHost`

  Specifies a comma-separated list of regular expression patterns that the remote client's hostname (as returned by `[java.net.]Socket.getInetAddress().getHostName()`) is compared to. If this property is specified, the remote hostname *must* match for this request to be accepted. If this property is not specified, all requests are accepted *unless* the remote hostname matches a `denyRemoteHost` pattern.

- `denyRemoteHost`

  Specifies a comma-separated list of regular expression patterns that the remote client's hostname (as returned by `[java.net.]Socket.getInetAddress().getHostName()`) is compared to. If this property is specified, the remote hostname must *not* match for this request to be accepted. If this property is not specified, request acceptance is governed solely by the `allowRemoteHost` property.Setting this property has no effect if the Apusic Application Server domain is accessed through a network listener that has the JK Listener option enabled.

- `authRealm`

  Specifies the name of an authentication realm, which overrides the server instance's default realm for standalone web applications deployed to this virtual server. A realm defined in a standalone web application's `web.xml` file overrides the virtual server's realm.

- `securePagesWithPragma`

  Set this property to `false` to ensure that for all web applications on this virtual server file downloads using SSL work properly in Internet Explorer.Individual web applications may override this setting by using the `sun-web-app` element of the `sun-web.xml` file.The default value is `true`.

- `alternatedocroot_`*n*

  Specifies an alternate document root (docroot), where *n* is a positive integer that allows specification of more than one. Alternate docroots allow web applications to serve requests for certain resources from outside their own docroot, based on whether those requests match one (or more) of the URI patterns of the web application's alternate docroots.If a request matches an alternate docroot's URI pattern, it is mapped to the alternate docroot by appending the request URI (minus the web application's context root) to the alternate docroot's physical location (directory). If a request matches multiple URI patterns, the alternate docroot is determined according to the following precedence order:Exact matchLongest path matchExtension matchFor example, the following properties specify three alternate docroots. The URI pattern of the first alternate docroot uses an exact match, whereas the URI patterns of the second and third alternate docroots use extension and longest path prefix matches, respectively. `<property name="alternatedocroot_1" value="from=/my.jpg dir=/srv/images/jpg"/> <property name="alternatedocroot_2" value="from=*.jpg dir=/srv/images/jpg"/> <property name="alternatedocroot_3" value="from=/jpg/* dir=/src/images"/>` The `value` of each alternate docroot has two components: The first component, `from`, specifies the alternate docroot's URI pattern, and the second component, `dir`, specifies the alternate docroot's physical location (directory). Spaces are allowed in the `dir` component.Individual web applications may override this setting by using the `sun-web-app` element of the `sun-web.xml` file.

- `contextXmlDefault`

  Specifies the location, relative to *domain-dir*, of the `context.xml` file for this virtual server, if one is used. For more information about the `context.xml` file, see The Context Container (`http://tomcat.apache.org/tomcat-5.5-doc/config/context.html`).

- `allowLinking`

  If `true`, resources that are symbolic links will be served for all web applications deployed on this virtual server. Individual web applications may override this setting by using the `sun-web-app` property `allowLinking` in the `sun-web.xml` file: `<sun-web-app> <property name="allowLinking" value="{true|false}"/> </sun-web-app>`.The default value is `false`.**Caution:**Setting this property to true on Windows systems exposes JSP source code.

- `send-error_`*n*

  Specifies custom error page mappings for the virtual server, which are inherited by all web applications deployed on the virtual server. A web application can override these custom error page mappings in its `web.xml` deployment descriptor. The value of each `send-error_`*n* property has three components, which may be specified in any order:The first component, `code`, specifies the three-digit HTTP response status code for which the custom error page should be returned in the response.The second component, `path`, specifies the absolute or relative file system path of the custom error page. A relative file system path is interpreted as relative to the *domain-dir* `/config` directory.The third component, `reason`, is optional and specifies the text of the reason string (such as `Unauthorized` or `Forbidden`) to be returned.For example: `<property name="send-error_1" value="code=401 path=/myhost/401.html`

`reason=MY-401-REASON"/>` This example property definition causes the contents of `/myhost/401.html` to be returned with 401 responses, along with this response line: `HTTP/1.1 401 MY-401-REASON`

- `redirect_`*n*

  Specifies that a request for an old URL is treated as a request for a new URL. These properties are inherited by all web applications deployed on the virtual server. The value of each `redirect_` *n* property has two components, which may be specified in any order:The first component, `from`, specifies the prefix of the requested URI to match.The second component, `url-prefix`, specifies the new URL prefix to return to the client. The `from` prefix is simply replaced by this URL prefix.For example: `<property name="redirect_1" value="from=/dummy url-prefix=http://etude"/>`

- `valve_` *n*

  Specifies a fully qualified class name of a custom valve, where *n* is a positive integer that allows specification of more than one. The valve class must implement the `org.apache.catalina.Valve` interface from Tomcat or previous Apusic Application Server releases, or the `org.Apusic Application.web.valve.Apusic ApplicationValve` interface from the current Apusic Application Server release. For example: `<property name="valve_1" value="org.Apusic Application.extension.Valve"/>` You can set this property for a specific web application.

- `listener_` *n*

  Specifies a fully qualified class name of a custom Catalina listener, where *n* is a positive integer that allows specification of more than one. The listener class must implement the `org.apache.catalina.ContainerListener` or `org.apache.catalina.LifecycleListener` interface. For example: `<property name="listener_1" value="org.Apusic Application.extension.MyLifecycleListener"/>` You can set this property for a specific web application.

- `errorReportValve`

  Specifies a fully qualified class name of a custom valve that produces default error pages for applications on this virtual server. Specify an empty string to disable the default error page mechanism for this virtual server.

4.12.8.3 Edit Virtual Server

Use the Edit Virtual Server page to modify the settings for a virtual server.

The Edit Virtual Server page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Id

  Internally visible virtual server identifier. It is not exposed to HTTP clients. The host names that are exposed to HTTP clients must be specified in the Hosts field.

- Hosts

  The host name or names for the machine on which the server is running. Use either actual or virtual host names that are registered with the DNS server for your network (and, on a UNIX system, in your `/etc/hosts` file). The default value is the system property value `${com.sun.aas.hostName}`.

- State

  The desired state for the server. The value may be any of the following:OnThe virtual server is active. This value is the default.OffThe virtual server is inactive. Attempts to access the server will return the error code 404 (resource not available).DisabledThe virtual server is disabled. Attempts to access the server will return the error code 403 (refused to fulfill the request).

- SSO

  Specifies whether single sign-on behavior is inherited from the HTTP Service, enabled, or disabled. By default, it is inherited from the HTTP Service. If this option is enabled, single sign-on is enabled for web applications on this virtual server that are configured for the same realm. If this option is disabled, single sign-on is disabled for this virtual server, and users must authenticate separately to every application on the virtual server.

- Network Listeners

  The network listener or listeners associated with this server.

- Default Web Module

The deployed web module (if any) that is to respond to all requests that cannot be mapped to other web modules deployed to the virtual server. If a Default Web Module is not specified, the web module that has an empty context root is used. If there is no web module with an empty context root, a system default web module is created and used.

- Log File

  The path name of the file where logging messages from this virtual server will appear. By default, logging messages will appear in *domain-dir* `/logs/server.log`..

- Docroot

  The absolute path to the root document directory for the server. The default value is *domain-dir* `/docroot`.

- Access Logging

  Specifies whether access logging is inherited from the HTTP Service, enabled, or disabled. By default, it is inherited from the HTTP Service.

- Directory

  The absolute directory path to the server access logs. The default is *domain-dir* `/logs/access`.

- Access address white list

  IP addresses allowed to access. Multiple IP addresses are separated by commas. Wildcards can be used in regular expressions.

- Access address blacklist

  Forbidden IP addresses. Multiple IP addresses are separated by commas. Wildcards can be used in regular expressions.

- Access the domain name white list

  Allowed domain name addresses. Multiple domain name addresses are separated by commas. Wildcards can be used in regular expressions.Enabled 'DNS Lookup' before setting the domain name.

- Access domain name blacklist

  Forbidden domain name addresses. Multiple domain name addresses are separated by commas. Wildcards can be used in regular expressions.Enabled 'DNS Lookup' before setting the domain name.

- Access Control Times

  Setting that can be access between begin and end, format HH:mm:ss(24h).

- Additional Properties

  Additional properties for the virtual server.

4.12.9 Web Container

4.12.9.1 General Properties

Use the General Properties page to add or modify properties for the web container.

The web container hosts web applications. It provides the environment in which servlets and JavaServer Pages (JSP) files run.

The following additional general property is available for the web container.

- `dispatcher-max-depth`

  Prevents recursive include or forward statements from creating an infinite loop by setting a maximum nested dispatch level. If this level is exceeded, the following message is written to the server log: `Exceeded maximum depth for nested request dispatches` The default value is 20.

For each property, the General Properties page contains the following information.

- Name

  The name of the property.

- Value

  The value of the property.

- Description

  A description of the property.

4.12.9.2 Session Properties

Use the Session Properties page to add or modify properties for the web container session.

The Session Properties page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Session Timeout

  Specifies the maximum number of seconds that an inactive session remains valid. If the value is set to 0 or less, sessions never expire. The default value is 1800.

- Additional Properties

  Additional properties for web container sessions.

The following properties are available for a web container session.

- `enableCookies`

  If set to `true`, uses cookies for session tracking. The default value is `true`.

- `enableURLRewriting`

  Enables URL rewriting. This provides session tracking by means of URL rewriting when the browser does not accept cookies. You must also use an `encodeURL` or `encodeRedirectURL` call in the servlet or JavaServer Pages (JSP) page. The default value is `true`.

- `idLengthBytes`

  Specifies the number of bytes in this web module's session ID. The default value is 128.

4.12.9.3 Manager Properties

The session manager provides the means to configure how sessions are created and destroyed, where the session state is stored, and the maximum number of sessions that are available.

Use the Manager Properties page to add or modify session manager properties for the web container.

The Manager Properties page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Reap Interval

  The number of seconds before inactive session data is deleted from the store. The default value is 60.Set this value lower than the frequency at which session data changes. For example, this value should be as low as possible (1 second) for a hit counter servlet on a frequently accessed web site; otherwise, you could lose the last few hits each time you restart the server.

- Max Sessions

  The maximum number of sessions that can be in cache. A value of -1 indicates no limit to the number of sessions. The default value is -1.After the limit is reached, an attempt to create a new session causes an `IllegalStateException` to be thrown.

- Session ID Length

Set the length of Session ID，not less than 10.The default value is 14.

- Session File Name

  The absolute or relative path to the directory in which the session state is preserved between application restarts, if preserving the state is possible. If this path is not set, the session state is not preserved. By default, this path is not set.A relative path is relative to the temporary directory for this web application.

- Session ID Generator Class Name

  The custom class for generating unique session IDs. Only one session ID generator class per server instance is permitted, and all instances in a cluster must use the same session ID generator to prevent session key collision. The class must be in the Application Server classpath. By default, a custom class is not set.

- Additional Properties

  Additional properties for the session manager. Apusic Application Server does not define any additional properties for the session manager.

4.12.9.4 Store Properties

Use the Store Properties page to add or modify session persistence (storage) properties for the web container.

The Store Properties page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Directory

  The absolute or relative pathname of the directory into which individual session files are written. The default value is *domain-dir* `/generated/jsp/j2ee-apps/` *appname* `/` *appname* `_war`. A relative path is relative to the temporary work directory for this web application.

- Additional Properties

  Additional properties for the store. The Apusic Application Server does not define any additional properties for the store.

4.12.10 EJB Container

4.12.10.1 EJB Settings

Enterprise Java Beans (EJB)

Use the Enterprise Java Beans (EJB) page to configure settings for the EJB container.

In addition to general settings for enterprise beans, you can configure pool settings and cache settings.

- Pool settings apply only to stateless session beans. By default, the container maintains a pool of enterprise beans in order to respond to client requests without the performance hit that results from creating the beans.

  If you experience performance problems in an application that uses deployed enterprise beans, you can help improve the applications's performance by creating a pool or by increasing the number of beans maintained by an existing pool.

- Cache settings apply only to stateful session beans. The container maintains a cache of enterprise bean data for the most used enterprise beans. This allows the container to respond more quickly to requests from other application modules for data from the enterprise beans.

  Cached enterprise beans are in one of three states: active, idle, and passivated. An active enterprise bean is currently being accessed by clients. An idle enterprise bean's data is currently in the cache, but no clients are accessing the bean. A passivated bean's data is temporarily stored and is read back into the cache if a client requests the bean.

The Enterprise Java Beans (EJB) page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Session Store Location

The directory where passivated stateful session beans and persisted HTTP sessions are stored on the file system. The default value is `${com.sun.aas.instanceRoot}/session-store`.Passivated beans are stateful session beans that have had their state written to a file on the file system. Passivated beans typically have been idle for a certain period of time, and are not currently being accessed by clients. Similarly, persisted HTTP sessions are individual web sessions that have had their state written to a file on the file system.

- Commit Option

  Specifies how the container caches passivated bean instances between transactions. If you select Option B (the default), the container caches a ready instance between transactions. If you select Option C, the container does not cache a ready instance between transactions.

- Initial and Minimum Pool Size

  The minimum number of beans to be maintained in the pool. The default value is 0.

- Maximum Pool Size

  The maximum number of beans that the container can maintain in the pool at one time. The default value is 32.

- Pool Resize Quantity

  The number of beans that are removed from the pool if they are idle for more than the time specified in the Pool Idle Timeout field. The value must be at least 0 and less than the Maximum Pool Size value. The default value is 8.

- Pool Idle Timeout

  The time, in seconds, that a bean in the pool can remain idle before it is removed from the pool. After this amount of time, the bean is destroyed. A value of 0 means that a bean can remain idle indefinitely. The default value is 600.

- Max Cache Size

  The maximum number of beans to be held in the cache. Increase the maximum number of beans to cache to eliminate the overhead of bean creation and destruction. However, if the cache is increased, the server consumes more memory and resources. Be sure your operating environment is sufficient for your cache settings. A value of 0 indicates an unbounded cache. The default value is 512.

- Cache Resize Quantity

  The number of beans to passivate when the cache is full. When the maximum number of cached beans is reached, the container removes a number of passivated beans from the backup store. This value also specifies the number of beans to be created if a request arrives when the pool has no available beans. The value must be greater than 1 and less than the Max Cache Size value. The default value is 32.

- Removal Timeout

  The number of seconds after which a passivated bean is removed from the session store. A value of 0 specifies that the container does not remove inactive beans automatically. If the Removal Timeout value is less than or equal to the Cache Idle Timeout value, beans are removed immediately without being passivated. The default value is 5400.

- Removal Selection Policy

  The policy the container uses to remove stateful session beans from the cache. The choices are as follows:Not recently used (NRU)Removes a bean that hasn't been used recently. This value is the default.First in, first out (FIFO)Removes the oldest bean in the cache.Least recently used (LRU)Removes the least recently accessed bean.

- Cache Idle Timeout

  The maximum number of seconds that a bean can remain idle in the cache. After this amount of time, the container can passivate this bean. A value of 0 specifies that beans never become candidates for passivation. The default value is 600.

- Additional Properties

  Additional properties for the EJB container.

The following properties are available for configuring the EJB container.

- `thread-core-pool-size`

  Specifies the number of core threads in the EJB container's common thread pool. The default value is `16`.

- `thread-max-pool-size`

Specifies the maximum number of threads in the EJB container's common thread pool. The default value is `32`.

- `thread-queue-capacity`

  Specifies the size of the thread pool queue, which stores new requests if more than `thread-core-pool-size` threads are running. The default value is the `Integer.MAX_VALUE`.

- `thread-keep-alive-seconds`

  Specifies the time, in seconds, past which threads in excess of `thread-core-pool-size` are terminated. The default value is `60`.

- `allow-core-thread-timeout`

  If set to `true`, all threads, even core threads, are subject to termination after `thread-keep-alive-seconds`. The default value is `false`.

- `prestart-all-core-threads`

  If set to `true`, all core threads in the EJB container's common thread pool are started, causing them to idly wait for work. If set to `false`, threads are not started until new requests arrive. The default value is `false`.

- `disable-nonportable-jndi-names`

  If set to `true`, Apusic Application Server specific JNDI names for an EJB module are disabled. The default is `false`.Because the EJB 3.1 specification defines portable EJB JNDI names, there is less need for Apusic Application Server specific JNDI names. By default, Apusic Application Server specific default JNDI names are applied automatically for backward compatibility.

4.12.10.2 MDB Settings

MDB Default Pool Settings

Use the MDB Default Pool Settings page to configure the MDB pool.

The MDB Default Pool Settings page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Initial and Minimum Pool Size

  The minimum number of message beans to be maintained in the pool. The default value is 0.

- Maximum Pool Size

  The maximum number of beans that the container can maintain in the pool at one time. The default value is 32.

- Pool Resize Quantity

  The number of beans that are removed from the pool if they are idle for more than the time specified in the Pool Idle Timeout field. The value must be at least 0 and less than the Maximum Pool Size value. The default value is 8.

- Pool Idle Timeout

  The time, in seconds, that a bean in the pool can remain idle before it is removed from the pool. After this amount of time, the bean is destroyed. A value of 0 means that a bean can remain idle indefinitely. The default value is 600.

- Additional Properties

  Additional properties for the MDB pool. The Apusic Application Server does not define any additional properties for the MDB pool.

4.12.10.3 EJB Timer Service

Use the EJB Timer Service page to configure the EJB timer service.

The timer service is a persistent and transactional notification service that is provided by the enterprise bean container and is used to schedule notifications or events used by enterprise beans. All enterprise beans except stateful session beans can receive notifications from the timer service. Persistent timers set by the service are not

destroyed when the server is shut down or restarted.

The EJB Timer Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Minimum Delivery Interval

  The minimum number of milliseconds allowed before the next timer expiration for a particular timer can occur. Setting this interval too low can cause server overload. The default value is 1000.

- Maximum Redeliveries

  The maximum number of attempts the timer service makes to deliver a timer expiration due for exception or rollback. The default value is 1.

- Redelivery Interval

  The interval, in milliseconds, between redelivery attempts. The default value is 5000.

- Timer Datasource

  The JNDI name of the JDBC resource that will be used as the timer datasource. If specified, this value overrides the default value of `jdbc/__TimerPool` for the timer service system application.

4.12.11 Java Message Service

Use the Java Message Service page to configure settings to be used by the default JMS provider, Apusic Application Server Message Queue.

The Java Message Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Ping

  Button to verify that the JMS service is running on the default JMS host. If the JMS service is up and running, a message reports that the ping succeeded.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Type

  The type of access for the JMS service. Available choices are:EMBEDDEDAccess the JMS service on the local host. The JMS service is started in the same JVM machine as Apusic Application Server. Lazy initialization starts the default embedded broker on the first access of JMS services rather than at Apusic Application Server startup. This is the default type of JMS Service.LOCALAccess the JMS service on the local host. The JMS provider is started along with the Apusic Application Server.REMOTEAccess the JMS service on another system. If you choose REMOTE, the JMS service is not started by Apusic Application Server the next time the server starts. Instead, the JMS service is started and managed by Message Queue, so you must start the Message Queue broker separately. If you choose this value and are using a remote host, follow the instructions in To Edit a JMS Host to specify the name of the remote host.

- Startup Timeout

  The number of seconds Apusic Application Server waits for the JMS service to start before aborting the startup. On a slow or overloaded system, increase the value from the default. The default value is 60.

- Start Arguments

  Arguments to customize the JMS service startup. Use any arguments available through the *as-install-parent* `/mq/bin/imqbrokerd` command.

- Reconnect

If the Reconnect Enabled checkbox is selected, the JMS service attempts to reconnect to a message server (or the list of addresses in the AddressList) when a connection is lost. This option is enabled by default.

- Reconnect Interval

  The number of seconds between reconnect attempts. This interval applies for attempts on each address in the AddressList and for successive addresses in the list. If it is too short, this time interval does not give a broker time to recover. If it is too long, the reconnect might represent an unacceptable delay. The default value is 5 seconds.

- Reconnect Attempts

  The number of attempts to connect (or reconnect) for each address in the AddressList before the client runtime tries the next address in the list. A value of -1 indicates that the number of reconnect attempts is unlimited (the client runtime attempts to connect to the first address until it succeeds). The default value is 3.

- Default JMS Host

  The name of the default JMS host. The default value is `default_JMS_host`.

- Master Broker

  The name of the Apusic Application Server clustered instance whose associated Message Queue broker is to be used as the master broker in the Message Queue broker cluster. Leave this field blank to enable Apusic Application Server to designate a master broker automatically.If you change this value, you must restart the Apusic Application Server cluster that uses the configuration.

- Address List Behavior

  The order of connection attempts. Available choices are: `random` Select an address from the AddressList randomly. If there are many clients attempting a connection using the same connection factory, specify `random` to prevent them from all being connected to the same address. This option is the default. `priority` The reconnect always tries to connect to the first server address in the AddressList and uses another one only if the first broker is not available.

- Address List Iterations

  The number of times the JMS service iterates through the AddressList in an effort to establish (or reestablish) a connection. A value of -1 indicates that the number of attempts is unlimited. The default value is 3. The maximum value is 2147483647.

- MQ Scheme and MQ Service

  The Message Queue address scheme name and the Message Queue connection service name if a non-default scheme or service is to be used.

- Additional Properties

  Additional properties for the JMS Service.For information on available properties, see Properties Specific to the JMS Service.

4.12.11.1 Properties Specific to the JMS Service

The following Oracle Message Queue broker configuration properties are available for the JMS Service.

- `instance-name`

  Specifies the full Message Queue broker instance name. The default is `imqbroker`.

- `instance-name-suffix`

  Specifies a suffix to add to the full Message Queue broker instance name. The suffix is separated from the instance name by an underscore character (`_`). For example, if the instance name is `imqbroker`, appending the suffix `xyz` changes the instance name to `imqbroker_xyz`.

- `append-version`

  If `true`, appends the major and minor version numbers, preceded by underscore characters (`_`), to the full Message Queue broker instance name. For example, if the instance name is `imqbroker`, appending the version numbers changes the instance name to `imqbroker_8_0`. The default is `false`.

- `user-name`

  The user name for creating the JMS connection. Needed only if the default username/password of `guest` / `guest` is not available in the broker. The default value is `guest`.

- `password`

The password for creating the JMS connection. Needed only if the default username/password of `guest` / `guest` is not available in the broker. The default value is `guest`.

4.12.11.2 JMS Hosts

Use the JMS Hosts page to configure JMS service hosts.

For each host, the following information is provided.

- Name

  The name of the JMS host.

- Host

  The name or Internet Protocol (IP) address of the system where the JMS host is running.

- Port

  The port number of the JMS service on the host.

The Hosts table also contains the following options.

- New

  Button to create a new JMS host.

- Delete

  Button to delete one or more selected JMS hosts.

4.12.11.2.1 NEW JMS HOST

Use the New JMS Host page to create a new JMS host.

The New JMS Host page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the host (for example, `NewJmsHost`).

- Host

  The name or Internet Protocol (IP) address of the system where the JMS host will run (`localhost` or the name of the local or remote system).

- Port

  The port number of the JMS service. Change this field only if the JMS service to be used is running on a non-default port. The default port is 7676.

- Admin Username and Admin Password

  The Message Queue broker user name and password. These are different from the Apusic Application Server user name and password. Edit these fields only if the Message Queue broker values have been changed using the *as-install-parent* `/mq/bin/imqusermgr` command. The default values are `admin` and `admin`.

4.12.11.2.2 EDIT JMS HOST

Use the Edit JMS Host page to modify the settings for a JMS host.

The Edit JMS Host page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the host. The Name field is a read-only field. You can only specify the Name field when you create a new JMS host.

- Host

  The name or Internet Protocol (IP) address of the system where the JMS host will run (`localhost` or the name of the local or remote system).

- Port

  The port number of the JMS service. Change this field only if the JMS service to be used is running on a non-default port. The default port is 7676.

- Admin Username and Admin Password

  The Message Queue broker user name and password. These are different from the Apusic Application Server user name and password. Edit these fields only if the Message Queue broker values have been changed using the *as-install-parent* `/mq/bin/imqusermgr` command. The default values are `admin` and `admin`.

4.12.12 ORB

Use the ORB page to configure the Object Request Broker.

The ORB provides the required infrastructure to identify and locate objects, handle connection management, deliver data, and request communication. Remote clients of enterprise beans (EJB modules) communicate with the Apusic Application Server using the ORB.

The ORB page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Thread Pool ID

  The thread pool used by the ORB. The ORB uses thread pools to respond to requests from remote clients of enterprise beans and other clients that communicate by using RMI-IIOP. The default thread pool for the ORB is `thread-pool-1`, which uses defaults appropriate for RMI/IIOP requests.

- Max Message Fragment Size

  The maximum fragment size for IIOP messages. Messages larger than this size are fragmented. The default value is 1024 bytes.

- Total Connections

  The maximum number of incoming connections for all IIOP listeners. The default value is 1024.

- IIOP Client Authentication

  If the Enabled checkbox is selected, the server rejects unauthenticated requests and inserts an authentication-required bit in Interoperable Object References (IORs) sent to clients. IIOP client authentication can be enabled or disabled. This option is disabled by default.

- Additional Properties

  Additional properties for the ORB. The Apusic Application Serverdoes not define any additional properties for the ORB.

4.12.12.1 IIOP Listeners

Use the IIOP Listeners page to configure IIOP listeners.

A number of IIOP listeners can be configured for an ORB, each accepting connections on different network address and port combinations. Listeners can be configured differently for secure communications and client authentication.

By default, three IIOP listeners are configured: `orb-listener-1`, `SSL`, and `SSL_MUTUALAUTH`.

For each IIOP listener, the IIOP Listeners page contains the following information.

- Name

The name of the IIOP listener.

- Enabled

  True if the IIOP listener is enabled, or false if the listener is not enabled.

- Network Address

  The Internet Protocol (IP) address or DNS resolvable host name for the IIOP listener.

- Listener Port

  The port number for the IIOP listener.

4.12.12.1.1 NEW IIOP LISTENER

Use the New IIOP Listener page to create a new IIOP listener.

The New IIOP Listener page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique name that identifies the IIOP listener.

- Network Address

  The network address of the IIOP listener. The network address can be an IP address or a DNS resolvable host name.

- Listener Port

  The port number upon which the IIOP listener is to listen.

- Listener

  The status of the listener. If the Enabled checkbox is selected, ORB IIOP inbound connections to the Apusic Application Server are enabled. This option is enabled by default.

- Security

  If the Enabled checkbox is selected, SSL is enabled for the IIOP listener. This option is disabled by default.

- Additional Properties

  Additional properties for the IIOP listener. The Apusic Application Serverdoes not define any additional properties for IIOP listeners.

4.12.12.1.2 EDIT IIOP LISTENER

Use the Edit IIOP Listener page to modify the settings for an IIOP listener.

The Edit IIOP Listener page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique name that identifies the IIOP listener. The Name field is a read-only field. You can only specify the name when you create an IIOP listener.

- Network Address

  The network address of the IIOP listener. The network address can be an IP address or a DNS resolvable host name.

- Listener Port

The port number upon which the IIOP listener is to listen.

- Listener

  The status of the listener. If the Enabled checkbox is selected, ORB IIOP inbound connections to the Apusic Application Server are enabled. This option is enabled by default.

- Security

  If the Enabled checkbox is selected, SSL is enabled for the IIOP listener. This option is disabled by default.

- Additional Properties

  Additional properties for the IIOP listener. The Apusic Application Serverdoes not define any additional properties for IIOP listeners.

4.12.12.1.3 SSL

Use the SSL page to modify the SSL settings for an IIOP listener.

The SSL page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- SSL3

  If this checkbox is selected, the SSL3 protocol is enabled for the IIOP listener. This option is enabled by default.

- TLS

  If this checkbox is selected, the TLS protocol is enabled for the IIOP listener. This option is enabled by default.

- Client Authentication

  If this checkbox is selected, clients must identify themselves to the server on every request. This option is disabled by default.

- Certificate Nickname

  The nickname of the server certificate in the certificate database or the PKCS#11 token. In the certificate, the name format is *tokenname* `:` *nickname*. Including the *tokenname* `:` part of the name in this attribute is optional.

- Key Store

  The name of the keystore file (for example, `keystore.jks`).

- Trust Algorithm

  The name of the trust management algorithm (for example, PKIX) to use for certification path validation.

- Max Certificate Length

  The maximum number of non-self-issued intermediate certificates that can exist in a certification path. This field is used only if the Trust Algorithm field is set to PKIX. A value of 0 implies that the path can only contain a single certificate. A value of -1 implies that the path length is unconstrained (there is no maximum). Setting a value less than -1 causes an exception to be thrown.

- Trust Store

  The name of the truststore file (for example, `cacerts.jks`).

- Cipher Suites

  An area where you can add or remove cipher suites. If you do not add any cipher suites, all cipher suites will be used.

4.12.13 System Properties

The System Properties page displays a list of the Java system properties that are defined in the selected named configuration. Java system properties are passed to the Java application launcher through the `-D` option of the Java application launcher whenApusic Application Server is started

These properties define the default values of port settings for all Apusic Application Server instances that reference the configuration.

For information about predefined port settings in a configuration, see Predefined System Properties.

For each property, the following information is displayed:

- Instance Variable Name

  The name of the system property.

- Default Value

  The value that is set for the property in the named configuration that the instance references. This field is read only.

- Instance Values

  A link to the Instance Values page for the property.

The Additional Properties table also contains the following options.

- Add Property

  Button to add a property. Clicking this button adds a row to the Additional Properties table.

- Delete Properties

  Button to delete one or more selected properties. Any property that is deleted reverts to its default value or, if no default value is set, is undefined.

The System Properties page also contains the following options:

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Dynamic Reconfiguration

  If this option is enabled, changes to the configuration are applied to the instances that reference the configuration without the need to restart the instances.

- Clusters and/or Instances Using This Configuration

  A list of the clusters and instances that reference the configuration with an indication of whether each cluster or instance is running:If the cluster or instance is running, a check mark (✓) is displayed.If the cluster or instance is stopped, an exclamation point (`!`) is displayed.

4.12.13.1 Predefined System Properties

Apusic Application Server defines the following configuration system properties:

- `ASADMIN_LISTENER_PORT`

  This property specifies the port number of the HTTP port or HTTPS port through which the DAS connects to the instance to manage the instance. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `HTTP_LISTENER_PORT`

  This property specifies the port number of the port that is used to listen for HTTP requests. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `HTTP_SSL_LISTENER_PORT`

  This property specifies the port number of the port that is used to listen for HTTPS requests. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `IIOP_LISTENER_PORT`

  This property specifies the port number of the port that is used for IIOP connections. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `IIOP_SSL_LISTENER_PORT`

  This property specifies the port number of the port that is used for secure IIOP connections. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `IIOP_SSL_MUTUALAUTH_PORT`

  This property specifies the port number of the port that is used for secure IIOP connections with client authentication. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `JAVA_DEBUGGER_PORT`

  This property specifies the port number of the port that is used for connections to the Java Platform Debugger Architecture (JPDA) (`http://www.oracle.com/technetwork/java/javase/tech/jpda-141715.html/`) debugger. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `JMS_PROVIDER_PORT`

  This property specifies the port number for the Java Message Service provider. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `JMX_SYSTEM_CONNECTOR_PORT`

  This property specifies the port number on which the JMX connector listens. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

- `OSGI_SHELL_TELNET_PORT`

  This property specifies the port number of the port that is used for connections to the Apache Felix Remote Shell (`http://felix.apache.org/site/apache-felix-remote-shell.html`). This shell uses the Felix shell service to interact with the OSGi module management subsystem. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

4.12.13.2 Instance Values

Use the Instance Values page to edit the values of a configuration system property for the Apusic Application Server instances that reference the selected named configuration.

The Instance Values page displays a list of instances that reference the selected configuration, in which the property is set.

- Instance Name

  The name that was assigned to the instance when the instance was created. Clicking this name opens the General Information page for the instance.

- Cluster Name

  The name of the cluster of which the instance is a member, if any. This field is read only.

- Default Value

  The value that is set for the property in the named configuration that the instance references. This field is read only.

- Override Value

  The value of the property that is set for the selected instance. This value overrides the default value.

4.12.14 Security

Use the Security page to set security properties for the selected Apusic Application Server instance or cluster.

The Security page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Security Manager

  If this option is selected, the security manager for the domain is enabled. This option is disabled by default.When this option is enabled, a JVM option, `-Djava.security.manager` , will be added to the JVM setting of the Apusic Application Server. You must restart the server to enable this change.Ensure that you have granted correct permissions for all applications. You can turn off the security manager to enhance performance.

- Default Realm

The active (default) realm that the server uses for authentication. Applications use this realm unless their deployment descriptor specifies a different realm. All configured realms appear in the list. The default value is `file`.

- JACC

  The class name of a configured JACC provider. The default value is `default`.

- Hostname Checker

  Whether to validate the host name.

  - Nocheck: Do not enable the hostname verification function.
  - Check: Enable the hostname verification function.
  - Custom: Set the custom hostname validator. At this time, you need to fill in the "Hostname Checker Class".

- Additional Properties

  Additional security properties for the server.Valid properties are dependent on the type of realm selected in the Default Realm field and are typically specified when you edit a realm.

4.12.14.1 Realms

Use the Realms page to create or set security realm properties for the Apusic Application Server instance or cluster.

The Realms page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the realm.Click the name go to edit page.

- Class Name

  The class name of the realm.

The Realms table also contains the following options.

- New

  Button to create a new realm.

- Delete

  Button to delete one or more selected realms.

4.12.14.1.1 NEW REALM

Use the New Realm page to create a new realm.

The New Realm page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique name that identifies the realm.

- Class Name

  The class name of the realm.Valid properties are dependent on the type of realm selected in the Default Realm field.

4.12.14.1.2 EDIT REALM

Use the Edit Realm page to modify the settings for a realm.

The Edit Realm page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Realm Name

  A unique name that identifies the realm.This field is read only.

- Class Name

  The class name of the realm.This field is read only.

- Properties specific to this Class

  Valid properties are dependent on the type of realm selected in the Default Realm field and are typically specified when you edit a realm.

4.12.14.1.3 PROPERTIES SPECIFIC TO THE `FILEREALM` CLASS

The following properties are required for a `file` realm.

- JAAS Context

  The JAAS (Java Authentication and Authorization Service) context (the identifier for the login module to use for this realm). The only valid value is `fileRealm`.

- Key File

  Full path and name of the file where the server will store all user, group, and password information for this realm. The default value is *domain-dir* `/config/keyfile` for the `file` realm and *domain-dir* `/config/admin-keyfile` for the `admin-realm` realm.

The key file for the `file` realm is initially empty, so users must be added on the File Users page before the `file` realm is used.

The key file for the `admin-realm` realm initially contains the administrator user name, the administrator password in an encrypted format, and the group to which this user belongs (`asadmin` by default).

Note:

Users in the group `asadmin` in the `admin-realm` are authorized to use the Administration Console and `asadmin` tools. Add only users to this group that have server administrative privileges.

The following optional property is available for a `file` realm.

- Assign Groups

  A comma-separated list of group names. All clients who present valid certificates are assigned to these groups, for example, `employee,manager`, where these are the names of user groups.

4.12.14.1.4 FILE USERS

Use the File Users page to manage `file` realm users. Users and groups in the `file` realm are listed in the key file, whose location is specified by the `file` property.

A user in the `file` realm can belong to a *Java group*, a category of users classified by common traits. For example, customers of an e-commerce application might belong to the `CUSTOMER` group, but the big spenders would belong to the `PREFERRED` group. Categorizing users into groups makes it easier to control the access of large numbers of users.

Initially after installation of the Apusic Application Server, the only user is the administrator entered during installation. By default, this user belongs to the group `asadmin`, in the realm `admin-realm`, which gives rights to modify the Apusic Application Server. Any users assigned to this group will have administrator privileges, that is, they will have access to the `asadmin` tool and the Administration Console.

The File Users page displays a list of users for the selected realm. For each user, the following information is provided:

- User ID

  The user name.

- Group List

  The groups to which the user belongs.

The File Users table also contains the following options.

- New

  Button to create a new user.

- Delete

  Button to delete one or more selected users.

4.12.14.1.5 PROPERTIES SPECIFIC TO THE `CERTIFICATEREALM` CLASS

The `certificate` realm supports SSL authentication. This realm sets up the user identity in the Apusic Application Server's security context, and populates it with user data obtained from cryptographically verified client certificates in the truststore and keystore files. Add users to these files using `keytool` or `certutil`.

With the `certificate` realm, Java containers handle authorization processing based on each user's Distinguished Name (DN) from his or her certificate. The DN is the name of the entity whose public key the certificate identifies. This name uses the X.500 standard, so it is intended to be unique across the Internet. For more information on key stores and trust stores, refer to the keytool documentation (`http://java.sun.com/j2se/1.4.2/docs/tooldocs/solaris/keytool.html`).

The following optional property is available for the `certificate` realm.

- Assign Groups

  A comma-separated list of group names. All clients who present valid certificates are assigned to these groups, for example, `employee,manager`, where these are the names of user groups.

4.12.14.1.6 PROPERTIES SPECIFIC TO THE `JDBCREALM` CLASS

To protect your web pages or web applications, you can set the security so that only registered users can access them. This is known as the authentication facility. This type of realm involves storing the credentials of your users inside a database. The Apusic Application Server uses the database information and the enabled JDBC realm option inside the configuration file.

The following properties are required for a JDBC realm.

- JAAS Context

  The JAAS (Java Authentication and Authorization Service) context (the identifier for the login module to use for this realm). The only valid value is `jdbcRealm`.

- JNDI

  The JNDI name for this realm. The default value is `jdbc/security`.

- User Table

  The table that contains a list of authorized users for this realm. The default value is `usertable`.

- User Name Column

  The name of the column that contains the list of users inside the user table. The default value is `userid`.

- Password Column

  The name of the column that contains the respective user's password in the user table. The default value is `password`.

- Group Table

  The name of the group table in the database. The default value is `grouptable`.

- Group Name Column

  The name of the group name column in the database's group table. The default value is `groupid`.

The following optional properties are available for a JDBC realm.

- Assign Groups

  A comma-separated list of group names. All clients who present valid certificates are assigned to these groups, for example, `employee,manager`, where these are the names of user groups.

- Database User

Allows you to specify the database user name in the realm instead of the `jdbc-connection-pool`. This prevents other applications from looking up the database, getting a connection, and browsing the user table. By default, the `jdbc-connection-pool` configuration is used.

- Database Password

  Allows you to specify the database password in the realm instead of the `jdbc-connection-pool`. This prevents other applications from looking up the database, getting a connection, and browsing the user table. By default, the `jdbc-connection-pool` configuration is used.

- Digest Algorithm

  (Optional) Specifies the digest algorithm. The default is `SHA-256`. You can use any algorithm supported in the JDK, or none.**Note:**In versions of Apusic Application Server prior to 4.0, the default algorithm was `MD5`. If you have applications that depend on the `MD5` algorithm, you can override the default `SHA-25` algorithm by using the `asadmin set` subcommand: `asadmin> **set server.security-service.property.default-digest-algorithm=MD5**` You can use the `asadmin get` subcommand to determine what algorithm is currently being used: `asadmin> **get server.security-service.property.default-digest-algorithm**` Also note that, to maintain backward compatibility, if an upgrade is performed from Apusic Application Server v2.*x* or v3.0.*x* to Apusic Application Server 4.0, the default algorithm is automatically set to `MD5` in cases where the digest algorithm had not been explicitly set in the older Apusic Application Server version.

- Password Encryption Algorithm

  The algorithm for encrypting passwords stored in the database.**Note:**It is a security risk not to specify a password encryption algorithm.

- Encoding

  The encoding. Allowed values are `Hex` and `Base64`. If `digest-algorithm` is specified, the default is `Hex`. If `digest-algorithm` is not specified, by default no encoding is specified.

- Charset

  The charset for the digest algorithm.

4.12.14.1.7 PROPERTIES SPECIFIC TO THE `LDAPREALM` CLASS

The following properties are required for an LDAP realm.

- JAAS Context

  The JAAS (Java Authentication and Authorization Service) context (the identifier for the login module to use for this realm). The only valid value is `solarisRealm`.

- Directory

  The LDAP URL for your server.

- Base DN

  The LDAP base distinguished name (DN) for the location of user data. This base DN can be at any level above the user data, since a tree scope search is performed. The smaller the search tree, the better the performance.

The following optional property is available for an LDAP realm.

- Assign Groups

  A comma-separated list of group names. All clients who present valid certificates are assigned to these groups, for example, `employee,manager`, where these are the names of user groups.

The following additional properties are available for an LDAP realm.

- `search-filter`

  The search filter to use to find the user. The default is `uid=%s` (`%s` expands to the subject name).

- `group-base-dn`

  The base DN for the location of group data. By default, it is same as the Base DN value, but it can be tuned, if necessary.

- `group-search-filter`

The search filter to find group memberships for the user. The default is `uniquemember=%d` (`%d` expands to the user element DN).

- `group-target`

  The LDAP attribute name that contains group name entries. The default is `CN` .

- `search-bind-dn`

  An optional DN used to authenticate to the directory for performing the `search-filter` lookup. Only required for directories that do not allow anonymous search.

- `search-bind-password`

  The LDAP password for the DN given in `search-bind-dn` .

4.12.14.1.8 PROPERTIES SPECIFIC TO THE `SOLARISREALM` CLASS

The following property is required for a Solaris realm.

- JAAS Context

  The JAAS (Java Authentication and Authorization Service) context (the identifier for the login module to use for this realm). The only valid value is `solarisRealm` .

The following optional property is available for a Solaris realm.

- Assign Groups

  A comma-separated list of group names. All clients who present valid certificates are assigned to these groups, for example, `employee,manager` , where these are the names of user groups.

4.12.14.1.9 PROPERTIES SPECIFIC TO THE `PAMREALM` CLASS

The following property is required for a PAM realm.

- JAAS Context

  The JAAS (Java Authentication and Authorization Service) context (the identifier for the login module to use for this realm). The only valid value is `pamRealm` .

The following optional property is available for a PAM realm.

- Assign Groups

  A comma-separated list of group names. All clients who present valid certificates are assigned to these groups, for example, `employee,manager` , where these are the names of user groups.

4.12.14.1.10 PROPERTIES SPECIFIC TO THE `CUSTOMREALM` CLASS

If you want to customize a realm,specify the Class Name .

4.12.14.2 Audit Modules

Use audit modules to develop an audit trail of all authentication and authorization decisions.

The Audit Modules page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the audit module.Click the name go to the edit page.

- Class Name

  The class name of the audit module.

The Modules table also contains the following options.

- New

Button to create a new audit module.

- Delete

  Button to delete one or more selected audit modules.

4.12.14.2.1 NEW AUDIT MODULE

Use the New Audit Modules page to create a new audit module.

The New Audit Modules page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique name that identifies the audit module.

- Class Name

  The class name of the audit module.

- Additional Properties

  Additional properties for the audit module.

4.12.14.2.2 EDIT AUDIT MODULE

Use the Edit Audit Modules page to modify a audit module.

The Edit Audit Modules page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  A unique name that identifies the audit module.This field is read only.

- Class Name

  The class name of the audit module.

- Additional Properties

  Additional properties for the audit module.

4.12.14.3 JACC Providers

Use JACC Providers to manage Java Authorization Contract for Containers (JACC) providers.

The JACC Providers page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name of the JACC Provider.Click the name go to the edit page.

- Policy Provider

  The policy provider of the JACC provider.

The JACC Provider table also contains the following options.

- New

  Button to create a new JACC provider.

- Delete

  Button to delete one or more selected JACC paoviders.

4.12.14.3.1 NEW JACC PROVIDER

Use the New JACC Providers page to create a new JACC Provider.

The New JACC Providers page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  JACC provider name; must contain only alphanumeric, underscore, dash, or dot characters.

- Policy Configuration

  Class that implements policy configuration factory.

- Policy Provider

  Class that implements policy factory.

- Additional Properties

  Additional properties for the JACC Provider.

4.12.14.3.2 EDIT JACC PROVIDER

Use the Edit JACC Provider page to modify a audit module.

The Edit JACC Provider page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Name

  The name that identifies the JACC Provider.This field is read only.

- Policy Configuration

  Class that implements policy configuration factory.

- Policy Provider

  Class that implements policy factory.

- Additional Properties

  Additional properties for the JACC Provider.

4.12.14.4 Message Security

Use Message Security Configurations to configure message security providers; enables server to perform end-to-end authentication of web service invocations and responses at the message layer.

The Message Security Configurations page contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Authentication Layer

  The authentication layer of the message security provider.Click the name go to the edit page.

- Default Provider

The default provider of the message security provider.

- Default Client Provider

  The default client provider of the message security provider.

4.12.14.4.1 EDIT HTTPSERVLET AUTHENTICATION LAYER

Use the Edit Message Security Configuration page to modify the HttpServlet authentication layer .

The Edit Message Security Configuration contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Authentication Layer

  The authentication layer of the message security provider.Here is HttpServlet.

- Default Provider

  The default provider of the message security provider.

- Default Client Provider

  The default client provider of the message security provider.

4.12.14.4.2 EDIT SOAP AUTHENTICATION LAYER

Use the Edit Message Security Configuration page to modify the SOAP authentication layer .

The Edit Message Security Configuration contains the following options.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Authentication Layer

  The authentication layer of the message security provider.Here is SOAP.

- Default Provider

  The default provider of the message security provider.

- Default Client Provider

  The default client provider of the message security provider.

4.12.15 Transaction Service

Use the Transaction Service page to configure transaction recovery, timeouts, and logging.

The Transaction Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- On Restart

  If the Enabled checkbox is selected, the Apusic Application Server attempts to recover incomplete transactions when the server is restarted. This option is disabled by default.

- Transaction Timeout

  Number of seconds the server waits before rolling back a transaction that has not completed. The default value is 0, meaning that the server waits indefinitely for a transaction to complete.

- Retry Timeout

  Number of seconds the Apusic Application Server tries to connect to an unreachable server. The default value is 600 (10 minutes).

- Transaction Log Location

  The directory where server logs are kept. Transaction logs are kept in the `tx` subdirectory of the directory specified by this field. The default value is the directory specified by the Log Root field of the Domain Attributes page, which is located under the Advanced tab of the Apusic Application Server node.

- Heuristic Decision

  Whether transactions that involve unreachable servers are to be committed or rolled back. The default value is Rollback. Committing indeterminate transactions can compromise the data integrity of your application.

- Keypoint Interval

  The number of transactions between keypoint operations, which compress the transaction log file. The default value is 65,536.

- Additional Properties

  Additional properties for the Transaction Service.For a description of available properties, see Properties Specific to the Transaction Service.

4.12.15.1 Properties Specific to the Transaction Service

The following properties are available for configuring the Transaction Service.

- `oracle-xa-recovery-workaround`

  If set to `true`, the Oracle XA Resource workaround is used in transaction recovery. The default value is `true`.

- `disable-distributed-transaction-logging`

  If set to `true`, disables transaction logging, which might improve performance. If the On Restart Enabled checkbox is selected, this property is ignored. The default value is `false`.

- `xaresource-txn-timeout`

  Changes the `XAResource` timeout. In some cases, the `XAResource` default timeout can cause transactions to be aborted, so it is desirable to change it. The default value is specific to the `XAResource` used.

- `pending-txn-cleanup-interval`

  Specifies the interval, in seconds, at which an asynchronous thread checks for pending transactions and completes them. If this property is not specified, there is no default. If this property is present but has no value, the default value is 60.

- `use-last-agent-optimization`

  If set to `true`, enables last agent optimization, which improves the throughput of transactions. If one non-XA resource is used with XA resources in the same transaction, the non-XA resource is the last agent. The default value is `true`.

- `delegated-recovery`

  If set to `true`, cluster-wide delegated recovery is enabled. The default value is false.

- `wait-time-before-recovery-insec`

  Specifies the wait time, in seconds, after which an instance starts the recovery for a dead instance.

- `db-logging-resource`

  Specifies the JNDI name of the JDBC resource for the database to which transactions are logged. FThere is no default value.

- `xa-servername`

  Specifies the host name that the transaction service uses to identify transactions being managed by the installed Apusic Application Server. This can sometimes be useful for recovering transactions from the log file that was created on a different host running the Apusic Application Server.

4.12.16 Connector Service

Use the Connector Service page to modify general settings for the Connector Service, which governs resource adapters.

The Connector Service page contains the following options.

- Load Defaults

  Button to restore settings that have default values to their default values. Settings that do not have default values are not changed.

- Configuration Name

  The name of the configuration to which the settings on this page apply. This field is read only.

- Shutdown Timeout

  The maximum number of seconds allowed during Apusic Application Server shutdown for the `ResourceAdapter.stop` method of a connector module's instance to complete. Resource adapters that take longer to shut down are ignored, and Apusic Application Server shutdown continues. The default value is 30 seconds.

- Connector Classloading Policy

  The policy to be used for loading classes. Available choices are: `derived` Indicates that the resource adapters are provided according to an application's references to any of the resource adapter's resources. This policy is the default. `global` Indicates that all resource adapters will be visible to all applications.This policy does not apply to the preinstalled JMS system resource adapter, `jmsra` , which is always visible to all applications.

4.13 Security Configuration for Security Role

The Security module can be managed by logging in to the control platform through the security administrator.

4.13.1 System configuration

Use the System configuration page to config the server's password policy,session,log,backup,message validation related and so on.

The System configuration page contains the following options.

Password Strategy

The password policy configuration item mainly configures the length and complexity of the user's password, with the following specific attributes:

- Password Length

  Specifies the password length requirement for all users, including administrators, which must be greater than or equal to the set value; the default is 8 characters.

- Password expire

  Specifies the number of days that a password is valid from the date of modification. If the password exceeds the validity period, it needs to be modified again. The default is 30 days.

- ValidHour

  Specify the valid hours of the account, starting from the first login. If the account exceeds the set valid hours, it will be disabled; the default is 720 hours. Users with the security role are not affected.

- Max login attempts

  The number of times a user is allowed to retry their password when logging in. If the number of times exceeds this limit, the user will be locked out; the default is 5 times.

- Password complexity

  Set the complexity of password,the default value is Common. Commom, must be a combination of uppercase and lowercase English letters, numbers, and special characters; Complex, must be a combination of uppercase and lowercase English letters, numbers, and special characters;.

- Password restore day

  The period of time after a user fails to log in that they are locked out. If the lockout time exceeds this value, they will automatically be unlocked. The default is 15 minutes.

- Weak password setting

Create a file weak-password under ${DOMAIN_HOME}/config/. The password set in the file is the weak password. When users set passwords, verification is not allowed. Multiple passwords are separated by English, such as setting abcd1234, Abcd123, Apusic321 as weak passwords. When users set passwords of abcd1234, Abcd123, Apusic321, verification is not allowed.

Session configuration

The session configuration item mainly configures the total number of sessions for user login and the session validity time. The specific configuration attributes are as follows:

- Session expire

  Specifies the maximum number of seconds an inactive session remains active.The default value is 1800 seconds.

- Max sessions

  The maximum number of users allowed to log in simultaneously (the total number of sessions for all users in the system, not for individual accounts; when a single account is repeatedly logged in, it is necessary to log out of other sessions before logging in again).The default value is -1 that represents unlimited.

Security log configuration

The audit and operation configuration items mainly configure the quantity and retention time of audit logs and operation logs, respectively. The specific configuration attributes are as follows:

- Max Autdit log Num

  The number limit in the audit log storage database. If it exceeds this limit, the oldest records will be overwritten.The default value is 10000.

- Max audit log preserved day

  The retention time of audit log records, in days..The default value is 180 days.

- Max Operate log Num

  The number limit in the operation log storage database, if exceeded, the oldest records will be overwritten.The default value is 10000.

- Max operate log preserved day

  The operation log record retention time, in days.The default value is 180 days.

- Scheduled backup cycle

  The regular backup cycle set for audit and operation logs, with the default save directory being ${DOMAIN_HOME}/audit/timing, and the unit being month/time.

- Scheduled backup retention days

  The retention time for the audit and operation logs of the timed backup, in days. -1 indicates permanent retention.

Business configuration backup

Business configuration backup mainly configures the default instance business backup directory and scheduled backup cycle. The specific configuration attributes are as follows:

- Share backup directory

  The directory where the business configuration information is backed up is backed up to the shared backup directory, which supports local and remote backups. Remote backups require file sharing to be set up first. The default backup location is ${com.apusic.aas.instanceRoot}/backup/config.

- Scheduled backup cycle

  The cycle of business configuration information backup, in months/time. The default is 1 month/time.

- Scheduled backup Enabled

  Whether to enable the timing backup business configuration. After enabling, the configuration information will be backed up according to the "timing backup cycle". The default is not to enable.

Email validation configuration

The main function of mail authentication configuration is to perform secondary verification for login users on the console.

- Mailbox switch

  The value is True/False; set whether to enable the email verification function.

- Send mail server host

  Set the email server for sending emails, typically in the format smtp.***.com.

- Send mail server port

  Set the port of the send mail server, and use the ssl port.

- Mail sender email

  Set the email address for sending emails.

- Send mail password

  Enter the password of the sending email address. If client authorization password is enabled, enter the authorization code.

- Mail recipient mailbox

  Set the email address to receive the verification code.

After enabling the email verification setting, the user clicks on login and enters the verification code input page.

After entering the verification code, click "OK" to enter the console operation page.

If you want to return to the login page, click "Return".

When you need to obtain the verification code again, click "Didn't receive the verification code, resend it".

4.13.2 User configuration

Use the User configuration page to create and manage users. Unable to modify or delete reserved users including admin,secure, and audit.

The information of users is stored in `${DOMAIN_HOME}/mydomain/config/admin-keyfile`.

The User configuration page contains the following options.

- User name

  The user name of the application server, which needs to be unique.

- Status

  Displays the status of the user.

- roles

  Displays the user's role.

- Password Expire

  Displays the expiration time of the password. The default is 30 days after initializing the user. Modify the "Password Expiration Time" in System Management to synchronize the time changes.

- Allow start(yyyy-MM-dd HH:mm:ss)

  The start time for users to access the console, in the format HH: MM.

- Allow End(yyyy-MM-dd HH:mm:ss)

  The end time for users to access the console, in the format HH: MM.

- Allowed IPs

  The IP that the user is allowed to access, which is the IP of the browser.

- Secret Level

  The confidentiality level of the user.

- PRIVILEGE

  Whether it is an initialized user. When it is TRUE, the user cannot be deleted directly.

The User table also contains the following options.

- Add

  Button to create a new user.

- Delete

  Button to delete one or more selected users.

4.13.2.1 New User

Use the New System User page to create a new user.

The New System User page contains the following options.

- User Name

  The user name of the application server. The name can contain up to 255 characters and can only contain alphanumeric, underscore, dash, or dot characters. It needs to be unique.

- Role Name

  Select the role of the user. Different roles have different permissions. There are security, sysadmin, auditor, and publish_role, which represent users with restricted access to create, edit, and other operations.

- New Password

  Set the user's password. In the "System configuration" section, when the "Password Complex" setting is set to "Commom", the password must contain at least two combinations of letters, numbers, and special symbols; when it is set to "Complex", the password must contain letters, numbers, and special symbols.

- Confirm New Password

  You need to enter the same password as the "New Password".

4.13.2.2 Edit User

Use the Edit System User page to modify a user's information.

The Edit System User page contains the following options.

- User Name

  The name of the user, this field is read only.

- Status

  Set the user status. NORMAL indicates normal; LOCKED indicates locked, which is locked for 15 minutes by default; DISABLED indicates disabled.

- Role Name

  Set the role of the user.

- New Password

  Set the user's password. In the "System Management" section, when the "Password Complexity" setting is set to "Normal", the password must contain at least two combinations of letters, numbers, and special symbols; when it is set to "Complex", the password must contain letters, numbers, and special symbols. The password cannot be repeated within 5 changes.

- Confirm New Password

  Please enter the same password as the "New Password".

- Allow start(yyyy-MM-dd HH:mm:ss)

  The start time for users to access the console, in the format HH: MM.

- Allow End(yyyy-MM-dd HH:mm:ss)

  The end time for users to access the console, in the format HH: MM.

- Allowed IPs

  The IP that the user is allowed to access, which is the IP of the browser.

- Secret Level

  The confidentiality level of the user.The user can only access applications or resources of the same or lower level. Currently, there are three levels of confidentiality: secret, confidential, and top secret. The permission level is: secret < confidential < top secret

Note:

After editing user information, the user's session status will be refreshed synchronously, and the session that is currently logged in will be invalidated and logged out.

4.13.2.3 Reset Password

The user can refer to the following methods to reset the password.

I. Resetting the password for current users

The current user logs in to the control platform and enters the [Administrator Password] page through the homepage to modify the password.

II. Resetting User Password Through the Security Administrator

The security administrator logs in to the control platform, goes to [User Management], clicks on the user's name, and enters the user's edit page to set a password for that user.

III. Resetting User Password by the System Administrator

The system administrator logs in to the control platform, goes to Configuration-> "server-config" -> Security -> Realm -> "admin-realm" -> "Manage Users", clicks on the user's name, and enters the edit page where they can modify the user's password.

IV. Resetting the user password when forgetting the password

Method 1: If a non-security management user, such as admin, forgets their password, they can log in to the control platform through the security administrator to reset their password.

Method 2: If an administrator including the security administrator and system administrator forgets the password, but other users know the password, you can copy the password of other users in mydomain/config/admin-keyfile under the installation path and replace the password of the user who needs to change the password. It takes effect after the system is restarted.

Method 3: If an administrator including the security administrator forgets the password, but remembers the password of the system administrator, use the system administrator such as admin to log in to the management, and enter the page of [Configuration Management] - [server-config] - [Security Services] - [Security Domain] - [admin-realm] - [Manage Users]. Click on the user name to edit the user, and you can modify the user password.

Method 4: Modify the password of the corresponding user in the mydomain/config/admin-keyfile file, replace the following password with `{SSHA256}iU4Ef2uGWYh3V+BQjpW5f8BTbgGWzKp7pfrNU020Nu2l9YLEwCWOpA==` , reset it to an empty password, and restart the system. You need to reset the password again.

4.13.2.4 Modify the initial user name

Generally, the system comes with three default user names for initialization, namely admin/secure/audit. It is not recommended to modify them by default. If you need to modify them, there are two solutions.

The first type: AAS has not been initialized. At this time, you can enter `${DOMAIN_HOME}/mydomain/config/admin-keyfile` to modify the corresponding username to the one you need to set, save it, and start AAS.

The second method: AAS has been initialized. You need to delete the database `${DOMAIN_HOME}/mydomain/database/userDataBase` ;modifythecorrespondingusernamein `${DOMAIN_HOME}/mydomain/config/admin-keyfile` to the user name that needs to be set, and save it. Start AAS, and you need to re-enter the user password at this time.

4.13.3 Role configuration

Use the Role configuration page to create and manage roles. The AAS comes with roles including sysadmin, security, and auditor, which cannot be edited or deleted. The developer role publish_role is also included, which allows editing or deletion.

The Role configuration page contains the following options.

- Role Name

  The name of the role.

- The user to which the role belongs

  The role to which the current role belongs. The "roles" setting in User configuration will be updated in real time.

- Role of resources

  The resources owned by the role.

- PRIVILEGE

  Whether it is an initialized user. When it is TRUE, the user cannot be deleted directly.

The role table also contains the following options.

- Add

  Button to create a new role.

- Delete

  Button to delete one or more selected roles.

4.13.3.1 New Role

Use the New Role page to create a new role.

The New Role page contains the following options.

- Role Name

  The name of the role. The name can contain up to 255 characters and can only contain alphanumeric, underscore, dash, or dot characters. It needs to be unique.

- Developer Role

  The developer role does not have some permissions to modify configurations, resources, and configuration modules.The default value is false.

- Role of resources

  Set resources for the current role. The Shift key can be used for continuous multi-selection, and the Ctrl key can be used for multiple selections. It is recommended to select "index" if needed.

4.13.3.2 Edit Role

Use the Edit Role page to modify a role's information.After modification, users need to log in to the control platform again to take effect.

The Edit Role page contains the following options.

- Role Name

  The name of the role. This field is read only.

- Developer Role

  The developer role does not have some permissions to modify configurations, resources, and configuration modules.The default value is false.

- Role of resources

  Set resources for the current role. The Shift key can be used for continuous multi-selection, and the Ctrl key can be used for multiple selections. It is recommended to select "index" if needed.

4.14 Log for Auditor Role

The Log module can be managed by logging in to the control platform through the auditor administrator.

The system will generate corresponding logs for important operations, sensitive data operations, or some illegal operations, and auditors can audit these logs.

4.14.1 Operate Logs

Log in to the control system as an audit administrator and switch to the operate Logs page.

The following operations can be performed in the Operate Log page:

- You can view all the operation logs of the system and filter them based on the criteria you add.
- Click the Backup button to backup the operation log. The backup log file is stored in the `${APUSIC_HOME}/domain_name/audit` directory.

4.14.2 Audit Logs

Log in to the control system as an auditor and switch to the Audit Logs page.

The following operations can be performed in the Audit Logs page:

- You can view all the audit logs of the system and filter and query them based on the added information.
- Click the Backup button to backup the audit log. The backup log file is stored in the `${APUSIC_HOME}/domain_name/audit` directory.

4.15 Transaction Management

The running transaction list can view the running time, status and other information of the transaction. It has the function of manually rolling back the running transaction.

Before you can view monitoring data, you must configure monitoring. Set the level of Transaction Service to HIGH or LOW.

When a transaction is active, the page will display transaction information and wait for it to run. After the transaction ends, it will no longer be displayed in the list.

Select the transaction list and click "Rollback" to manually cancel the transaction.

4.16 Lifecycle Modules

Use the Lifecycle Modules page to configure lifecycle modules.

A lifecycle module performs tasks when it is triggered by one or more events in the server's lifecycle. Possible trigger server events are: initialization, startup, ready to service requests, and shutdown. Lifecycle modules are not part of the Java EE specification, but are an enhancement to the Apusic Application Server.

For each lifecycle module, the following information is provided.

- Name

  The name of the lifecycle module.

- Enabled or Status

  A check mark if the application is enabled, or an X if the application is disabled, if only the default server instance, `server`, exists.If clusters or other standalone server instances exist, select Enabled on All Targets to view the targets on which the application is deployed.

- Load Order

  A value specifying the relative order in which the lifecycle module should be loaded.

The Lifecycle Modules table also contains the following options.

- New

  Button to create a new lifecycle module.

- Delete

  Button to delete one or more selected lifecycle modules.

- Enable

  Button to enable one or more selected lifecycle modules. Present if only the default server instance, `server`, exists.

- Disable

Button to disable one or more selected lifecycle modules. Present if only the default server instance, `server`, exists.

4.16.1 New Lifecycle Module

Use the New Lifecycle Module page to define a new lifecycle module.

The New Lifecycle Module page contains the following options.

- Name

  The name of the lifecycle module.

- Class Name

  The fully qualified name of the lifecycle module's class file. The class must implement the `com.sun.appserv.server.LifecycleListener` interface.

- Classpath

  The classpath for the lifecycle module. The classpath specifies where the lifecycle module is located. The default value is *domain-dir* `/applications`. If the module is already in this directory (that is, in the server classpath), this field can be left blank.

- Load Order

  The order in which this lifecycle module is to be loaded at startup. Modules with smaller integer load order values are loaded sooner. The value can range from 101 to the operating system's `MAXINT`. Values from 1 to 100 are reserved.

- Description

  A description of the lifecycle module.

- Status

  If this option is selected, the lifecycle module is enabled. This option is enabled by default.

- On Load Failure

  If this option is selected, the server will be shut down if the lifecycle module fails to load. This option is disabled by default.

- Targets

  Clusters and standalone instances for the lifecycle module. Move desired targets to the Selected Targets column using the Add and Add All buttons. Move any unneeded targets to the Available Targets column using the Remove and Remove All buttons.This option is displayed only if clusters or standalone instances have been created in the domain.

4.16.2 Edit Lifecycle Module

Use the Edit Lifecycle Module page to modify settings for a lifecycle module.

The Edit Lifecycle Module page contains the following options.

- Name

  The name of the lifecycle module. The name is a read-only field. You can only specify a name when you create a new lifecycle module.

- Class Name

  The fully qualified name of the lifecycle module's class file. The class must implement the `com.sun.appserv.server.LifecycleListener` interface.

- Classpath

  The classpath for the lifecycle module. The classpath specifies where the lifecycle module is located. The default value is *domain-dir* `/applications`. If the module is already in this directory (that is, in the server classpath), this field can be left blank.

- Load Order

  The order in which this lifecycle module is to be loaded at startup. Modules with smaller integer load order values are loaded sooner. The value can range from 101 to the operating system's `MAXINT`. Values from 1 to 100 are reserved.

- Description

A description of the lifecycle module.

- Status

  If this option is selected, the lifecycle module is enabled. This option is enabled by default.

- On Load Failure

  If this option is selected, the server will be shut down if the lifecycle module fails to load. This option is disabled by default.

4.16.3 Lifecycle Module Targets

Use the Lifecycle Module Targets page to view target clusters and standalone server instances on which the lifecycle module can be enabled.

The Lifecycle Module Targets page contains the following information.

- Target Name

  The name of the cluster or standalone server instance.

- Enabled

  Displays `true` if the lifecycle module is enabled on the target, or `false` if the lifecycle module is disabled.

The Targets table also contains the following options.

- Manage Targets

  Button to manage lifecycle module targets.

- More Actions

  Drop-down list of the following actions.EnableAction to enable the lifecycle module on the selected target.DisableAction to disable the lifecycle module on the selected target.

4.17 Patch

The Apusic application server supports patch upgrades through the control platform. The patch function is divided into two modules: patch management, which manages the uploading and removal of patches through the management console, and lists the corresponding effects of patches; and instance upgrade, which displays the patch allocation and effectiveness of all instances in the instance upgrade interface, and the background will transfer or remove patches to instances on remote nodes through the sftp protocol according to actual needs.

4.17.1 Processing Mechanism

The class loading mechanism of the AAS V10 application server is relatively complex. There are four common application server jar packages or bytecode paths that require patch replacement:

- ApusicAS/aas/modules/
- *ApusicAS/aas/lib/install/applications/*
- ApusicAS/aas/domains/${domain-name}/lib/*
- ApusicAS/aas/lib/install/applications/__admingui/WEB-INF/lib/

The files are processed by different class loaders. When the application server loads the patch, it identifies the class loader corresponding to each jar through the directory structure within the patch, such as 1./modules 2./lib/install/applications/ 3./lib 4./console, and prioritizes the processing of the path at the corresponding stage.

4.17.2 Patch naming

In order to better manage patches and clarify the relationship between patches, the naming of patches is standardized.

The specific naming rules are as follows:

VersionNumber	Description	Example
First version number	Indicates software major version updates and replacements	V**10**.0.1.P1
Second version number	Indicating a planned major feature upgrade of the software	V10.**0**.1.P1
Third version number	Indicates a planned small feature upgrade for the software	V10.0.**1**.P1

Fourth version number	Version number released	V10.0.1.P1
Fith version number	Minor version number released	V10.0.1.P1.001
Add T after the version number	Temporary version number	V10.0.1.P1T
Add T after the minor version number	Temporary Minor version number	V10.0.1.P1.001T

4.17.3 Filtering and sorting of patches

In the future, there may be many temporary and formal patches that do not require users to care about and organize. The patch tool will automatically determine which patches should take effect.

Both major and minor version patches are cumulative patches. In other words, at most one major and minor version patch can be in effect at the same time. If the major version patch is more recent than the minor version patch, the minor version patch will not be in effect.

example:

Patch type:

1. Cumulative patch package.
2. Temporary cumulative patch package.
3. Cumulative minor version patch package.
4. Temporary minor version patch package.

Note: The cumulative patch includes all patches, and the minor version patch includes the transition patch package.

For example:

A V10.0.1.P1 contains a.jar b.jar c.jar

B temporary V10.0.1.P1.001T includes e.jar

C V10.0.1.P1.001 includes e.jar

D temporary V10.0.1.P1.002T includes e.jar f.jar

E V10.0.1.P1.002 includes e.jar and f.jar

F (temporary) V10.0.1.P2T includes a.jar b.jar c.jar e.jar f.jar

G V10.0.1.P2 includes a.jar, b.jar, c.jar, e.jar, and f.jar

Loading mechanism:

When A and B exist simultaneously, load A and B

When A, B, and C exist simultaneously, load A and C

When A, B, C, and D exist simultaneously, load A and D

When A, B, C, D, and E are present simultaneously, load A and E

When A, B, C, D, E, and F are present simultaneously, load F

When A, B, C, D, E, F, and G are present simultaneously, load G

4.17.4 Patch Manage

You can access the Patch -> Patch Manage through the control platform to upload patch files. The patch files need to be `zip` packages, and they need to be packaged according to the processing mechanism and patch naming.

After uploading, the uploaded files are saved under patches/.

If the patch file has already been specified, the patch file will be extracted to patch/. Deleting the patch file on the Patch Manage page will not affect the files under patch/, but there will no longer be a patch zip package in patches/

4.17.5 Instance Upgrade

Instance upgrades can be performed through the control platform. The Instance Upgrade page displays all instances in the current AAS, and allows for the specified patching and rollback of patches for the instances.

- Name

  The name of the instance.

- Clusters

  The instance type,singleton instance or cluster.

- Nodes

  The node where the instance is located.

- Location

  The location where the patch file is stored.

- Patches Specified

  The patches that have been specified for the instance.

- Specified time

  The time specified by the patch.

- Effective Time

  The time when the patch becomes effective.

- Rollback Patches

  Enter the rollback patch page and select the patch file to roll back.

- Specify Patches

  Enter the specified patch page and select the patch file to specify.

- Choose Cluster

  You can select the instance of the cluster to view instance upgrade infomations.

Note:

Please backup the application server before upgrading the patch.

4.17.5.1 Specify Patches

Select the instance that requires a patch to be specified,Click Specify Patches Button enter the page for specifying the patch.

Use the Specify Patches page to upgrade instances.

Select the patch, and click Enable button. After selecting, you will be prompted to restart the instance for the patch to take effect.

4.17.5.2 Rollback Patches

Select the instance that needs to be rolled back, click Rollback Patches button enter the rollback patch page.

Use the Rollback Patches page to upgrade instances.

Select the patch, and click Rollback Patch button. After selecting, you will be prompted to restart the instance for the rollback to take effect.

5 License Authorization

The Apusic Application Server requires a corresponding license to function properly. Typically, Kingdee Apusic Cloud Computing Co., Ltd will provide the corresponding license based on the product version purchased by the user.

The license location is `${APUSIC_HOME}/license.xml`.

If you encounter issues such as an expired or invalid license during usage, it is recommended that you first contact the designated Apusic service personnel to re-apply for the corresponding license. When re-applying for the license, please ensure to provide the product's authorization code (auth code) to the Apusic service personnel for processing.

If there is an authorization issue with the node, you can view the log, which will print out the corresponding auth code. Copy the auth code and provide it to the Apusic docking personnel to re-apply for the corresponding license.

Starting from December 2023, the method for obtaining feature codes will be modified.

You can run the command startserv in the `ApusicAS/aas/bin` directory to obtain it, similar to the format:

`startserv -ac [ethname or ip]`

Where [ethname or ip] indicates that you can enter the name of an IP address or MAC address, such as the commonly used IP address 172.24.1.116 or eth0 as a parameter to obtain software signature codes (applicable to Windows and Linux):

```
startserv -ac 172.24.1.116
```

The output format is as follows, with Auth Code= the right side representing the software feature code:

`Auth Code=SZTY1640356587`

You can also enter the network card name to obtain the signature code (applicable to Linux, you can obtain the network card name through ifconfig)

```
startserv -ac eth0
```

The output format is as follows, Auth Code=the right side is the software feature code.

6 GM/T Certificates Configuration

The Apusic Application Server supports the configuration of national cryptography certificates GM/T Certificates, which usually refer to national commercial cryptography algorithm certificates. This means that the application server can integrate and use encryption certificates based on Chinese national commercial cryptography standards such as SM2, SM3, SM4, and other algorithms to enhance its security performance and data protection capabilities.

The national cryptography certificate is issued by a certificate authority (CA) that meets the requirements of the National Cryptography Administration of China. They use China's self-developed cryptographic algorithms for signing, encryption, decryption, and other operations. This is particularly important for applications deployed within China, as the Chinese government encourages and supports the use of national cryptography algorithms to protect national security and information security.

6.1 Configuring GM/T Certificates

Before configuring the GM/T certificate, the certificate file needs to be prepared in advance.Such as keystore.p12,truststore.jks,CA.cert.pem.

The Apusic Application Server enables the support function of GM/T certificate.

1.Modifying the `domain.xml` :

```
<jvm-options>-Dcom.apusic.security.ssl.EnableGMTLS=true</jvm-options>
```

2.Copy the file keystore.p12 to `${APUSIC_DOMAIN}/config/`

3.Modifying the attributies of `domain.xml` :

```
<protocol name="http-listener-2" security-enabled="true">
<http max-connections="250" default-virtual-server="server">
<file-cache></file-cache>
</http>
<ssl classname="com.sun.enterprise.security.ssl.ApusicSSLImpl"
sm11-enabled="true"
key-store-type="PKCS12"
key-store="${com.apusic.aas.instanceRoot}/config/keystore.p12"
key-store-password="fdjaio%^&2123f"
client-auth="want"></ssl>
</protocol>
```

Attribute description:

- security-enabled

  When the value is true, enable HTTPS mode.

- sm11-enabled

  When the value is true, enable TLS for national encryption.

- key-store

  Set the path of keystore.

- key-store-type

  Set the type of keystore, generally is JKS, PKCS12

- key-store-password

  Set the password of keystore.

- trust-store

  Set the path of truststore.

- trust-store-type

Set the type of the truststore, generally is JKS.

- trust-store-password

  Set the password of truststore.

- client-auth-enabled

  When the value is true, client authentication is enabled.

4. Start the Apusic Application Server.

Note:

You can configure the GM/T certificate through the control platform,Configurations-[server-config]-Network Config-Protocols-[jttp-listener-2]-SSL,enabled sm11Enabled button and configure the informations of certificate.

7 JMX APIs

JMX (Java Management Extensions) connectivity is a standard way to manage and monitor Java applications. JMX allows developers and administrators to monitor and manage Java applications by defining a set of standard interfaces and patterns. These interfaces and patterns allow different tools to interact with Java applications in a unified way.

The Apusic Application Server supports JMX APIs,you can connect as follow:

7.1 Setting by Console

Set the address and port through the console. [Configurations] -> [server-config] -> [Admin Service] -> [JMX Connector].

- Address

  Specifies the IP address or host name; name must contain only alphanumeric, underscore, dash, or dot characters.For example, `172.20.140.11`.

- Port

  Specifies the port of the naming service (RMIRegistry) where the stub for JMX Connector Service is bound.The default value is 6886.

After modify,restart the Apusic Application Server.

7.2 Setting by `domain.xml`

Set the address and port through the `domain.xml`. Modifying as the following:

```
<jmx-connector address="172.20.140.41" port="6886" name="system" auth-realm-name="admin-realm" security-enabled="false">
```

After modify,start the Apusic Application Server.

7.3 Connect through tools

Open a connection tool, such as Jconsole, and enter the console's IP and port in the remote process, such as `172.20.140.11:6886`; enter the username for the console, such as `admin`; enter the password for the console's corresponding username; click "Connect" to complete the connection.

8 RESTful APIs

RESTful(Representational State Transfer) APIs are primarily used by other systems to retrieve monitoring information from application servers.

The Apusic Application Server supports RESTful APIs.

Usually, accessing https://IP:6848/monitoring/domain allows entry to the monitoring page.Add `.json` after the contextroot, and the data in json format will appear.

Note:

Before you can view monitoring data, you must configure monitoring. See To Configure the Monitoring Servicefor details.

9 Configuring Parameters through Environment Variables

The Apusic Application Server supports setting thread pools and JVM parameters through environment variables. Parameters can be set directly in the profile file or referenced from a file.

Note:

1. It is necessary to set the environment variable export AAS_CONFIG for the settings to take effect. A value of true indicates enablement, false indicates disablement, and a file path indicates referencing that file.
2. If environment variables are set, their values will be applied with priority.
3. When setting JVM parameters, multiple parameters should be separated by spaces. To remove a JVM parameter, prefix it with [-]. JVM parameters will be referenced directly but will not modify the displayed values of JVM options in the domain.xml configuration file.
4. When referencing external files, the file format can be .conf, .properties, .txt, etc. Ensure proper file permissions are granted. The value of AAS_JVM_OPTS does not require quotes.
5. Setting the log file AAS_LOG_FILE with environment variables will not take effect immediately upon first reference; a restart of AAS is required for the change to take effect.

9.1 Settings in the profile file

You need to enable the reference environment variable, `export AAS_CONFIG=true` ; then set the parameters. Save the configuration file and make it effective. Start Apusic Application Server to make it effective.

For information on available properties, see Properties Specific to Environment Variables of Apusic Application Server.

9.2 Reference external file settings

Set the parameter in a file with a file format such as `.conf/.properties/.txt` , for example, properties.conf; then enable the reference environment variable in the format `export AAS_CONFIG=[file_path]` , for example, `export AAS_CONFIG=/opt/properties.conf` ; `[file_path]` requires open permissions; save the configuration file and make it effective, then Start Apusic Application Server to make it effective.

For information on available properties, see Properties Specific to Environment Variables of Apusic Application Server.

9.3 Properties Specific to Environment Variables of Apusic Application Server

Aviriables	Description
AAS_CONFIG	Whether to enable the environment variable setting parameter, true for enabling, false for disabling; the file path can be set.
AAS_HTTP_THREAD_POOL_MAX_QUEUE_SIZE	Configure the maximum queue size of the `http-thread-pool`
AAS_HTTP_THREAD_POOL_MAX_SIZE	Configure the maximum size of the `http-thread-pool`
AAS_HTTP_THREAD_POOL_MIN_SIZE	Configure the minimum size of the `http-thread-pool`
AAS_HTTP_THREAD_POOL_IDLE_TIMEOUT	Configure the idle timeout of the `http-thread-pool`
AAS_THREAD_POOL_1_MAX_QUEUE_SIZE	Configure the maximum queue size of the `thread-pool-1`
AAS_THREAD_POOL_1_MAX_SIZE	Configure the maximum size of the `thread-pool-1`
AAS_THREAD_POOL_1_MIN_SIZE	Configure the minimum size of the `thread-pool-1`
AAS_THREAD_POOL_1_IDLE_TIMEOUT	Configure the idle timeout of the `thread-pool-1`
AAS_ADMIN_THREAD_POOL_MAX_QUEUE_SIZE	Configure the maximum queue size of the `admin-thread-pool`
AAS_ADMIN_THREAD_POOL_MAX_SIZE	Configure the maximum size of the `admin-thread-pool`
AAS_ADMIN_THREAD_POOL_MIN_SIZE	Configure the minimum size of the `admin-thread-pool`
AAS_ADMIN_THREAD_POOL_IDLE_TIMEOUT	Configure the idle timeout of the `admin-thread-pool`
AAS_HTTP_LISTENER_1_PORT	Configure the port of the `http-listener-1`
AAS_HTTP_LISTENER_1_URI_ENCODE	Configure the URI encode of the `http-listener-1`

AAS_HTTP_LISTENER_1_MAX_POST_SIZE	Configure the maximum post size of the `http-listener-1`
AAS_HTTP_LISTENER_1_MAX_CONNECTIONS	Configure the maximum connection of the `http-listener-1`
AAS_HTTP_LISTENER_1_COMPRESS_ENABLE	Configure the whether to enable the compress of the `http-listener-1`
AAS_HTTP_LISTENER_1_COMPRESS_TYPE	Configure the compress type of the `http-listener-1`
AAS_HTTP_LISTENER_1_COMPRESS_SIZE	Configure the compress size of the `http-listener-1`
AAS_HTTP_LISTENER_1_DIABLE_METHOD	Configure whether to enable the method of the `http-listener-1`
AAS_HTTP_LISTENER_1_RELAXEDQUERYCHARS	Configure the relaxed query chars of the `http-listener-1`
AAS_HTTP_LISTENER_1_FILE_CACHE_ENABLE	Configure whether to enable the file cache of the `http-listener-1`
AAS_HTTP_LISTENER_1_FILE_CACHE_LIVE	Configure the file cache live time of the `http-listener-1`
AAS_HTTP_LISTENER_1_FILE_CACHE_MAX_SIZE	Configure the file cache maximum size of the `http-listener-1`
AAS_HTTP_LISTENER_1_FILE_CACHE_MAX_COUNT	Configure the file cache maximum count of the `http-listener-1`
AAS_JVM_OPTS	Set JVM options, multiple JVM options are separated by spaces, and delete operations are preceded by [-] before the parameter
AAS_TCP_BACKLOG	Configure the back log of the `tcp`
AAS_TCP_ACCEPT_THREADS	Configure the accept threads of the `tcp`
AAS_HTTP_LISTENER_2_PORT	Configure the port of the `http-listener-2`
AAS_HTTP_LISTENER_2_URI_ENCODE	Configure the URI encode of the `http-listener-2`
AAS_HTTP_LISTENER_2_MAX_POST_SIZE	Configure the maximum post size of the `http-listener-2`
AAS_HTTP_LISTENER_2_MAX_CONNECTIONS	Configure the maximum connection of the `http-listener-2`
AAS_HTTP_LISTENER_2_COMPRESS_ENABLE	Configure the whether to enable the compress of the `http-listener-2`
AAS_HTTP_LISTENER_2_COMPRESS_TYPE	Configure the compress type of the `http-listener-2`
AAS_HTTP_LISTENER_2_COMPRESS_SIZE	Configure the compress size of the `http-listener-2`
AAS_HTTP_LISTENER_2_DIABLE_METHOD	Configure whether to enable the method of the `http-listener-2`
AAS_HTTP_LISTENER_2_RELAXEDQUERYCHARS	Configure the relaxed query chars of the `http-listener-2`
AAS_HTTP_LISTENER_2_FILE_CACHE_ENABLE	Configure whether to enable the file cache of the `http-listener-2`
AAS_HTTP_LISTENER_2_FILE_CACHE_LIVE	Configure the file cache live time of the `http-listener-2`
AAS_HTTP_LISTENER_2_FILE_CACHE_MAX_SIZE	Configure the file cache maximum size of the `http-listener-2`
AAS_HTTP_LISTENER_2_FILE_CACHE_MAX_COUNT	Configure the file cache maximum count of the `http-listener-2`
AAS_ADMIN_LISTENER_PORT	Configure the port encode of the `admin-listener`
AAS_ADMIN_LISTENER_URI_ENCODE	Configure the URI encode of the `admin-listener`
AAS_ADMIN_LISTENER_MAX_POST_SIZE	Configure the maximum post size of the `admin-listener`
AAS_ADMIN_LISTENER_MAX_CONNECTIONS	Configure the maximum connection of the `admin-listener`
AAS_ADMIN_LISTENER_COMPRESS_ENABLE	Configure the whether to enable the compress of the `admin-listener`
AAS_ADMIN_LISTENER_COMPRESS_TYPE	Configure the compress type of the `admin-listener`
AAS_ADMIN_LISTENER_COMPRESS_SIZE	Configure the compress size of the `admin-listener`
AAS_ADMIN_LISTENER_DIABLE_METHOD	Configure whether to enable the method of the `admin-listener`
AAS_ADMIN_LISTENER_RELAXEDQUERYCHARS	Configure the relaxed query chars of the `admin-listener`
AAS_ADMIN_LISTENER_FILE_CACHE_ENABLE	Configure whether to enable the file cache of the `admin-listener`

AAS_ADMIN_LISTENER_FILE_CACHE_LIVE	Configure the file cache live time of the `admin-listener`
AAS_ADMIN_LISTENER_FILE_CACHE_MAX_SIZE	Configure the file cache maximum size of the `admin-listener`
AAS_ADMIN_LISTENER_FILE_CACHE_MAX_COUNT	Configure the file cache maximum count of the `admin-listener`
AAS_FILE_ROTATION_LIMIT	Configure the file rotation limit of the `logs`
AAS_FILE_MAXHISTORY_FILES	Configure the file max history size of the `logs`
AAS_LOG_FILE	Configure the directory of the `logs`

10 Configuring Circuit Breaker and Rate Limiting

In the realm of distributed systems, ensuring the resilience, stability, and reliability of services is paramount.

Circuit Breaker Pattern

The Circuit Breaker pattern is an architectural pattern used to prevent a failing component from causing cascading failures in a distributed system. It acts as a safeguard that automatically interrupts requests to a failing service, allowing the system to recover gracefully and prevent the failure from spreading to other parts of the system.

Rate Limiting

Rate Limiting is a technique used to control the rate of requests sent to a service, protecting it from being overwhelmed by excessive traffic. By limiting the number of requests a client can make within a given time frame, rate limiting helps to maintain the responsiveness and stability of the service.

10.1 Configuring Rules

There are three types of rules configured under the domain.xml root element:

- flow-rule: rate limiting rule
- degrade-rule: circuit breaker Rule
- system-rule: system rule

Configuring the rules as following:

```
<domain>
<flow-controls>
<system-rule cpu-usage="0.2" resource-name="/jsp/index.jsp" />
</flow-controls>
</domain>
```

10.2 Properties Specific to Circuit Breacker and Rate Limiting

flow-rule:

Attributies	Description	Default Value
resource-name	Configure the url that needs to be throttled, starting with / and including the context path; for example: /jsp/index.jsp, where jsp is the context path	
grade	Configure the value of grade, which can be set to qps or thread; it indicates whether to limit traffic based on qps or the number of threads	thread
count	Configuring the count of qps or thread	
control-behavior	Set the value of control-behavior, which can be set to reject,warm_up,rate_limiter or warm_up_rate_limiter reject: Means to reject the error directly if the limit is exceeded; warm_up: It needs to be used in conjunction with the warm_up_period_sec parameter; slowly rise to the specified tps within the time specified by this parameter; rate_limiter It needs to be used in conjunction with max_queue_time_ms,The traffic is passed at a constant speed according to the specified qps, and the part exceeding the qps is queued. An error is reported only if the queuing time exceeds the specified time of max_queue_time_ms; warm_up_rate_limiter Combines the functions of warm_up and rate_limiter	reject
max-queue-time-ms	Configuring the maxinum queue time of the request	
warm-up-period-sec	Configuring the warm up period of the request	

degrade-rule:

Attributies	Description	Default Value

resource-name	Configure the url that needs to be degrade, starting with / and including the context path; for example: /jsp/index.jsp, where jsp is the context path	
grade	Set the value of grade, which can be set to rt, exception_ratio, exception_count r: Indicates that the slow request ratio is used for circuit-breaking. The definition of slow requests is specified by count, with the unit being ms. The slow request ratio is specified by the slow-ratio-threshold parameter; exception_ratio: Configuring the value of exception ratio, with a value range of 0 to 1 exception_count: The number of exceptions, specified by count	rt
count	Configuring the value of count.The meaning of the value depends on grade	
time-window-seconds	The duration of the fuse, during which all requests will fail. After this time, some requests will be allowed to pass, and then the decision to remove the fuse or continue the fuse will be based on the response results of these requests	
min-request-amount	The minimum number of requests, the default value is 5. Requests within this value are all allowed	5
slow-ratio-threshold	Configuring the slow request ratio, with a value range of 0 to 1	1
stat-interval-ms	Configuring the statistical time period	1000ms

system-rule:

Attributies	Description	Default Value
resource-name	Configure the url that needs to be degrade, starting with / and including the context path; for example: /jsp/index.jsp, where jsp is the context path	
system-load	Configuring the value of system load. This corresponds to the load metric of the operating system's top command. When the system load reaches the set value, traffic will be throttled for the set URL. If the system load exceeds the specified value, it is also necessary to determine whether currentThread > maxSuccessQps * minRt/ 1000 before throttling. Where maxSuccessQps represents the maximum value of qps ever, and minRt represents the minimum response time ever. The product of these two represents the theoretical maximum load of the system	
cpu-usage	Configuring the CPU usage rate, with a value range of 0 to 1; when the CPU usage rate exceeds the set value, traffic will be limited to the set URL	
qps	Configuring the qps number; when the qps exceeds the set value, the set URL will be throttled	
average-rt-ms	Configuring the average response time; when the average response time exceeds the set value, traffic will be limited to the set URL	
max-threads	Configuring the maximum number of threads; when the number of threads exceeds the set value, the set URL will be throttled	5
slow-ratio-threshold	Configuring the slow request ratio, with a value range of 0 to 1; when the ratio of full requests exceeds the set value, traffic will be limited to the set URL	1
stat-interval-ms	Configuring the statistical time period	1000ms

全国统一服务热线

4008-555-800

金蝶天燕云计算股份有限公司(简称“金蝶天燕云”)成立于2000年，前身为“金蝶中间件公司”，是金蝶集团旗下新一代软件基础云平台服务商，云计算国家标准制定企业，国家信创产业核心软件企业。金蝶天燕是国家863重点研发计划与核高基重大专项承接企业，也是“两网一站四库十二金”国家重点工程的基础平台提供商，产品广泛应用于政府、军工、金融、能源等关键行业，累计服务客户总数超过10万家。

Apusic 金蝶天燕

云计算国家标准制定企业

金蝶集团旗下基础软件企业

信息技术应用创新核心企业

官网：www.apusic.com