Oracle® Fusion Middleware

Oracle API Gateway Administrator Guide
11g Release 2 (11.1.2.4.0)

March 2015

Oracle API Gateway Administrator Guide, 11g Release 2 (11.1.2.4.0)

Copyright © 1999, 2015, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services. This documentation is in prerelease status and is intended for demonstration and preliminary use only. It may not be specific to the hardware on which you are using the software. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to this documentation and will not be responsible for any loss, costs, or damages incurred due to the use of this documentation.

The information contained in this document is for informational sharing purposes only and should be considered in your capacity as a customer advisory board member or pursuant to your beta trial agreement only. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle Software License and Service Agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

13 March 2015

Contents

1. What's new 1
    New topics 1
    Updated topics 1

1. API Gateway administration .....

    1. Introduction to API Gateway administration 2
        Overview 2
        API Gateway form factors 2
        Who owns the API Gateway platform and how is it administered? 2
            Operations team 2
            Architecture team 3
        Where do you deploy an API Gateway? 3
        Where do you deploy API Gateway Analytics? 4
        Secure the last mile 4
        API Gateway administration lifecycle 4
    2. Plan an API Gateway system 6
        Overview 6
        Policy development 6
            Policy development guidelines 6
            Example policy requirements 6
        Traffic analysis 7
            Traffic analysis guidelines 7
        Load balancing and scalability 7
            Load balancing guidelines 8
        SSL termination 8
        High Availability and failover 9
            HA stand-by systems 10
            HA and failover guidelines 10
        Backup and recovery 10
            Disaster recovery guidelines 11
        Development staging and testing 11
            Staging and testing guidelines 12
        Hardening—secure the API Gateway 12
            Hardening guidelines 12
        Capacity planning example 13
            Example required throughput 13
            Example development process 13
    3. How API Gateway interacts with existing infrastructure 15
        Overview 15
        Databases 15
        Anti virus 15
        Operations and management 15
        Network firewalls 16
            Advantages over traditional application firewalls 16
            Firewall modes 16
        Application servers 16
        Enterprise Service Buses 17
            Similarities between API Gateways and ESBs 17
            Differences between API Gateways and ESBs 17
        Directories and user stores 17
            Simple inline user store deployment 17
            API Gateway in DMZ—LDAP in LAN 19
            Split deployment between DMZ and LAN 20

Access control 20
Public Key Infrastructure 21
Registries and repositories 22
Software Configuration Managment 22
2. Manage an API Gateway domain .....
1. Configure an API Gateway domain 23
Overview 23
Managedomain script 23
Register a host in a domain 24
Create an API Gateway instance 24
Test the health of an API Gateway instance 25
2. Manage domain topology in API Gateway Manager 26
Overview 26
Manage API Gateway groups 26
Create an API Gateway group 26
Delete an API Gateway group 27
Lock an API Gateway group 27
Manage API Gateway instances 27
Create API Gateway instances 27
Delete API Gateway instances 27
Start API Gateway instances 28
Stop API Gateway instances 28
Edit API Gateway tags 28
Deploy API Gateway configuration 29
Deploy a deployment package 29
Deploy policy and environment packages 29
3. Configure Admin Node Manager high availability and security 31
Overview 31
Hierarchy of SSL certificates in a domain 31
How SSL certificates are generated for domain processes 32
External Certificate Authority 33
Add the first Admin Node Manager to the domain 33
Sign Admin Node Manager certificate with system-generated CA key 33
Sign Admin Node Manager certificate with user-provided CA key 34
Sign Admin Node Manager certificate with external CA 34
Additional certificate generation options 34
Add a Node Manager to the domain 35
Sign Node Manager certificate with system-generated CA key 35
Sign Node Manager certificate with user-provided CA key 35
Sign Node Manager certificate with external CA 35
Additional options 36
Add an API Gateway instance to the domain 36
Sign API Gateway instance certificate with system-generated CA key 36
Sign API Gateway instance certificate with user-provided CA key 36
Sign API Gateway instance certificate with external CA 36
Change a Node Manager to an Admin Node Manager 36
Sign Node Manager certificate with system-generated CA key 37
Sign Node Manager certificate with user-provided CA key 38
Sign Node Manager certificate with external CA 38
Regenerate all SSL certificates in a domain 38
Sign certificates in domain with system-generated CA key 39
Sign certificates in domain with user-provided CA key 39
Sign certificates in domain with external CA 39
Reset passphrase for CA private key 40
Change domain SSL certificate expiry date 40
Admin Node Manager backup and disaster recovery 40
Location of SSL private keys and certificates 41

4. Managedomain command reference 43
Overview 43
Managedomain command interpreter mode 43
View help for a command 43
Run a command 44
Managedomain interactive mode 45
Host Management 45
API Gateway Management 47
Group Management 47
Topology Management 47
Deployment 48
Domain SSL certificates 49
Managedomain command mode 49
3. Manage API Gateway operation .....
1. Start and stop the API Gateway 51
Overview 51
Set passphrases 51
Start the Node Manager 51
Start the API Gateway instance 51
Startup options 52
Connect to the API Gateway in Policy Studio 53
Stop the API Gateway instance 53
Stop the Node Manager 54
2. Start the API Gateway tools 55
Overview 55
Before you begin 55
Launch API Gateway Manager 55
Start Policy Studio 55
3. Run API Gateway as non-root on UNIX/Linux 57
Overview 57
Linux capabilities 57
Before you begin 57
Modify API Gateway file ownership 57
SSL accelerators for HSM 58
Set the CAP_NET_BIND capability on vshell 58
Install the libcap2 package if required 58
API Gateway appliance version 7.1.0 or later 59
Add API Gateway library locations 59
Create the ld.so.conf file 59
Run ldconfig 60
Modify the init.d script to use sudo 60
Modify the jvm.xml file 60
Restart the API Gateway 61
Run API Gateway as non-root on Solaris 61
4. Configure API Gateway high availability 62
Overview 62
HA in production environments 62
Load Balancing 63
Java Message System 63
File Transfer Protocol 64
Remote Hosts 64
Distributed caching 64
External Connections 64
Embedded Apache ActiveMQ 65
Embedded Apache Cassandra database 65
5. API Gateway backup and disaster recovery 67
Overview 67

Components that must be backed up 67
Back up API Gateway 67
Back up API Gateway Analytics 68
Back up databases and third-party systems 68
Disaster recovery plan and tests 68
Example of creating an API Gateway disaster recovery site 68
Back up the production environment 68
Copy to the disaster recovery site 69
Further Information 70
6. Manage API Gateway settings 71
Overview 71
General settings 71
Cache 71
MIME/DIME 71
Namespaces 71
HTTP session 71
Logging settings 71
Transaction Audit Log 71
Transaction Access Log 71
Transaction Event Log 72
Messaging settings 72
Monitoring settings 72
Real Time Monitoring 72
Traffic Monitor 72
Security settings 72
Security Service Module 72
Kerberos 72
Tivoli 73
4. Manage API Gateway security .....
1. Configure an API Gateway encryption passphrase 74
Overview 74
Configure the passphrase in Policy Studio 74
Enter the passphrase when you edit configuration in Policy Studio 74
Provide the passphrase in a configuration file or at startup 75
Enter the Node Manager passphrase in a configuration file 75
Enter the API Gateway passphrase in a configuration file 75
Prompt for the passphrase at server startup 75
Provide the passphrase automatically at startup using a script 76
Promotion between environments 76
Further information 76
2. Manage certificates and keys 77
Overview 77
View certificates and keys 77
Certificate management options 78
Configure an X.509 certificate 78
Create a certificate 78
Import certificates 79
Configure a private key 79
Private key stored locally 80
Private key provided by OpenSSL engine 80
Private key stored on external HSM 81
Configure HSMs and certificate realms 81
Manage HSMs with keystoreadmin 81
Step 1—Register an HSM provider 82
Step 2—Create a certificate realm and associated keystore 82
Step 3—Start the API Gateway when using an HSM 83
Configure SSH key pairs 84

Add a key pair 84
Manage OpenSSH keys 85
Configure PGP key pairs 85
Add a PGP key pair 85
Manage PGP keys 86
Global import and export options 86
Import and export certificates and keys 86
Manage certificates in Java keystores 86
Further information 87
3. Manage API firewalling 88
Overview 88
Configure API firewalling 88
Enable threat protection 88
Configure threat protection rules 89
Monitor API firewalling 90
Further information 91
5. Deploy API Gateway configuration .....
1. Manage API Gateway deployments 92
Overview 92
Connect to a server in Policy Studio 92
Edit a server configuration in Policy Studio 92
Manage deployments in API Gateway Manager 93
Compare and merge configurations in Policy Studio 93
Manage Admin users in API Gateway Manager 93
Configure policies in Policy Studio 93
2. Deploy API Gateway configuration 94
Overview 94
Create a package in Policy Studio 94
Configure package properties in Policy Studio 94
Deploy packages in Policy Studio 95
Deploy a factory configuration in Policy Studio 96
Deploy currently loaded configuration in Policy Studio 96
Push configuration to a group in Policy Studio 96
View deployment results in Policy Studio 96
Deploy on the command line 97
Deploy packages in API Gateway Manager 97
6. Monitoring and reporting .....
1. Monitor services in API Gateway Manager 98
Overview 98
Ensure monitoring is enabled 98
View real-time monitoring 98
View traffic monitoring 99
Filter message traffic 99
View message content 99
View performance statistics 101
Filter performance statistics 101
Detect malformed messages 101
Monitor real-time metrics 102
Configure dynamic trace, logging, and monitoring 103
Monitor and report on services with API Gateway Analytics 103
2. Redact sensitive data from API Gateway Manager 104
Overview 104
Redaction configuration 104
Enable redaction for an API Gateway 105
Redact HTTP message content 105
URL path matching 105
Supported HTTP features 106

Example: redact an HTTP Basic authorization header 106
Redact JSON message content 107
JSON redactor configuration 107
Example: redact OAuth message tokens from a JSON message 107
Redact XML message content 108
XML redactor configuration 108
Example: redact a WS-Security username token from a XML message 109
Redact HTML form message content 110
Redact raw message content 110
Example: redact credit card details from raw text 111
Redact sensitive data from log files 111
3. Configure API Gateway for API Gateway Analytics 112
Overview 112
Connect to the API Gateway in Policy Studio 112
Configure the metrics database connection 112
Configure transaction audit logging to the metrics database 113
Configure the API Gateway to write to the transaction event log 113
Deploy the updated configuration to the API Gateway 113
Configure the Node Manager to process event logs and update the metrics database 114
Use the managedomain interactive menu 114
Use the managedomain command options 114
Configure additional options for event log processing in the Node Manager 115
Further information 116
4. Monitoring and reporting with API Gateway Analytics 117
Overview 117
Launch API Gateway Analytics 117
Monitor the API Gateway system 117
Systems 117
System resources 118
Monitor API services, methods, and clients 119
Example: API service performance 119
Monitor remote hosts 120
Monitor protocols 121
Audit transactions 122
Schedule custom reports 123
Create reports in a monitoring view 123
Using the reports view 123
Further information 124
5. Configure scheduled report settings 126
Overview 126
Database configuration 126
Scheduled reports configuration 126
SMTP configuration 126
6. Purge the metrics database 128
Overview 128
Run the dbpurger command 128
Example commands 128
7. Troubleshoot your API Gateway installation .....
1. Configure API Gateway logging and events 130
API Gateway logs and events 130
Domain management and diagnostics 130
Message transactions 130
Events and alerts 130
Configure audit logs per domain 130
View in API Gateway Manager 131
View the domain audit log file 131
Redact domain audit log output 132

Configure transaction audit log destinations 133
Configure transaction audit logs per filter 133
Message payload logging 134
Configure transaction event logs per API Gateway 134
Configure transaction access logs per path 134
Manage API Gateway events and alerts 134
Configure dynamic trace and log settings 135
Further information 136
2. Configure API Gateway diagnostic trace 137
Overview 137
View API Gateway trace files 137
Set API Gateway trace levels 137
Configure API Gateway trace files 138
Run trace at DEBUG level 139
Run trace at DATA level 141
Integrate trace output with Apache log4J 142
Get help with API Gateway 142
3. API Gateway performance tuning 144
Overview 144
General performance tuning 144
Minimize tracing 144
Disable real-time monitoring 145
Disable traffic monitoring 145
Disable transaction logging 146
Disable Access logging 146
Advanced performance tuning 146
Configure spill to disk 146
Configure database pooling 147
Configure HTTP keep alive 148
Configure chunked encoding 148
Single test client and sever 149
JVM memory 149
Number of client threads on Linux 150
Multiple connection filters 150
8. Manage user access .....
1. Manage API Gateway users 151
Overview 151
API Gateway users 151
Add API Gateway users 151
API Gateway user attributes 151
API Gateway user groups 152
Add API Gateway user groups 152
Update API Gateway users or groups 152
2. Manage Admin users 153
Overview 153
Admin user privileges 153
Admin user roles 153
Add a new Admin user 154
Remove an Admin user 154
Reset an Admin user password 154
Manage Admin user roles 155
3. Configure Role-Based Access Control (RBAC) 156
Overview 156
Local admin user store 157
Admin user password storage algorithm 158
Configure an LDAP repository to store credentials 159
RBAC Access Control List 159

Configure RBAC users and roles 160
Management service roles and permissions 161
4. Active Directory for authentication and RBAC of management services 164
Overview 164
Step 1: create an Active Directory group 164
Step 2: create an Active Directory user 165
Step 3: create an LDAP connection 168
Step 4: create an LDAP repository 169
Step 5: create a test policy for LDAP authentication and RBAC 171
Step 6: use the LDAP policy to protect management services 178
Add an LDAP user with limited access to management services 178
5. OpenLDAP for authentication and RBAC of management services 180
Overview 180
Step 1: create an OpenLDAP group for RBAC roles 180
Step 2: add RBAC roles to the OpenLDAP RBAC group 181
Step 3: add users to the OpenLDAP RBAC group 184
Step 4: create an LDAP connection 185
Step 5: create an OpenLDAP repository 185
Step 6: create a test policy for LDAP authentication and RBAC 187
Step 7: use the OpenLDAP policy to protect management services 193
9. Manage network-level settings .....
1. Configure a DNS service with wildcards for virtual hosting 194
Overview 194
DNS workflow 194
BIND DNS software 194
Configure a wildcard domain 195
Configure DNS options 195
Configure default zones 196
Configure logging 196
Configure a wildcard domain 197
Configure domain zone files 197
10. Manage ActiveMQ messaging .....
1. Manage embedded ActiveMQ messaging 198
Overview 198
Manage messaging queues 198
Manage messages in a queue 198
Create a new message 199
View message contents 200
Manage messaging topics 200
Manage messaging subscribers 201
Create a new subscriber 201
Manage messaging consumers 202
11. API Gateway settings reference .....
1. General settings 203
Overview 203
Settings 203
2. MIME/DIME settings 206
Overview 206
Configuration 206
3. Namespace settings 207
Overview 207
SOAP Namespace 207
Signature ID Attribute 207
WSSE Namespace 208
4. HTTP Session settings 209
Overview 209
Configuration 209

5. Transaction audit log settings 210
Overview 210
Configure log output 210
Log to Text File 210
Log to XML File 211
Log to Database 211
Log to Local Syslog 212
Log to Remote Syslog 212
Log to System Console 212
6. Transaction access log settings 213
Overview 213
Access log format 213
Configure the access log 213
Redact sensitive details from the access log 214
7. Transaction event log settings 216
Overview 216
Transaction event log formats 216
Event log header entries 216
Event log system entries 217
Event log transaction entries 217
Configure the transaction event log 218
8. Embedded ActiveMQ settings 220
Overview 220
General messaging settings 220
SSL settings 220
Authentication settings 221
9. Traffic monitoring settings 222
Overview 222
Configuration 222
10. Real-time monitoring metrics 224
Overview 224
Enable monitoring 224
Configure real-time metrics 224

What's new

New topics

This release of the *API Gateway Administrator Guide* contains the following new topics:

- *Manage API firewalling*
- *Transaction event log settings*

Updated topics

This release of the *API Gateway Administrator Guide* contains the following updated topics:

- *Manage domain topology in API Gateway Manager*
- *Configure Admin Node Manager high availability and security*
- *Managedomain command reference*
- *Configure API Gateway high availability*
- *Manage API Gateway settings*
- *Configure an API Gateway encryption passphrase*
- *Manage certificates and keys*
- *Monitor services in API Gateway Manager*
- *Configure API Gateway for API Gateway Analytics*
- *Monitoring and reporting with API Gateway Analytics*
- *Configure scheduled report settings*
- *Configure API Gateway logging and events*
- *Transaction access log settings*
- *Traffic monitoring settings*
- *Real-time monitoring metrics*
- *Embedded ActiveMQ settings*

Introduction to API Gateway administration

Overview

Oracle API Gateway is a comprehensive platform for managing, delivering, and securing APIs. This guide explains how to administer and manage the API Gateway platform. This chapter introduces the main issues involved in API Gateway administration.

Tip

For an introduction to API Gateway features, tools, and architecture, see the *API Gateway Concepts Guide*.

API Gateway form factors

The API Gateway is available in software and hardware form factors. The available platforms are as follows:

- Software installation (for details on supported platform versions, see the *API Gateway Installation Guide*)

Who owns the API Gateway platform and how is it administered?

The API Gateway platform is administered by the following groups:

- Operations—Runtime management of message traffic, logs and alerts, and high availability is performed by Operations staff.
- Architecture—Design-time policy definition, which determines the behavior of the API Gateway platform, is performed by Security Architects and Systems Architects.

Operations team

Operations staff are responsible for making sure that the API Gateway platform is running correctly. They are concerned with the following problems:

- System status and health
- Network connectivity
- Security alerts
- System security
- Backups and recovery
- Maintenance of logs

The API Gateway platform provides a web-based console named API Gateway Manager that is dedicated to the Operations team. The API Gateway Manager includes the following features:

- Dashboard for monitoring overall system health and network topology.
- Real-time monitoring and metrics for all messages processed by the API Gateway.
- Traffic monitoring to quickly isolate failed or blocked message transactions and provide detailed information about each transaction, payload, and so on.
- API Gateway logs, trace, events, and alerts.
- Messaging based on Java Message System (JMS). This includes managing queues, topics, subscribers, consumers, and messages in Apache ActiveMQ.

- Dynamic system settings, user roles, and credentials.

As well as providing operational management functionality of its own, the API Gateway also interoperates with third-party network operations tools such as HP OpenView, BMC Control, and CA UniCenter. Finally, all functionality available in the API Gateway Manager is also available as a REST API.

Architecture team

System architects and security architects have an overarching view of enterprise IT infrastructures, and so are more concerned with using the API Gateway to help integrate and secure existing enterprise systems. Architects wish to create API policies and integrate with third-party systems.

Policy Studio is a rich API Gateway policy development tool that enables architects and policy developers to create and control API Gateway policies. You can use Policy Studio to visually define policy workflows in a drag-n-drop environment. This means that configuration is performed at the systems architecture level, without needing to write code.

For more details on using Policy Studio to create API Gateway policies, see the *API Gateway Policy Developer Guide*.

Where do you deploy an API Gateway?

API Gateways can be deployed in the Demilitarized Zone (DMZ) or in the Local Area Network (LAN) depending on policies or requirements, as shown in the following diagram:

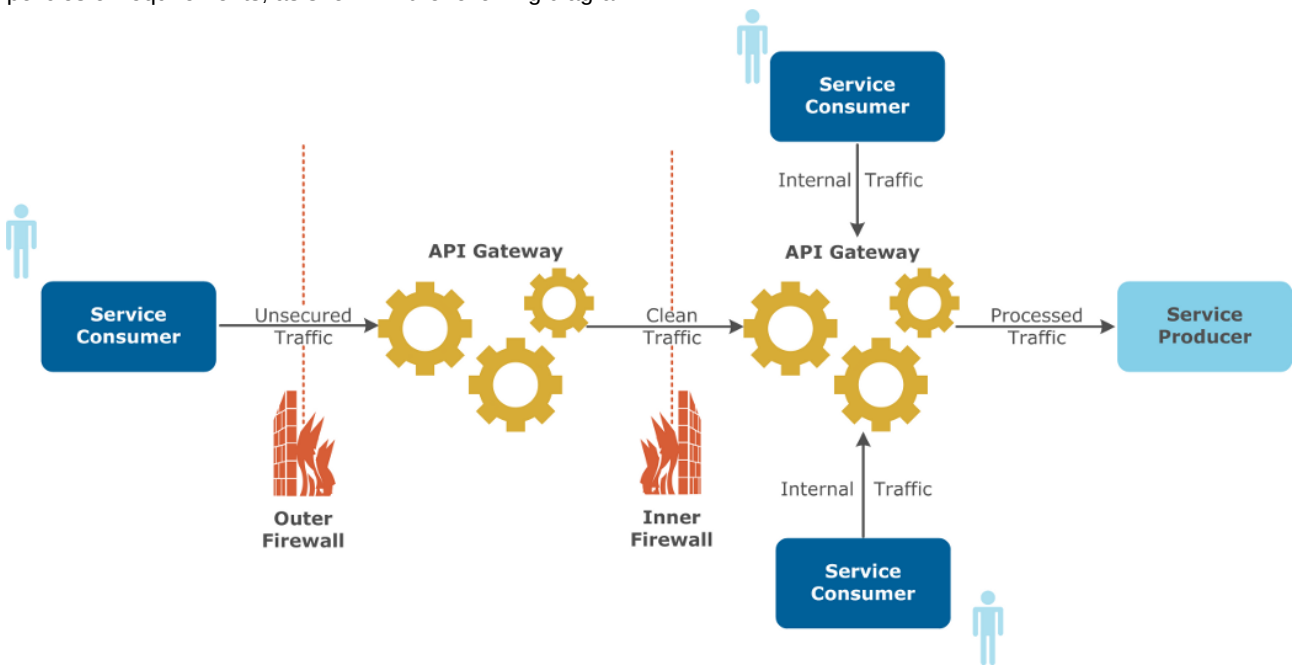

The following guidelines help you to decide where to deploy the API Gateway:

- If you are processing only traffic from external sources, consider locating the API Gateway in the DMZ. If the API Gateway is also processing internal traffic, consider locating it in the LAN.
- If you are processing traffic internally and externally, a combination of API Gateways in the DMZ and internally on the LAN is considered best practice. The reason for this is that different policies should be applied to traffic depending on its origin.
- Both internal and external traffic should be checked for threats and to make sure that they contain the correct parameters for REST API requests, or correspond to Web service definitions.

- External traffic carries a greater potential risk and should be scanned by the API Gateway located in the DMZ to make sure that it does not in any way affect the performance of internal applications.
- Internal traffic and pre-scanned external traffic should then be processed by the API Gateway located in the LAN. This type of checking includes:
  - Checking API service level agreements and enforcing throttle threshold levels
  - Integration with a wide range of third-party systems
  - Web service standards support

Where do you deploy API Gateway Analytics?

Although you can select the API Gateway Analytics component in the API Gateway installer, it is good practice to install API Gateway Analytics on a separate host from API Gateway installations. You should ensure that API Gateway Analytics runs on a dedicated host, or on a host that is not a running an API Gateway instance and/or Node Manager.

You can deploy API Gateway Analytics on any supported host platform (for example, Windows, Linux, or Solaris) depending on its availability in your architecture. For more details on supported platform versions, see the *API Gateway Installation Guide*.

Important

API Gateway Analytics supports a range of databases for storing historic reports and metrics (for example, Oracle, DB2, MySQL, and Microsoft SQL Server). It is not advised to install the database used for API Gateway Analytics in the DMZ. You should install this database in the LAN on a separate host from your API Gateway installations.

You can secure the connection to the API Gateway Analytics database by dedicating it to one IP address. For more details on configuring the API Gateway Analytics database, see the *API Gateway Installation and Configuration Guide*.

Secure the last mile

Securing the last mile refers to preventing internal users from directly accessing services without going through the API Gateway. This can be achieved in multiple ways. You should carefully choose which option is best for your use case, taking into account the security level you want to achieve, and the impact on performance the solution will have. You should choose from the following approaches:

- **Controlling traffic at the network level**: Services can only be accessed if the traffic is coming from pre-approved IP addresses. This is the simplest solution to put in place, is very secure, and has no impact on performance or existing applications.
- **Establishing a mutual SSL connection between API Gateways and services**: This solution is the easiest to put in place and has little to no impact on existing applications. However, it does have a non-negligible impact on latency.
- **Passing authentication tokens from API Gateways to back-end services**: This involves passing authentication tokens for WS-Security, Security Assertion Markup Language (SAML), and so on. This solution has a low impact on latency but requires some development because the target service container must validate the presence and the contents of the token.

For more details on configuring mutual SSL, and configuring WS-Security and SAML authentication tokens, see the *API Gateway Policy Developer Guide*.

API Gateway administration lifecycle

The main stages in the overall API Gateway administration lifecycle are as follows:

1. Planning an API Gateway system. This is described in the next topic.
2. Installing API Gateway components. See the *API Gateway Installation and Configuration Guide*.
3. Configuring a domain. This is described in Chapter 5, Deploy API Gateway configuration.
4. Operating and managing the API Gateway. This is described in the rest of this guide.
5. Upgrading an API Gateway. See the *API Gateway Installation and Configuration Guide*.

Plan an API Gateway system

Overview

One of the most important tasks when deploying an API Gateway system is confirming that the system is fit for purpose. Enterprise software systems are hugely valuable to the overall success of the business operation. For any organization, there are many implications of system downtime with important consequences to contend with. The most important factors to look at when architecting an API Gateway deployment are as follows:

- What type of API Gateway policies will run on the system?
- What type of message traffic will the API Gateway be required to process?
- What is the expected distribution of the traffic?
- How important is the system, and does it need to be always available?
- How to make changes to the system without incurring system outage
- Planning for outage and disaster recovery

Policy development

The functional characteristics of any given policy run by an API Gateway can have a huge effect on the overall system throughput and latency times. Depending on the purpose of a particular policy, the demand on valuable processing power will vary. The following guidelines apply in terms of processing power:

- Threat analysis and transport-based authentication tasks are relatively undemanding.
- XML processing such as XML Schema and WS-Security username/password authentication are slightly more intensive.
- Calling out to third-party systems is expensive due to network latency.
- Cryptographic operations like encryption and signing are processor intensive.

The key point is that API Gateway policy performance depends on the underlying requirements, and customers should test their policies before deploying them into a production environment.

Policy development guidelines

Architects and policy developers should adhere to the following guidelines when developing API Gateway policies:

- Decide what type of policy you need to process your message traffic. Think in terms of functional requirements instead of technologies. Oracle can help you to map the technologies to the requirements. Example functional requirements include the following:
  - *Only trusted clients should be allowed send messages into the network*
  - *An evidential audit trail should be kept*
- Think about what you already have in your architecture that could help to achieve these aims. Examples include LDAP directories, databases that already have replication strategies in place, and network monitoring tools.
- Create a policy to match these requirements and test its performance. Oracle provides an integrated performance testing tool (Oracle API Gateway Explorer) to help you with this process.
- Use the Oracle API Gateway Manager console to help identify what the bottlenecks are in your system. If part of the solution is slowing the overall system, try to find alternatives to meet your requirements.
- View information on historic network usage in Oracle API Gateway Analytics.
- Test the performance capability of the backend services.

Example policy requirements

Supplier A is creating a service that will accept Purchase Order (PO) documents from customers. The PO documents are formatted using XML. The functional requirements are:

- The service should not accept anything that will damage the PO system.
- Incoming messages must to be authenticated against a customer database to make sure they come from a valid customer account.
- The supplier already has an LDAP directory and would like to use it to store the customer accounts.
- The supplier must be able prove that the message came from the customer.

These requirements can be achieved using a policy that includes processing for Threatening Content and checking the XML Signature, which verifies the certificate against the LDAP Directory. For details on how to develop API Gateway policies using Policy Studio, see the *API Gateway Policy Developer Guide*.

Traffic analysis

In the real world, messages do not arrive in a continuous stream with a fixed size like a lab-based performance test. Message traffic distribution has a major impact on system performance. Some of the questions that need to be answered are as follows:

- Is the traffic smooth or does it arrive in bursts?
- Are the messages all of the same size? If not, what is the size distribution?
- Is the traffic spread out over 24 hours or only during the work day?

Traffic analysis guidelines

You should adhere to the following guidelines when analyzing message traffic:

- Use the **Traffic** tab in the API Gateway Manager web console to analyze message traffic. For more details, see *Monitor services in API Gateway Manager*.
- Use the API Gateway Analytics web console to analyze historical message traffic. For more details, see *Monitoring and reporting with API Gateway Analytics*.
- Take traffic distribution into account when calculating performance requirements.
- Take message size distribution into account when running performance tests.
- If traffic bursts cause problems for service producers, consider using the API Gateway to smooth the traffic (for example, using the Throttling filter). For more details, see the *API Gateway Policy Developer Guide*.

Load balancing and scalability

The API Gateway scales well both horizontally and vertically. Customers can scale horizontally by adding more API Gateways to a cluster and load balancing across it using a standard load balancer. API Gateways being load balanced run the same configuration to virtualize the same APIs and execute the same policies. If multiple API Gateway groups are deployed, load balancing should be across groups also.

For example, the following diagram shows load balancing across two groups of API Gateways deployed on two hosts:

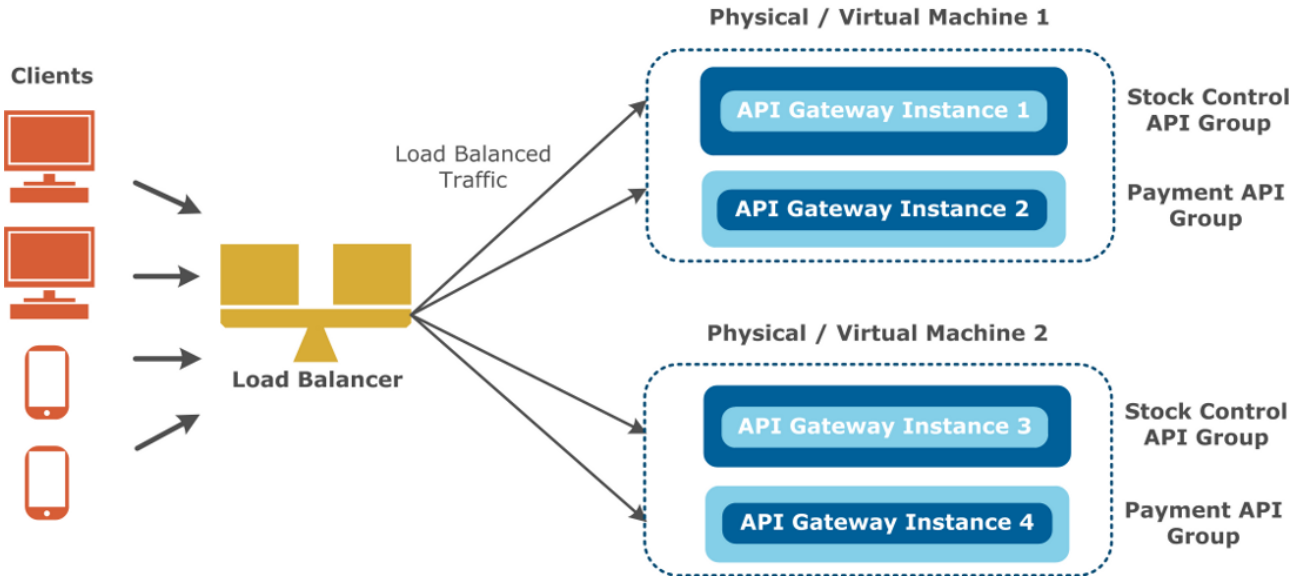

The API Gateway imposes no special requirements on load balancers. Loads are balanced on a number of characteristics including the response time or system load. The execution of API Gateway policies is stateless, and the route through which a message takes on a particular system has no bearing on its processing. Some items such as caches and counters are held on a distributed cache, which is updated on a per message basis. As a result, API Gateways can operate successfully in both sticky and non-sticky modes. For more details on caching, see the *API Gateway Policy Developer Guide*.

The distributed state poses a number of questions in terms of active/active and active/passive clustering. For example, if the counter and cache state is important, you must design your overall system so that at least one API Gateway is active at all times. This means that for a resilient HA system, a minimum of at least two active API Gateways at any one time, with a third and fourth in passive mode is recommended.

The API Gateway ensures zero downtime by implementing configuration deployment in a rolling fashion. For example, while each API Gateway instance in the cluster or group takes a few seconds to update its configuration, it stops serving new requests, but all existing in-flight requests are honored. Meanwhile, the rest of the cluster or group can still receive new requests. The load balancer ensures that requests are pushed to the nodes that are still receiving requests. For more details on deploying API Gateway configuration, see the *Deployment and Promotion Guide*.

Load balancing guidelines

Oracle recommends the following guidelines for load balancing:

- Use the Admin Node Manager and API Gateway groups to maintain the same policies on load-balanced API Gateways. For more details, see the *API Gateway Concepts Guide*.
- Configure alerts to identify when API Gateways and backend services are approaching maximum capacity and need to be scaled.
- Use the API Gateway Manager console to see which parts of the system are processing the most traffic.

SSL termination

Secure Socket Layer (SSL) connections can be terminated at the load balancer or API Gateway level. These options are described as follows:

- **SSL connection is terminated at load balancer**:

    - The SSL certificate and associated private key are deployed on the load balancer, and not on the API Gateway. The subject name in the SSL certificate is the fully qualified domain name (FQDN) of the server (for example, `oracle.com`).
    - The traffic between the load balancer can be in the clear or over a new SSL connection. The disadvantage of a new SSL connection is that it puts additional processing load on the load balancer (SSL termination and SSL establishment).
    - If mutual (two-way) SSL is used, the load balancer can insert the client certificate into the HTTP header. For example, the F5 load balancer can insert the entire client certificate in `.pem` format as a multi-line HTTP header named `XClient-Cert` into the incoming HTTP request. It sends this header to the API Gateway, which uses it for validation and authentication.
- **Load balancer is configured for SSL pass-through, and all traffic passed to the API Gateway**:
  With SSL pass-through, the traffic is encrypted so the load balancer cannot make any layer seven decisions (for example, if HTTP 500 is returned by API Gateway, route to HA API Gateway). To avoid this problem, you can configure the API Gateway so that it closes external ports on defined error conditions. In this way, the load balancer is alerted to switch to the HA API Gateway.

Note

The API Gateway can also optionally use a cryptographic accelerator for SSL termination. For more details, see the *API Gateway Policy Developer Guide*.

High Availability and failover

API Gateways are used in high value systems, and customers typically deploy them in High Availability (HA) mode to protect their investments. The API Gateway architecture enables this process as follows:

- The Admin Node Manager is the central administration server responsible for performing all management operations across an API Gateway domain. It provides policy synchronization by ensuring that all API Gateways in an HA cluster have the same policy versions and configuration. For details on Admin Node Manager HA and API Gateway group-based architecture, see the *API Gateway Concepts Guide*.
- API Gateway instances are stateless by nature. No session data is created, and therefore there is no need to replicate session state across API Gateways. However, API Gateways can maintain cached data, which can be replicated using a peer-to-peer relationship across a cluster of API Gateways. For more details, see the *API Gateway Policy Developer Guide*.
- API Gateway instances are usually deployed behind standard load balancers which periodically query the state of the API Gateway. If a problem occurs, the load balancer redirects traffic to the hot stand-by machine.
- If an event or alert is triggered, the issue can be identified using API Gateway Manager or API Gateway Analytics, and the active API Gateway can then be repaired.

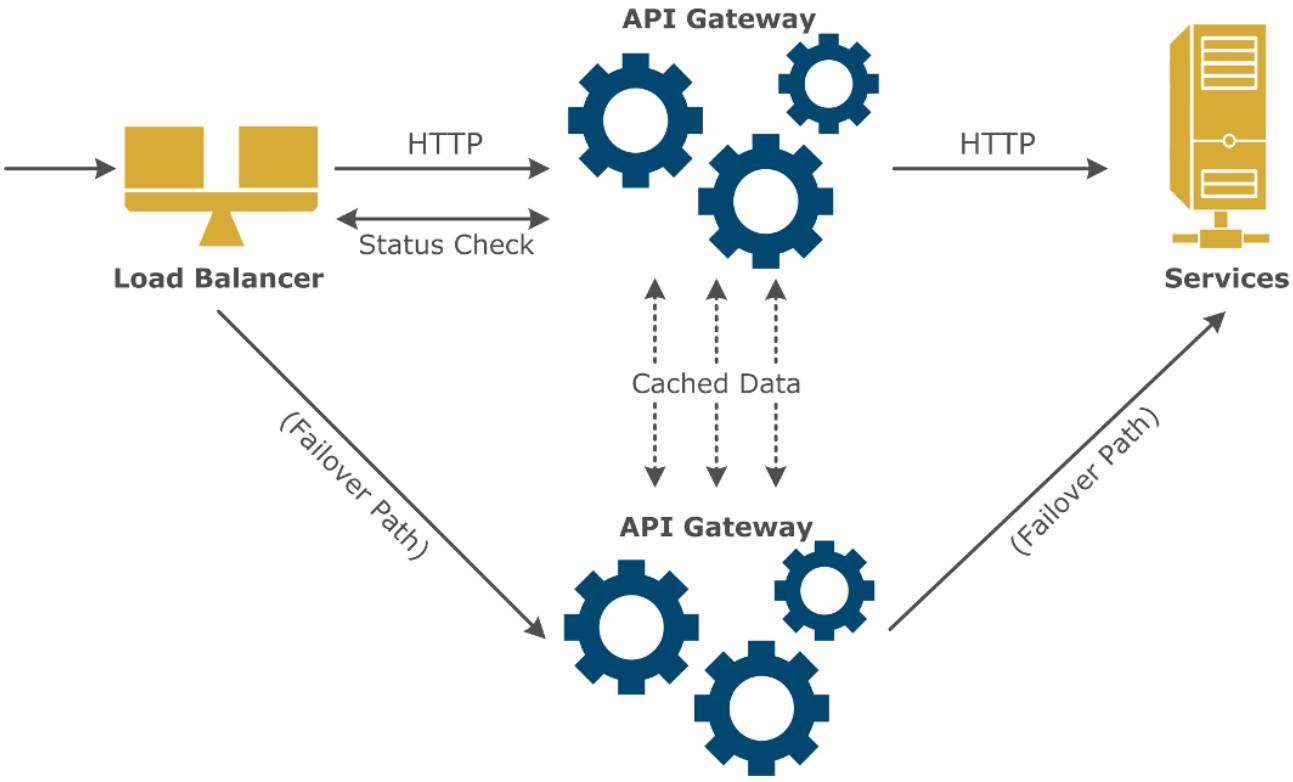

HA stand-by systems

High Availability can be maintained using hot, cold, or warm stand-by systems. These are described as follows:

- **Cold stand-by**: System is turned off.
- **Warm (passive) stand-by**: System is operational but not containing state.
- **Hot (active) stand-by**: System is fully operational and with current system state.

HA and failover guidelines

Oracle recommends the following guidelines for HA stand-by systems:

- For maximum availability, use an API Gateway in hot stand-by for each production API Gateway.
- Use API Gateways to protect against malicious attacks that undermine availability.
- Limit traffic to backend services to protect against message flooding. This is particularly important with legacy systems that have been recently service-enabled. Legacy systems may not have been designed for the traffic patterns to which they are now subjected.
- Monitor the network infrastructure carefully to identify issues early. You can do this using API Gateway Manager and API Gateway Analytics. Interfaces are also provided to standard monitoring tools such as syslog and Simple Network Management Protocol (SNMP).

Backup and recovery

Most customers have a requirement to keep a mirrored backup and disaster recovery site with full capacity to be able to recover from any major incidents. These systems are typically kept in a separate physical location on cold stand-by until

the need arises for them to be brought into action.

Disaster recovery guidelines

The following applies for backup and recovery:

- The backup and disaster site must be a full replica of the production site (for example, with the same number of API Gateway instances).
- The Admin Node Manager helps get backup and recovery sites up and running fast by synchronizing the API Gateway policies in the backup solution.
- Remember to also include any third-party systems in backup and recovery solutions.

Development staging and testing

The most common reason for system downtime is change. Customers successfully alleviate this problem through effective change management as part of a mature software development lifecycle. A software development lifecycle controls change by gradually pushing it through a series of stages until it reaches production.

Each customer will have their own approach to staging depending on the value of the service and the importance of the data. Staging can be broken into a number of different milestones. Each milestone is intended to isolate a specific type of issue that could lead to system downtime. For example:

- The development stage is where the policy and service are created.
- Functional testing makes sure the system works as intended.
- Performance testing makes sure the system meets performance requirements.
- System testing makes sure the changes to the system do not adversely affect other parts.

In some cases, each stage is managed by a different group. The number of API Gateways depends on the number of stages and requirements of each of these stages. The following diagram shows a typical environment topology that includes separate API Gateway domains for each environment:

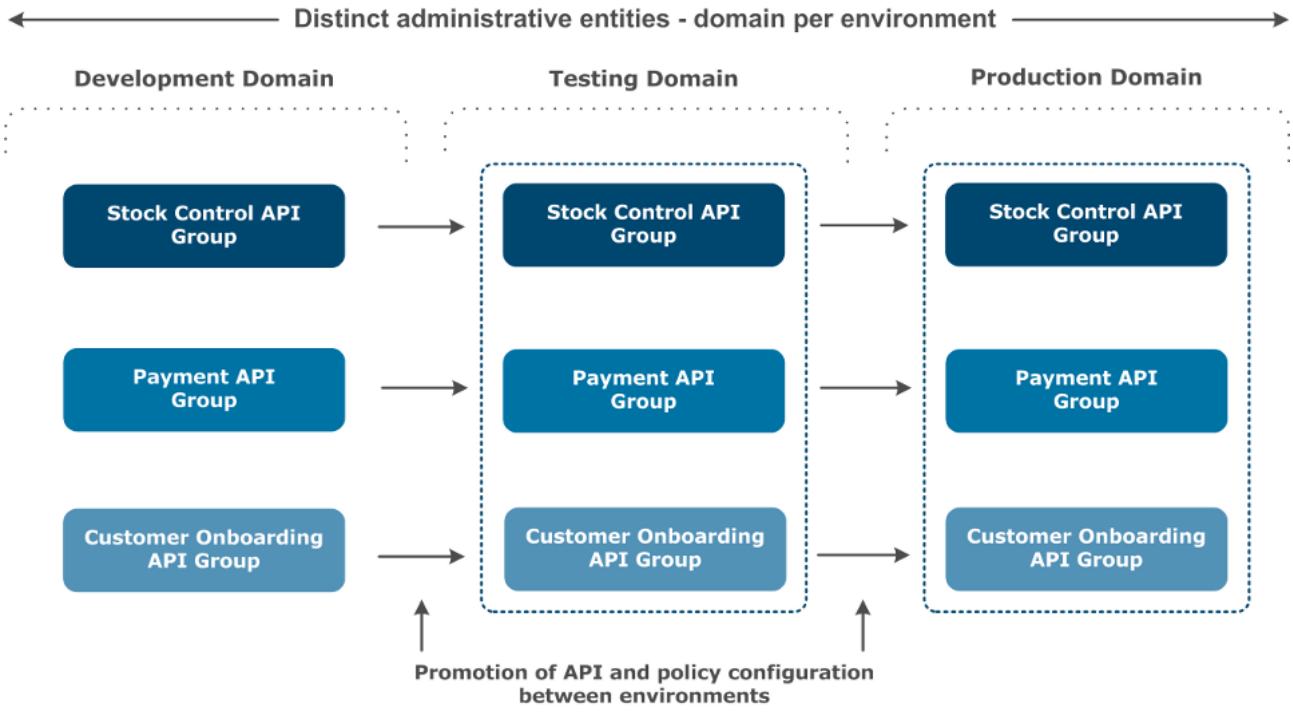

Staging and testing guidelines

The following guidelines apply to development staging and testing:

- Use API Gateway configuration packages (`.fed`, `.pol`, and `.env`) to control the migration of policies from development through to production.
- Keep an audit trail of all system changes.
- Have a plan in place to roll back quickly in the event of a problem occurring.
- Test all systems and policy updates before promoting them to production.
- Test High Availability and resiliency before going into production.

For more details, see the *API Gateway Deployment and Promotion Guide*.

Hardening—secure the API Gateway

The API Gateway platform is SSL-enabled by default, so you do not need to SSL-enable each API Gateway component. However, in a production environment, you must take additional precautions to ensure that the API Gateway environment is secure from external and internal threats.

Hardening guidelines

The following guidelines apply to securing the API Gateway:

- You must change all default passwords (for example, change the password for Policy Studio and API Gateway Manager). For more details, see *Manage Admin users*.
- The default X.509 certificates used to secure API Gateway components are self-signed (for example, the certificate used by the API Gateway Manager on port `8090`). You can replace these self-signed certificates with certificates issued by a Certificate Authority (CA). For more details, see *Manage certificates and keys*.
- By default, API Gateway management ports bind on all HTTP interfaces (IP address set to `*`). You can change

these ports or restrict them to bind on specific IP addresses instead. For more details, see the chapter on "Configuring HTTP Services" in the *API Gateway Policy Developer Guide*.

- By default, API Gateway configuration is unencrypted. You can specify a passphrase to encrypt API Gateway instance configuration. For more details, see *Configure an API Gateway encryption passphrase*.
- You can configure user access at the following levels:
  - Policy Studio, API Gateway Manager, and Configuration Studio users—see *Manage Admin users*.
  - API Gateway users—see *Manage API Gateway users*.
  - Role-based access—see *Configure Role-Based Access Control (RBAC)*.

Capacity planning example

The following is an example formula for deciding how big your API Gateway deployment will need to be. This example is purely intended for illustration purposes only:

```
numberOfGateways =
  (ceiling (requiredThroughput x factorOfSafety / testedPerformance) x HA) x (1 + numDR)
             + staging + eng))
```

This formula is described as follows:

- *factorOfSafety* has a value of 2 for normal.
- Typically High Availability (*HA*) has a value of 2
- *ceiling* implies that the number is rounded up. It is very important that this is not under provisioned by accident.
- Typically customers have a single backup and disaster recovery site (*numDR*=1).
- There is a large variation in performance between API Gateways depending on the policy deployed so it important to performance test policies prior to final capacity planning.
- *staging* is the amount of stage testing licenses. Customers with very demanding applications should have a mirror of their production system. This is the only way to be certain.

Example required throughput

For example, a software system is being designed that will connect an automobile manufacturer's purchasing system with a supplier. The system is designed to meet the most stringent load and availability requirements. An API Gateway system is being used to verify the integrity of the purchase orders.

The system as a whole is required to process 10,000,000 transactions per day. Each of these transactions will result in 8 individual request/responses. 90% of the load will happen between 9am in the morning and 5pm in the evening. The customer wishes to have a factor of safety of 2 for load spikes.

```
Overall throughput = (8 x 10,000,000 x 0.9)/8/60/60= 2,500 transactions per second (tps)
```

Example development process

The development process for the example system should be as follows:

1. A Signature Verification policy is developed by one of the system engineers using a software license on his PC. When he is satisfied with the policy, he pushes it forward to the test environment.
2. Test engineers deploy the policy to the interoperability testing environment. They make sure that all of the systems work correctly together, and test the system as a whole. The API Gateway must show it can process the signature and integrate with the service backend and PKI systems. When the system as a whole is working correctly, the testing moves to system testing.
3. System testing will prove that the system as a whole is reliable and has the capacity to meet the performance requirements. It will also make sure that the system interoperates with monitoring software and behaves properly dur-

ing failover and recovery. During system test the API Gateway is found to be capable of processing 1,000 messages per second for this policy.

4. The customer will have the following infrastructure:
   - 1 backup and disaster site which is a full replica of the production site
   - A full HA system on hot standby
   - 2 API Gateways for stage testing
   - A factor of safety on the load of 2
5. The number of API Gateways should be as follows:

   ```
   ((ceiling(2500/1000)*2)*2)*(1+1)+2+1 = 27
   ```

   The resulting architecture provisioned should be 27 API Gateways made up of the following:
   - 12 production licenses (6 live, 6 standby)
   - 12 backup and disaster recovery licenses
   - 2 staging licenses
   - 1 engineering license

How API Gateway interacts with existing infrastructure

Overview

As part of API Gateway policy execution, the API Gateway needs to interact with various components of the existing infrastructure. For example, these include external connections to systems such as databases, LDAP servers, third-party identity management products, and so on. This topic describes how the API Gateway interacts with these components in your existing network infrastructure.

Databases

API Gateway interoperates with a range of commonly used databases for a wide variety of purposes. For example, this includes managing access control credentials, authorization attributes, identity and role information, and logging and reporting. API Gateway uses Java Database Connectivity (JDBC) to manage database connections. The following databases are supported:

- MySQL
- Oracle
- DB2
- Microsoft SQL Server

For details on supported versions, see the *API Gateway Installation Guide*.

Anti virus

API Gateway supports Anti-Virus (AV) scanning using virus scanning engines from third-party vendors such as McAfee, Sophos, and ClamAV. Depending on the engine, this scanning can happen in-process or remotely using a client SDK or Internet Content Adaptation Protocol (ICAP). The API Gateway takes message attachments, passes them to the AV scanner, and then acts on the decision. The following conditions apply:

- AV signature distribution and updates are performed using mechanisms of the antivirus vendor.
- Licensing for the AV engine is performed through the AV vendor distribution channels.

For more details, see the AntiVirus Integration Guides in the *API Gateway Interoperability Guide*.

Operations and management

API Gateway has a number of different options for operations and management:

- Detailed transaction logs can be sent to syslog (UDP), relational databases, or flat files. These contain detailed records of processed messages, their contents, how long processing took, and decisions taken during message processing. This type of logging can also include information alerts about policy execution failures and breached Service Level Agreements (SLAs), and information about critical events such as connection or disk failures. Logging levels can be controlled for an API Gateway or policy as a whole or on a filter-by-filter basis. Auditing information can be viewed in real time in the API Gateway Manager console, and can be pushed to a database for later analysis using API Gateway Analytics. For more details on logging, see Chapter 7, Troubleshoot your API Gateway installation.
- The API Gateway Manager console is used to provide a current snapshot of how the API Gateway is behaving. For example, it displays how many messages are being processed, and what services are under the most load. API Gateway Manager displays what is happening now on API Gateway instances, and can be viewed by pointing a browser at an Admin Node Manager. For more details, see Chapter 6, Monitoring and reporting.
- Flexible alerts can be sent out to email, SNMP, OPSEC, syslog, Twitter, and Windows Event Log based on a condition being met in a policy. An example might be to email a service owner for every 1000 failures or to generate an alert if a service is processing more than 10000 messages per second. They can also be used to generate alerts on client usage. For more details, see the *API Gateway Policy Developer Guide*.

- Service Level Agreement filters can be used to perform a statistical measure of the services quality of service. They are used to make sure that the amount of network connection errors, response times and server errors are below a certain threshold. For more details, see the *API Gateway Policy Developer Guide*.

Network firewalls

When deployed in a Demilitarized Zone (DMZ) or perimeter network, the API Gateway sits behind network firewalls. Network Address Translation (NAT) firewall functionality is used on the network firewall to provide the API Gateway with a publicly routable address in the DMZ. This allows the API Gateway to route traffic internally to a local IP address range. API Gateways may be dual-homed to pass messages between the DMZ and the internal trusted network.

The API Gateway can then perform security processing on the incoming message traffic. For example, this includes the following:

- Looking for known attack patterns.
- Checking the validity and structure of the message for anomalies.
- Looking for traffic patterns that suggest a Denial-of-Service (DoS) attack.

Advantages over traditional application firewalls

API Gateways have a number of advantages over traditional application firewalls for message processing:

- Can understand the structure of the traffic, and can detect subtle attack mechanisms such as entity expansion attacks and external reference attacks.
- Can consume information in the messages such as security and platform-specific tokens.
- Can use well understood standards such as JSON or XML Schema, JSON Path or XPath, WSDL, and OAuth to properly content filter the traffic.

If an attack or unusual traffic pattern is detected, the API Gateway can notify the firewall to block the traffic at source using OPSEC or some other notification mechanism.

Firewall modes

You can configure the API Gateway in the following modes:

- Block unidentified traffic, which is the default setting. If there is no policy configured for this traffic, block it, and (depending on configuration), raise an alert. In this way, rogue message traffic is detected and blocked.
- Pass unidentified traffic.

You can configure the API Gateway to act as a network endpoint or a network proxy, or both in tandem.

Application servers

Application servers are the infrastructure alongside which API Gateways are most commonly deployed. For example, API Gateways are often deployed in production systems with IBM WebSphere, Oracle WebLogic Server, Microsoft Biztalk Server, Fiorano SOA Platform, TIBCO ActiveMatrix BusinessWorks, and many others.

API Gateways interact with application servers in a number of modes, for example:

- Intercepting application server traffic by acting as a proxy.
- Offloading processing handed over by the application server (for example, to offload message transformation, encryption, or signature validation).

For increased integration with such application server systems, the API Gateway supports numerous transport protocols including HTTP, HTTPS, JMS, email (SMTP/POP), FTP, and so on.

Enterprise Service Buses

API Gateways and Enterprise Service Buses (ESBs) have similar functionality and complement each other very well. Example ESB systems include Oracle Enterprise Service Bus, IBM WebSphere ESB, Progress Sonic ESB, Software AG WebMethods, and JBoss ESB. API Gateways are primarily used to perform the following tasks:

- Protect ESBs and downstream systems from traffic surge, potential DoS attacks, and threats.
- Offload expensive operations such as message validation and cryptography operations from ESBs.

Similarities between API Gateways and ESBs

API Gateways and ESBs typically both perform the following tasks:

- Protocol mediation
- Message routing and transformation
- Service composition
- Message processing

Differences between API Gateways and ESBs

The main differences between API Gateways and ESBs are as follows:

- API Gateways can be used for simple composite services (chained), but do not support Business Process Execution Language (BPEL), and are not suitable for long duration composite services.
- ESBs are usually delivered with backend adapters for systems such as CICS, IMS, Siebel, or SAP.
- API Gateways are stateless and cannot maintain transaction state.
- API Gateways are targeted at performance and application acceleration.
- API Gateways have been designed to provide superior security capabilities, without impacting on performance.

Directories and user stores

API Gateway supports a wide variety of user stores including LDAP, Active Directory, and access control products such as CA SiteMinder and Oracle Access Manager. User stores contain some of the most valuable information in an organization. For example, this includes private identity information such as phone numbers, addresses, email addresses, medical plan IDs, usernames and passwords, certificates, organization structures, and so on. API Gateways must be able to interact with user stores without compromising them.

The API Gateway can use LDAP directories to retrieve user information such as the following:

- Authentication using username/password.
- Retrieve certificates and checking signatures.
- Authorization of clients based on attribute values.
- Retrieval of attributes for placing into SAML assertions.
- Checking certificate validity using Certificate Revocation Lists (CRL) retrieved from user stores.

For more details on integration with LDAP servers, see the *API Gateway Integration Guide*.

Simple inline user store deployment

The following diagram shows a simple inline user store deployment:

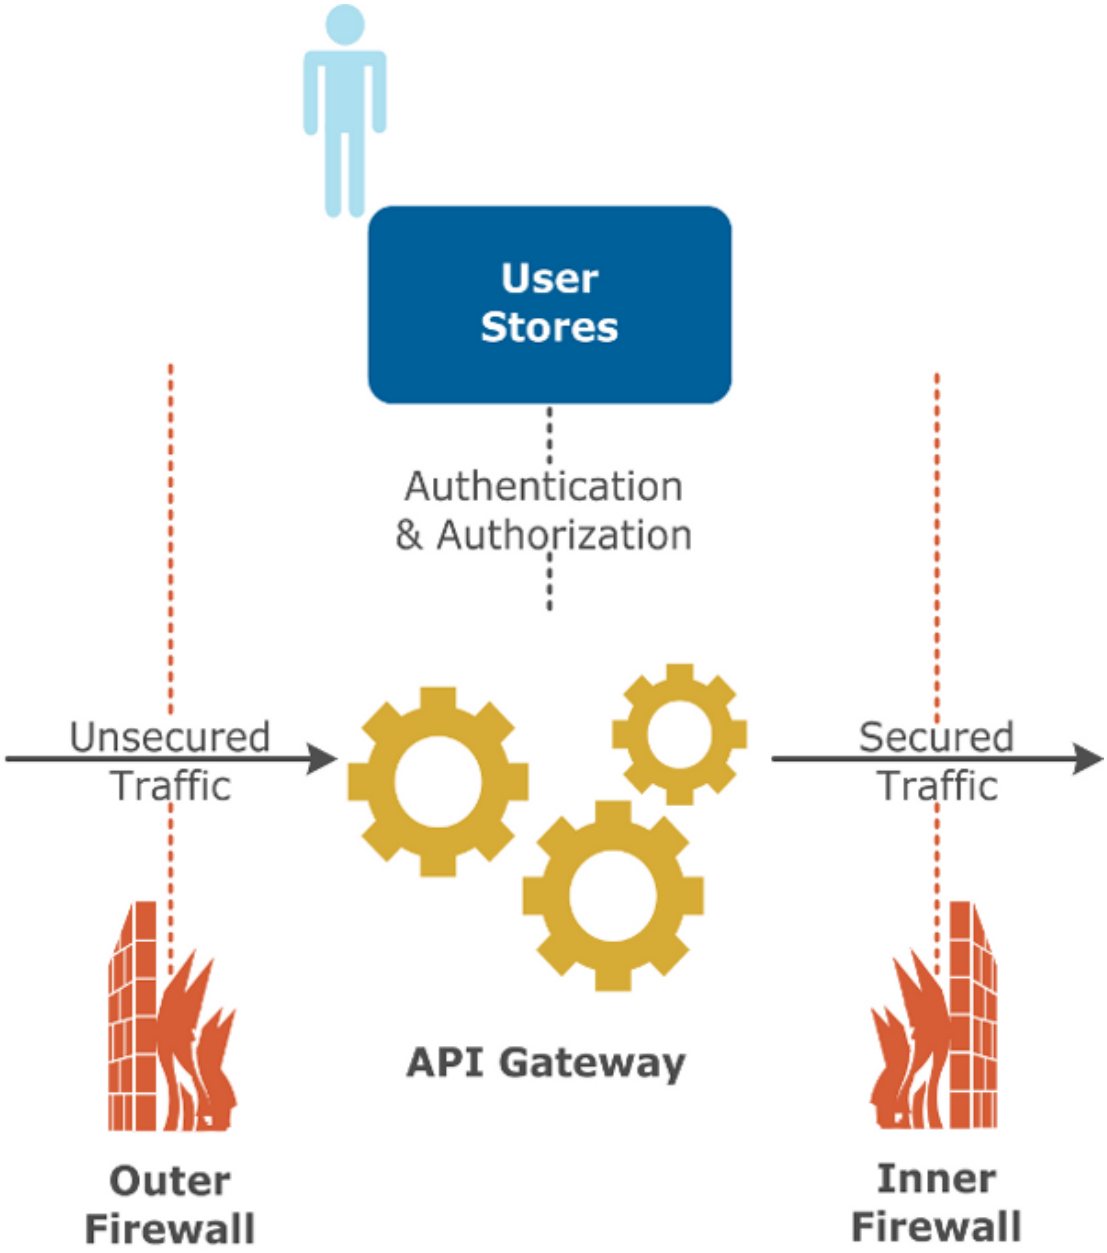

Advantages
This architecture has the following advantages:

- This is simple and easy to set up.
- There is a single entry point through the DMZ for all message traffic.
- This is the least expensive option.

Disadvantages
This architecture has the following disadvantages:

- Exposing important user information in the DMZ is a potential security risk.
- The LDAP server is only being used for external traffic.

API Gateway in DMZ—LDAP in LAN

This is a very common setup where the API Gateway is located in the DMZ, and where it communicates with a user store located in the LAN:

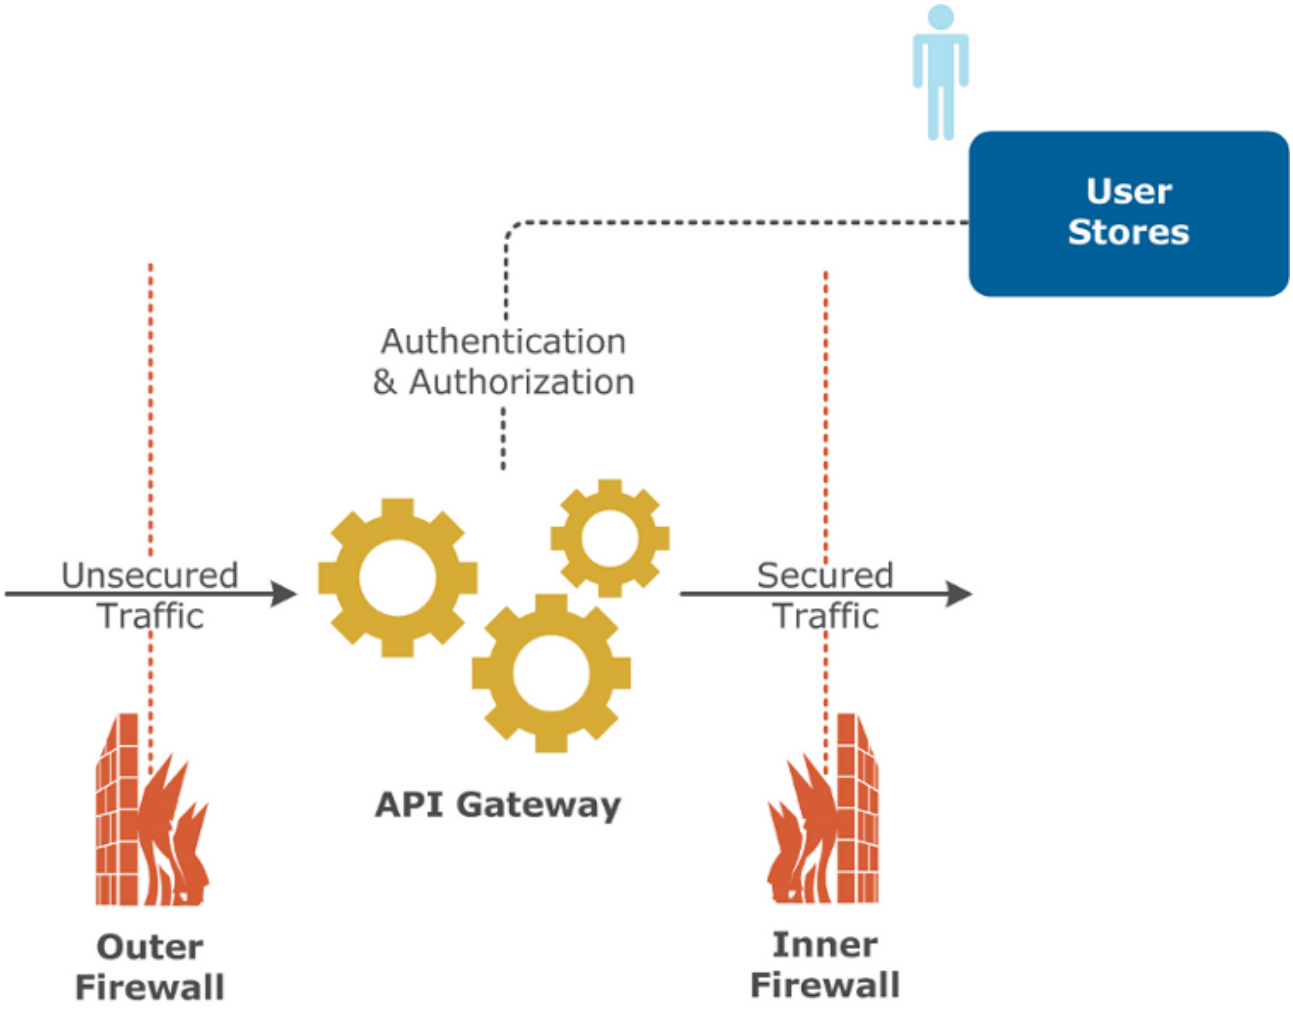

Advantages
This architecture has the following advantages:

- The user store is protected from external access.
- This option is only moderately expensive.

Disadvantages
This architecture has the following disadvantages:

- Administrators must maintain two entry points into the LAN from the DMZ for a single application.

- The user store is addressable from the DMZ, which is contrary to the security policies enforced by many organizations.

Split deployment between DMZ and LAN

In this scenario, two API Gateways are used to split the security checks across the DMZ and the LAN. For example, threats and JSON and XML schema validity are performed in the DMZ, while authentication and/or authorization is performed in the LAN:

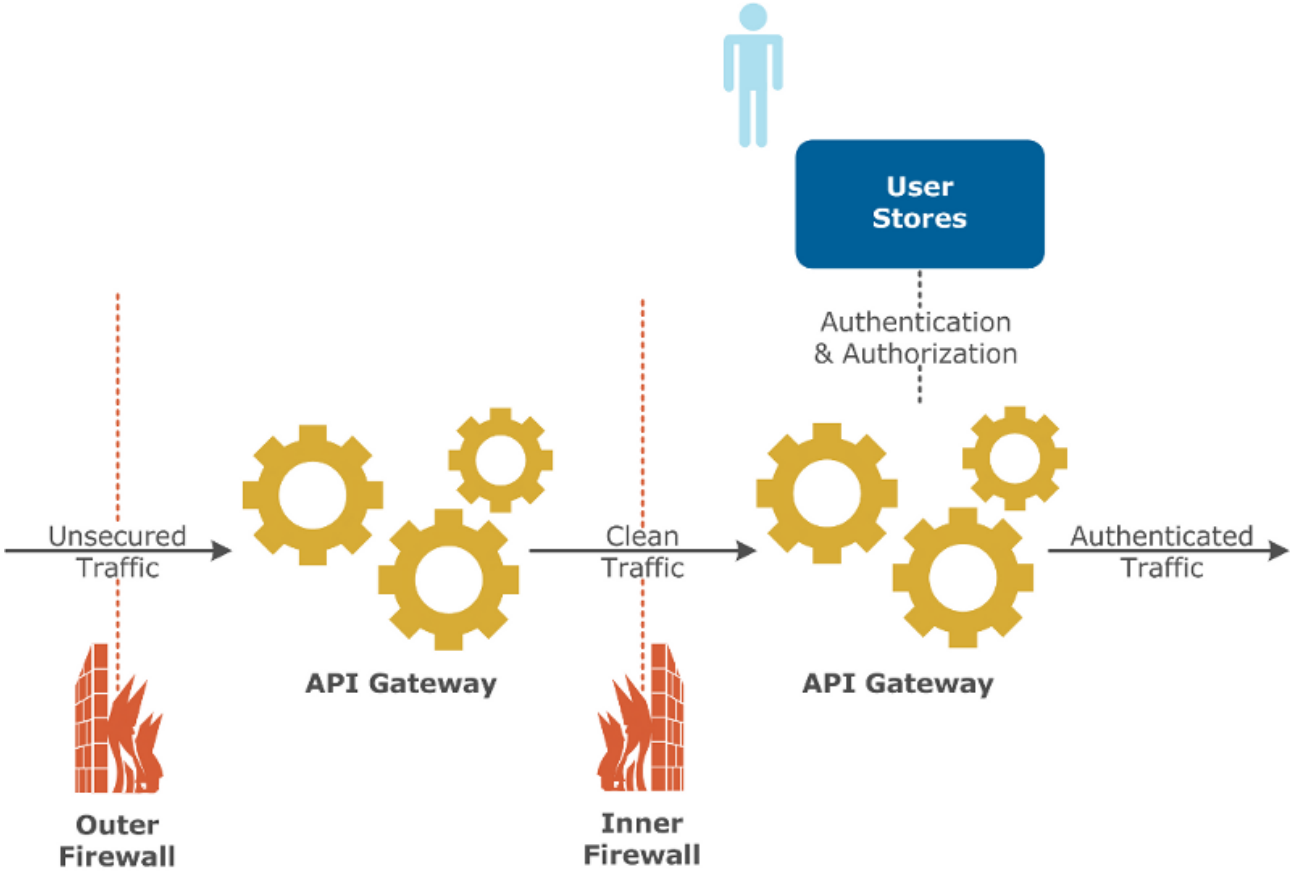

Advantages
This architecture has the following advantage:

- This is the most secure deployment available.
- By separating threats from access control, it is easy to run separate security policies for internal and external traffic.

Disadvantages
This architecture has the following disadvantages:

- This is a more expensive option involving multiple API Gateways.

Access control

API Gateway interoperates with third-party Access Control and Identity Management products at a number of different levels. For example:

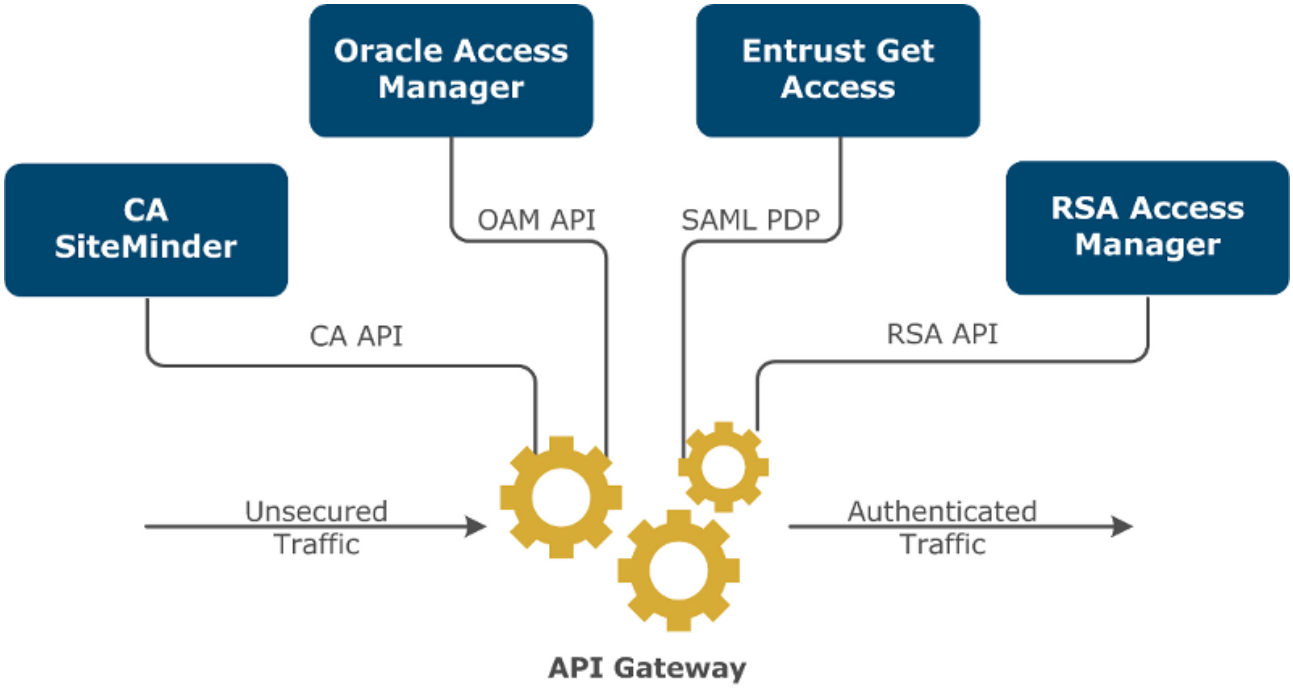

The options shown in the diagram are described as follows:

- API Gateway can directly connect into the Identity Management system as an agent. This solution is currently available for third-party products such as Oracle Access Manager, Oracle Entitlements Server, RSA Access Manager, CA SiteMinder, and IBM Tivoli Access Manager. The Identity Management policy is defined in the Identity Management product to which the API Gateway delegates the authentication and authorization.
- API Gateway can connect to the Identity Manager using the XML Access Control Markup Language (XACML) and Security Assertion Markup Language (SAML) protocols. API Gateway can request an authorization decision from a SAML Policy Decision Point (PDP) for an authenticated client using the SAML Protocol (SAMLP), which is defined in XACML. In such cases, the API Gateway presents evidence to the PDP in the form of user credentials, such as the Distinguished Name of a client X.509 certificate, or a username/password combination.
- The PDP decides whether a user is authorized to access the requested resource. It then creates an authorization assertion, signs it, and returns it to the API Gateway in a SAMLP response. The API Gateway can then perform a number of checks on the response, such as validating the PDP signature and certificate, and examining the assertion. It can also insert the SAML authorization assertion into the message for consumption by a downstream service. This enables propagation of the access control decision to occur.

Public Key Infrastructure

API Gateway supports Secure Sockets Layer (SSL) and Transport Layer Security (TLS) for transport-based authentication, encryption, and integrity checks. The API Gateway can interact with Public Key Infrastructure (PKI) systems in the following ways:

- Connecting to Online Certificate Status Protocol (OCSP) and XML Key Management Specification (XKMS) to query certificate status online.
- Using a Certificate Revocation List (CRL) retrieved from a directory, file, or LDAP to check certificate status.
- Checking a certificate chain.

Certificates and keys can be stored in the API Gateway keystore, or in a network or an optional Hardware Security Module (HSM). For more details, see *Manage certificates and keys*.

Registries and repositories

API Gateway includes the following support for service registries and repositories:

- Architects and policy developers can use Policy Studio to pull Web service definitions (WSDL) from UDDI directories or HTTP-based repositories. These WSDL files are then used to generate the required security policies. The API Gateway can update UDDI registries with updated WSDL files or can serve them directly to the client. For more details, see the *API Gateway Policy Developer Guide*.

Software Configuration Managment

The API Gateway does not include its own built-in Software Configuration Managment or Source Code Management (SCM) features. It is recommended that you use your existing SCM tools to manage your API Gateway configuration packages (`.fed`, `.pol`, and `.env` files) and API Gateway policy or general configuration exports (`.xml` files). For example, commonly used source code management tools include Concurent Version System (CVS), Subversion (SVN), and Git.

In addition, you can use the properties metadata associated with API Gateway configuration packages (`.fed`, `.pol`, and `.env` files) to associate version information with the packages that are provided to upstream enviroments. You can also use the API Gateway Compare/Merge features to determine changes between different configuration packages.

For more details on API Gateway configuration packages and their properties metadata, see the *API Gateway Deployment and Promotion Guide*. For details on exporting API Gateway policies and general configuration, and on Compare/Merge, see the *API Gateway Policy Developer Guide*.

Configure an API Gateway domain

Overview

This topic describes how to use the `managedomain` command in interactive mode to configure a managed API Gateway. It shows how to register a host in a new domain, and create a new API Gateway instance. These are the minimum steps required to configure a domain.

Important

This chapter assumes that you have already installed the API Gateway (see the *API Gateway Installation Guide*). To use the API Gateway, you must have a domain configured in your installation. If you did not already configure a domain when installing the API Gateway, you must configure a domain using `managedomain`.

A single API Gateway installation supports a single API Gateway domain only. If you wish to run API Gateways in different domains on the same host, you need separate installations for each domain. For details API Gateway domains and groups, see the *API Gateway Concepts Guide*.

You can also use the topology view in the web-based API Gateway Manager tool to manage a newly created domain. For example, you can perform tasks such as create or delete API Gateway groups and instances, and start or stop API Gateway instances. For more details, see *Manage domain topology in API Gateway Manager*

Managedomain script

When configuring a domain, the managedomain script enables you to perform tasks such as the following:

- Host management (registering and deleting hosts, or changing Admin Node Manager credentials)
- API Gateway management (creating and deleting API Gateway instances, or adding Windows and Linux/Solaris services)
- Group management (editing or deleting API Gateway groups)
- Topology management (viewing topologies)
- Deployment (deploying to a group, listing deployments, creating or downloading deployment archives, and editing group passphrases)
- Domain SSL certificates (regenerating SSL certificates on localhost)

For example, you can use the `managedomain` script to register a host in a domain and create a new API Gateway instance. These are the minimum tasks required to create a new domain, and which are documented in this topic.

Further information
For details on selecting specific options, enter the `managedomain` command in the following directory, and follow the instructions at the command prompt:

Windows	`INSTALL_DIR\apigateway\Win32\bin`
UNIX/Linux	`INSTALL_DIR/apigateway/posix/bin`

Note

To register an API Gateway instance as a service on Windows or Linux/UNIX, you must run the `man-`

agedomain command as Administrator on Windows or root on Linux/UNIX.

For more details on managedomain options, see *Managedomain command reference*. For details on how to use the managedomain command to configure SSL certificates and Admin Node Manager high availability, see *Configure Admin Node Manager high availability and security*.

Register a host in a domain

To register a host in a managed domain, perform the following steps:

1. Change to the following directory in your API Gateway installation:

Windows	INSTALL_DIR\apigateway\Win32\bin
UNIX/Linux	INSTALL_DIR/apigateway/posix/bin

2. Enter the following command:

```
managedomain --menu
```

3. Enter **1** to register your host, and follow the instructions when prompted. For example, if this is the first host in the domain, enter **y** to configure an Admin Node Manager on the host. Alternatively, to add the host to an existing domain, enter **n** to configure a local Node Manager that connects to the Admin Node Manager in the existing domain.
4. Enter **q** to quit when finished.
5. Enter the following command to start the Admin Node Manager or local Node Manager on the registered host:

```
nodemanager
```

Important

Before registering multiple hosts in a domain, you must first ensure that a licensed API Gateway is installed on each host machine. Then to register each host, you must select option 1 on each host machine.

You must ensure the Admin Node Manager is running in the domain to enable monitoring and management of API Gateway instances.

Create an API Gateway instance

To create an API Gateway instance, perform the following steps:

1. Open a new command window.
2. Change to the following directory in your API Gateway installation:

Windows	INSTALL_DIR\apigateway\Win32\bin
UNIX/Linux	INSTALL_DIR/apigateway/posix/bin

3. Enter the following command:

```
managedomain --menu
```

4. Enter **5** to create a new API Gateway instance, and follow the instructions when prompted. You can repeat to create multiple API Gateway instances on local or remote hosts.
5. Enter **q** to quit when finished.
6. Use the `startinstance` command to start the API Gateway, for example:

```
startinstance -n "my_server" -g "my_group"
```

Note

You can add an API Gateway instance on any registered host in the domain, not just the local host. However, if you are creating Windows or UNIX services for the API Gateway, you must run `managedomain` on same host.

You must run `startinstance` on the host on which you intend to start the instance. On UNIX/Linux, you must ensure that the `startinstance` file has execute permissions. Running `startinstance` without any arguments lists all API Gateway instances available on the host.

Test the health of an API Gateway instance

You can test the connection to the new API Gateway instance by connecting to the Health Check service. For example, enter the following default URL in your browser:

```
http://HOST:8080/healthcheck
```

This should display a simple `<status>ok</status>` message.

You can view the newly created API Gateway instance on the API Gateway Manager dashboard. For example, the default URL is as follows:

```
https://HOST:8090
```

The port numbers used to connect depend on those entered when configuring the domain using `managedomain`, and are available from the localhost only.

Alternatively, you can also connect to the new API Gateway instance in Policy Studio. For more details, see *Start the API Gateway tools*.

Manage domain topology in API Gateway Manager

Overview

This topic describes how to use the topology view in the web-based API Gateway Manager tool to manage an existing API Gateway domain. For example, you can perform tasks such as the following:

- Create and delete API Gateway groups
- Lock and unlock API Gateway groups
- Create and delete API Gateway instances
- Start and stop API Gateway instances
- Create and edit API Gateway configuration tags
- Deploy API Gateway configuration

Note

When using API Gateway Manager to manage an existing domain, you must ensure that the host was first registered in the domain using the `managedomain` script. For more details, see *Configure an API Gateway domain*.

The API Gateway Manager web console is available from the following URL:

```
https://HOST:8090/
```

For more details, see *Start the API Gateway tools*.

Manage API Gateway groups

You can use the topology view in API Gateway Manager tool to create, delete, and lock API Gateway groups. The following example shows the options available at the group level:

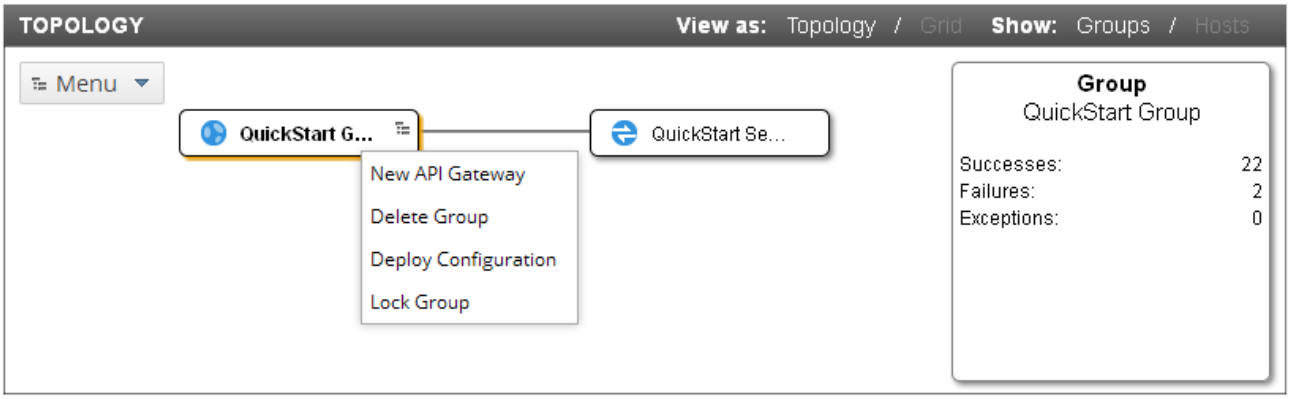

Create an API Gateway group

To use the API Gateway Manager to create an API Gateway group, perform the following steps:

1. Click the **Menu** button in the topology view on the **Dashboard** tab.
2. Select **Create New Group**.

3. Enter a group name (for example, `Engineering`).
4. Click **OK**.

Delete an API Gateway group

To delete a group, perform the following steps:

1. Ensure that the API Gateway instances in the group have been stopped.
2. Hover over the group in the topology view, and click the edit button on the right.
3. Select **Delete Group**.
4. Click **OK**.

Lock an API Gateway group

To lock a group and prevent other users from editing its configuration, perform the following steps:

1. Hover over the group in the topology view, and click the edit button on the right.
2. Select **Lock Group**.
3. Click **OK**.

Similarly to unlock an API Gateway group, select **Unlock Group**. Only admin users can unlock groups locked by other users.

Tip

You can also use Policy Studio to delete or lock an API Gateway group. In the Policy Studio topology view, right-click the API Gateway group, and select **Remove Group** or **Lock Group Access**.

Manage API Gateway instances

You can use the topology view in API Gateway Manager to create, delete, start, and stop API Gateway instances.

Create API Gateway instances

To use the API Gateway Manager to create an API Gateway instance, perform the following steps:

1. Hover over the API Gateway group in the topology view, and click the edit button on the right.
2. Select **New API Gateway**.
3. Configure the following fields:
   - **Name**: API Gateway instance name (for example, `Server2`).
   - **Management Port**: Local management port (for example, `8086`).
   - **Services Port**: External traffic port (for example, `8081`).
   - **Host**: Host address (for example, `127.0.0.1`).
   - **Domain CA Passphrase**: The passphrase for the domain SSL certificate authority.
4. Click **OK**.

Delete API Gateway instances

To delete an API Gateway instance, perform the following steps:

1. Ensure that the API Gateway instance has been stopped.
2. Hover over the API Gateway instance in the topology view, and click the edit button on the right.

3. Select **Delete Server**.
4. Click **OK**.

Tip

You can also use Policy Studio to create or delete an API Gateway instance. In the Policy Studio topology view, right-click the API Gateway group, and select **Create API Gateway**. Or right-click an existing API Gateway instance, and select **Remove API Gateway**.

Start API Gateway instances

To start an API Gateway instance, perform the following steps:

1. Ensure that the API Gateway instance has been stopped.
2. Hover over the API Gateway instance in the topology view, and click the edit button on the right.
3. Select **Start**.
4. Click **OK**.

The following example shows how to start a stopped API Gateway instance:

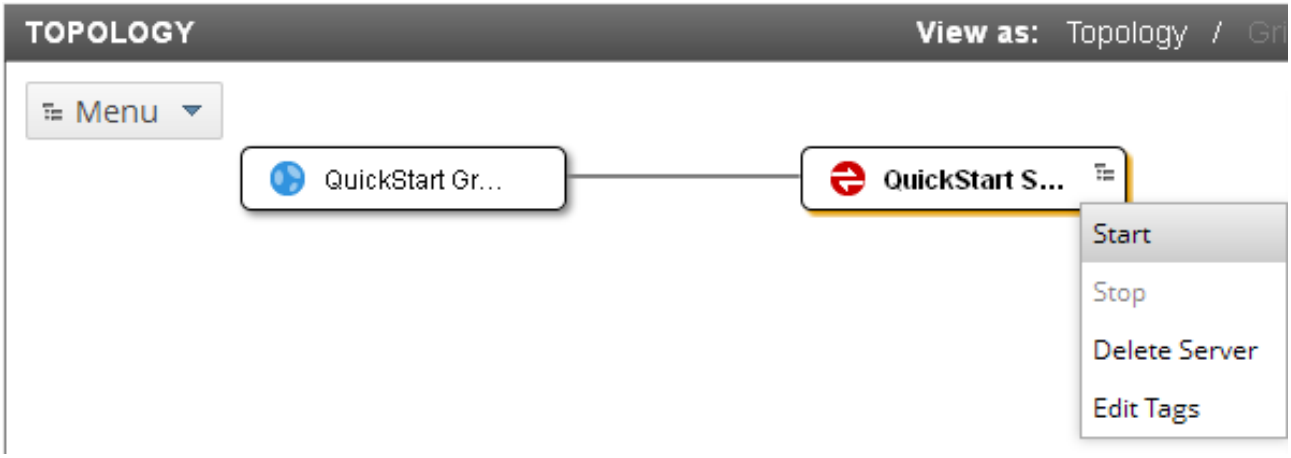

Stop API Gateway instances

To stop an API Gateway instance, perform the following steps:

1. Ensure that the API Gateway instance has been started.
2. Hover over the API Gateway instance in the topology view, and click the edit button on the right.
3. Select **Stop**.
4. Click **OK**.

Edit API Gateway tags

API Gateway tags are name-value pairs that you can add to API Gateway instances, which then enable you to filter for API Gateway instances by tag in the API Gateway Manager **Dashboard**. To add an API Gateway tag, perform the following steps:

1. Hover over the API Gateway instance in the topology view, and click the edit button on the right.

2. In the dialog, click the green plus icon to add a tag.
3. Enter a **Tag** name (for example, `Department`), and a **Value** (for example, `Engineering`).
4. Click **Apply**.

To view the newly created tag in the in the API Gateway Manager topology view, click **View as** > **Grid**. To filter this view, enter a tag name or value in the **Filter** field.

Deploy API Gateway configuration

You can use the API Gateway Manager web console to deploy API Gateway configuration packages to a group of API Gateway instances. This includes the following configuration packages:

- deployment package (`.fed`)
- policy package (`.pol`)
- environment package (`.env`)

For more details on configuration packages, see the *API Gateway Promotion and Deployment Guide*.

Deploy a deployment package

To deploy an existing deployment package to a group of API Gateways, perform the following steps:

1. In the **TOPOLOGY** view, right-click the API Gateway group to which to deploy the package, and select **Deploy Configuration**.
2. Select **I wish to deploy configuration contained in a single Deployment Package**.
3. Click **Browse for .fed**, and select the `.fed` file in the dialog.
4. Click **Next**.
5. Select **Deploy** in the wizard, and the deployment package is deployed to the API Gateway group.
6. Click **Finish**.

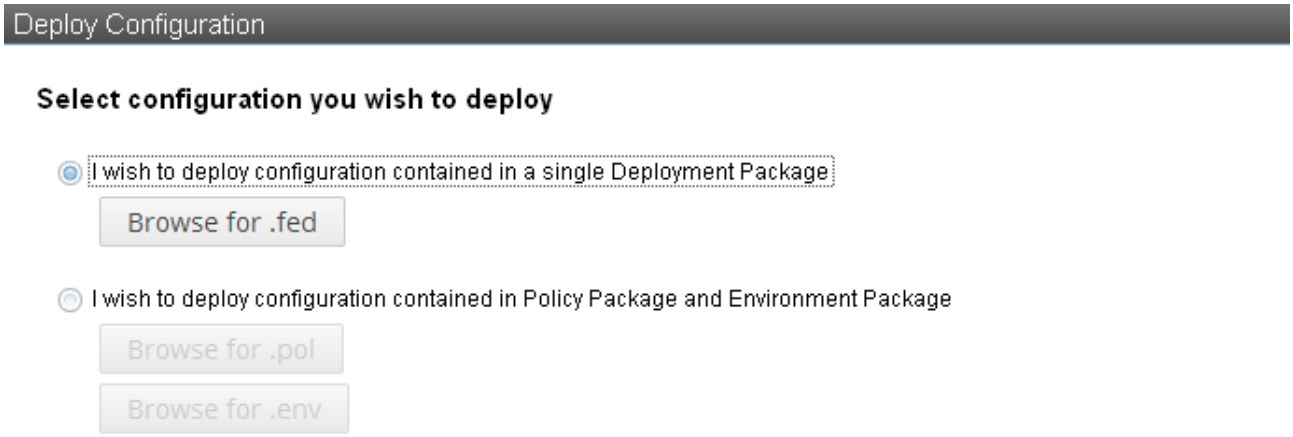

Deploy policy and environment packages

To deploy existing policy and environment packages to a group of API Gateways, perform the following steps:

1. In the **TOPOLOGY** view, right-click the API Gateway group to which to deploy the packages, and select **Deploy**

Configuration.

2. Select **I wish to deploy configuration contained in Policy Package and Environment Package**.
3. Click **Browse for .pol**, and select the `.pol` file in the dialog.
4. Click **Browse for .env**, and select the `.env` file in the dialog.
5. Click **Next**.
6. Select **Deploy** in the wizard, and the packages are deployed to the API Gateway group.
7. Click **Finish**.

Tip

You can also use Policy Studio to deploy API Gateway configuration. For details, see *Deploy API Gateway configuration*.

Configure Admin Node Manager high availability and security

Overview

The Admin Node Manager (ANM) is the central administration server for an API Gateway domain, and is responsible for performing management operations across the domain. The Node Manager (NM) on each machine manages all the local API Gateways on that machine, regardless of the group they are in. This includes collecting monitoring information, managing dynamic settings, deploying configuration, and so on.

In addition to managing the local API Gateways on its host, the Admin Node Manager communicates with the Node Managers in the domain to perform management operations across the domain. In this architecture, the Node Managers only communicate with the Admin Node Manager, and the API Gateway Manager, Policy Studio, and `managedomain` tools connect to the Admin Node Manager. For more details on Admin Node Manager architecture, see the *API Gateway Concepts Guide*.

Important

It is recommended that you configure at least two Admin Node Managers in an API Gateway domain for high availability. This topic describes how to use the using the `managedomain` command to configure and secure multiple Admin Node Managers in a domain.

Hierarchy of SSL certificates in a domain

API Gateway uses Secure Sockets Layer (SSL) for communications between all processes in a domain. This includes internal management traffic between the Admin Node Manager, Node Manager, and API Gateway instances, as shown in the diagram.

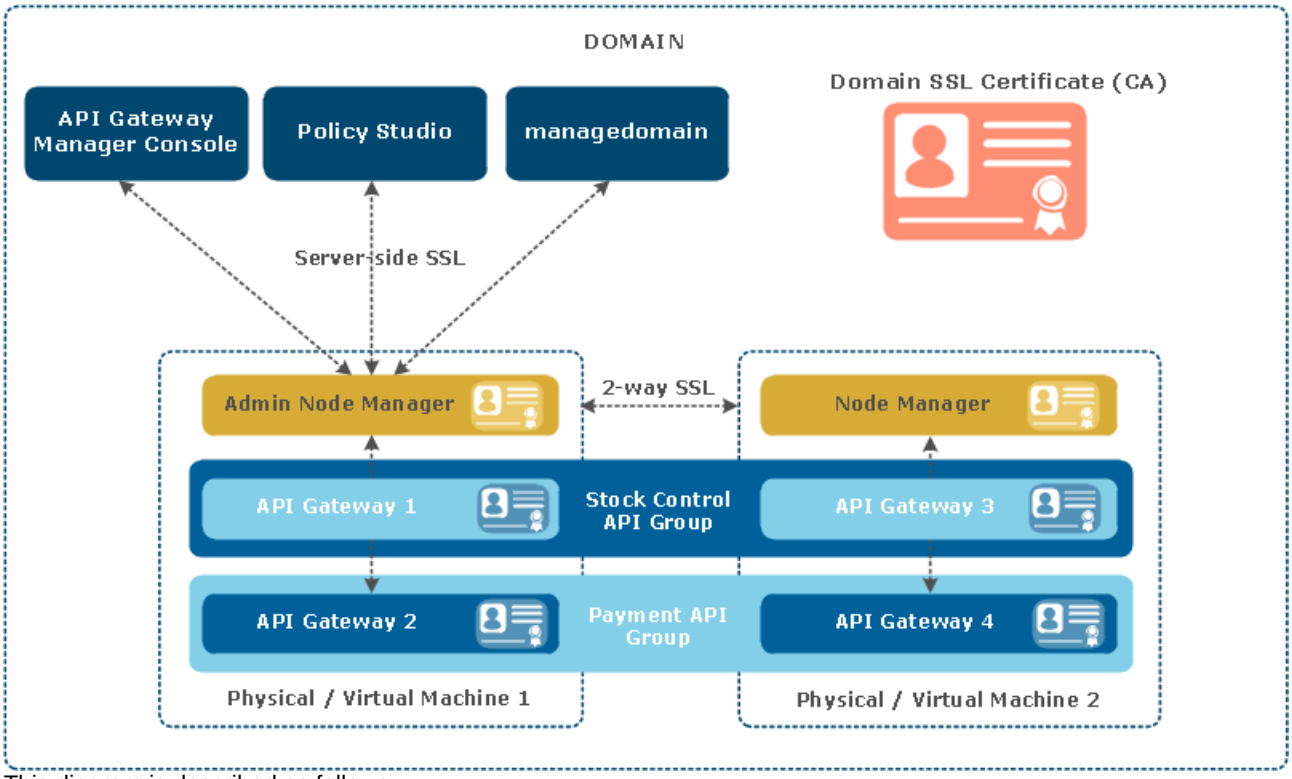

This diagram is described as follows:

- Server-side SSL authentication is used for communication between the Admin Node Manager and client applications (Policy Studio, API Gateway Manager, `managedomain`).
- Mutual SSL is used for communication between Admin Node Managers and Node Managers, and also between Node Managers and API Gateway instances.
- Each Admin Node Manager, Node Manager, and API Gateway process has an associated SSL certificate used in the SSL session for inter-process communication. These certificates are also used to determine whether a process is a Node Manager or an Admin Node Manager.
- The domain SSL certificate is the Certificate Authority (CA) used to sign the SSL certificates generated for each Admin Node Manager, Node Manager, and API Gateway process.

Note

The SSL inter-process communication discussed in this topic refers to system-level management communication only. This is not related to any SSL communication with back-end services, or any SSL ports that the API Gateway exposes for business traffic.

How SSL certificates are generated for domain processes

Each time you register a new Admin Node Manager or Node Manager, or create a new API Gateway instance, a new SSL certificate is generated for inter-process communication. Each SSL certificate associated with an Admin Node Manager, Node Manager, or API Gateway instance is signed by the same CA. You must use one of the following options to sign these SSL certificates:

Signing option	Description
System-generated CA key	By default, a system-generated CA key and certificate is used to sign all SSL certificates.

Signing option	Description
and certificate	The system creates a private key and self-signed certificate on the first Admin Node Manager to be registered. All subsequent requests for new SSL certificates for inter-process communication are processed by this Admin Node Manager which has the CA key.
User-provided CA key and certificate	Provide your own CA key by specifying a P12 file that contains a key and certificate. The system uses this CA key to sign all new SSL certificates for inter-process communication.
External CA	A private key and Certificate Signing Request (CSR) file are generated by the system, and then you can take this CSR to your chosen external CA. When you receive the certificate from the external CA, you must submit the certificate to `managedomain`, which then completes registering the new Admin Node Manager, Node Manager, or API Gateway instance associated with that certificate.

These options are available when using `managedomain` to register a new Admin Node Manager or Node Manager, or to create a new API Gateway instance, and when regenerating all SSL certificates for the domain. These options are also available when editing a host in `managedomain` (for example, when you change a Node Manager to an Admin Node Manager a new certificate is generated).

Note

Private keys for the SSL certificates are always generated and stored on the host that uses them. This applies to a system-generated CA key, a key for an Admin Node Manager or Node Manager SSL certificate, or a key for an API Gateway SSL certificate for inter-process communication.

External Certificate Authority

You can use a third-party external CA to sign the SSL certificates (for example, Verisign). This is the typical approach that is recommended in a production environment. For more details, see your external CA user documentation.

In a development environment, you can install API Gateway on a locked down host to act as an external CA for testing purposes. No license is required because only `managedomain` is running. This external CA host stores a system-generated CA private key and certificate. You can use `managedomain` to process CSRs and output certificates.

In this scenario, the CSR file is generated by an API Gateway installation on another host (with a Node Manager and API Gateways running), and transported to the external CA installation out-of-band. The certificate file generated by the external CA is transported back out-of-band to the installation that created the CSR and where the new Node Manager or API Gateway is created. The certificate is specified to `managedomain` using the submit certificate option to complete Node Manager host registration or API Gateway creation.

Add the first Admin Node Manager to the domain

Important

When registering the first Admin Node Manager, you must run `managedomain -i` on the host that will run the Admin Node Manager. This generates the private key for the Admin Node Manager, and private keys must always be generated where they are used. No SSL certificates are provided out-of-the-box .

Sign Admin Node Manager certificate with system-generated CA key

To register the first Admin Node Manager in the domain and sign its SSL certificate with a system-generated CA key, run the following command:

```
managedomain -i
```

This command generates the CA private key, and a self-signed certificate. It also generates the Admin Node Manager private key and certificate signed by self-signed CA certificate. This command uses the `--sign_with_generated` option by default, which means that the SSL certificates are signed with a system-generated CA private key and certificate. There are no Node Managers running in the domain to call yet.

Sign Admin Node Manager certificate with user-provided CA key

To register the first Admin Node Manager in the domain and sign its SSL certificate with a user-provided CA key, run the following example command:

```
managedomain -i --sign_with_user_provided --ca=/home/keys/test.p12
```

This command generates the Admin Node Manager private key and certificate signed by user-provided CA certificate.

Sign Admin Node Manager certificate with external CA

To register the first Admin Node Manager in the domain and sign its SSL certificate with an external CA, run the following example command:

```
managedomain -i --sign_with_external_ca
```

This command generates the Admin Node Manager private key and the CSR for the certificate. It then prompts you to take your CSR file to your external CA and return with the signed certificate.

When you have the signed certificate, use the following command to submit the certificate and finish creating the Admin Node Manager configuration:

```
managedomain --submit_cert --cert cert.pem
```

API Gateway as external CA

In a production environment, you can use the `--sign_with_external_ca` option, and take the generated CSR to an external CA (for example, Verisign). Alternatively, in a development environment, you can run the following command to generate the certificate on an isolated API Gateway installation:

```
managedomain --sign_csr --csr nodemanager.csr
```

This command creates a certificate signed using the system-generated CA key. It writes the certificate to a file, displays the file name so you can retrieve and copy to the host running `managedomain`, and complete registration of the new Node Manager.

Tip

If you wish to avoid using a system-generated self-signed certificate, you can manually copy a user-provided CA private key and certificate to the following locations:

```
apigateway/groups/certs/private/domain.p12
apigateway/groups/certs/private/domainkey.pem
apigateway/groups/certs/domaincert.pem
```

Additional certificate generation options

In addition to the example commands already shown, you can specify the following options when using `managedomain` to sign SSL certificates for inter-process communication:

Option	Description
`--domain_passphrase`	Encrypt a newly generated `domain.p12`, or unlock a user-provided `domain.p12`.
`--domain_name`	Specify the Distinguished Name (DName) for a system-generated CA certificate. Defaults to `CN=Domain`.
`--sign_alg`	Specify the certificate signing algorithm (for example, `sha1`, `sha256`, or `sha512`). Defaults to `sha1`.
`--subj_alt_name`	Specify a subject alternative name content for the Admin Node Manager certificate. This is used by web browsers during the SSL handshake exchange. You can specify this multiple times, for example: `--subj_alt_name "DNS.1=name.oracle.com" --subj_alt_name "DNS.2=othername.oracle.com" --subj_alt_name "IP.1=10.142.1.2"`
`--key_passphrase`	You can also encrypt temporary private key files stored on disk. For example, for the Admin Node Manager private key before it is consumed and written to the Node Manager configuration.

Add a Node Manager to the domain

The steps for adding a Node Manager or Admin Node Manager are different because there is now an Admin Node Manager running in the domain that must be invoked to add the new Node Manager.

Sign Node Manager certificate with system-generated CA key

To add a Node Manager to the domain and sign its SSL certificate with a system-generated CA key, run the following command:

```
managedomain -a --sign_with_generated --anm_host my_anm.com
```

This command generates the Node Manager private key, obtains the Node Manager certificate from the Admin Node Manager, and completes registration of the new Node Manager in the domain.

Sign Node Manager certificate with user-provided CA key

To add a Node Manager to the domain and sign its SSL certificate with a user-provided CA key, run the following command:

```
managedomain -a --sign_with_user_provided --ca=/home/keys/test.p12 --anm_host my_anm.com
```

This command signs the new Node Manager certificate with the user-provided CA key and certificate, and completes registration of the new Node Manager in the domain.

Sign Node Manager certificate with external CA

To add a Node Manager to the domain and sign its SSL certificate with an external CA, run the following command:

```
managedomain -a --sign_with_external_ca --anm_host my_anm.com
```

This command performs similar steps to adding the first Admin Node Manager to the domain and signing with an external CA.

When you have the signed certificate, use the following command to submit the certificate and finish creating the Node Manager configuration:

```
managedomain --submit_cert --cert cert.pem
```

See also the section called “API Gateway as external CA”.

Additional options

You can use the `--is_admin` option to specify whether the process is an Admin Node Manager or Node Manager. For more details, see the section called “Change a Node Manager to an Admin Node Manager”. See also the section called “Additional certificate generation options”.

Add an API Gateway instance to the domain

This section describes the options for signing an API Gateway instance SSL certificate when adding an API Gateway instance to the domain.

Sign API Gateway instance certificate with system-generated CA key

To add an API Gateway instance to the domain and sign its SSL certificate with a system-generated CA key, run the following example command:

```
managedomain -c -n APIGateway1 -g Group1
```

This command uses the `--sign_with_generated` option by default.

Sign API Gateway instance certificate with user-provided CA key

To add an API Gateway instance to the domain and sign its SSL certificate with a user-provided CA key, run the following example command:

```
managedomain -c -n APIGateway1 -g Group1 --sign_with_user_provided --ca=/home/keys/test.p12
```

Sign API Gateway instance certificate with external CA

To add an API Gateway instance to the domain and sign its SSL certificate with an external CA, run the following example command:

```
managedomain -c -n APIGateway1 -g Group1 --sign_with_external_ca
```

This command performs similar steps to adding the first Admin Node Manager to the domain and signing with an external CA.

When you have the signed certificate, use the following command to submit the certificate and finish creating the API Gateway configuration:

```
managedomain --submit_cert --cert cert.pem
```

See also the section called “API Gateway as external CA” and the section called “Additional certificate generation options”.

Change a Node Manager to an Admin Node Manager

Tip

When you add a Node Manager to the domain using the `-a` option, you can add it as a Node Manager or Admin Node Manager. For more details, see the section called “Add a Node Manager to the domain”. The section that follows describes how to change whether it is a Node Manager or Admin Node Manager at a later stage.

An existing Node Manager in a domain can only communicate with API Gateway instances on the same host machine, and only an Admin Node Manager can connect to a Node Manager. This means that if you have a single Admin Node Manager in a domain, this is a single point of failure for management capabilities. If the single Admin Node Manager fails, you cannot manage or monitor the domain or deploy configuration. However, you can change an existing Node Manager to act as an Admin Node Manager in the domain. A domain can run with any number of Admin Node Managers, but it is recommended that you have at least two Admin Node Managers in a domain.

You can change the admin capabilities of the Node Manager running on the host local to `managedomain` using the `--edit_host` option. This enables you to change a Node Manager to an Admin Node Manager, or change an Admin Node Manager to a Node Manager. Changing admin capabilities means that the tags on the process in the topology are updated, and a new SSL certificate is generated.

The same options for signing an SSL certificate for a Node Manager or API Gateway instance (described in the preceding sections) are available when changing a Node Manager to an Admin Node Manager, or an Admin Node Manager to a Node Manager.

Note

To change admin capabilities of an Admin Node Manager or Node Manager, you must run `managedomain` on the same host that the Admin Node Manager or Node Manager you wish to change is running. You can edit other aspects of a host remote from `managedomain`.

You cannot remove admin capabilities when there is only one Admin Node Manager in a domain because there must be at least one Admin Node Manager in a domain.

Sign Node Manager certificate with system-generated CA key

When `host1.com` has an Admin Node Manager, to change admin capabilities from Admin Node Manager to Node Manager, run the following example command:

```
managedomain --edit_host --host host1.com
```

In this case, because `host1.com` is an Admin Node Manager and the `--is_admin` parameter is not passed, `managedomain` interprets this command as a change of capabilities from Admin Node Manager to Node Manager. This command uses the `--sign_with_generated` option by default.

When `host1.com` has a Node Manager, to change admin capabilities from Node Manager to Admin Node Manager, run the following example command:

```
managedomain --edit_host --host host1.com --is_admin
```

In both cases, the command generates and updates the private key and certificate in the local Node Manager configuration, and prompts you to reboot the Node Manager.

Important

To remove admin capabilities from the first registered Admin Node Manager, you must first move the following directories to another Admin Node Manager:

```
apigateway/groups/certs
apigateway/groups/certs/private
```

Otherwise, you will not be able to remove admin capabilities. This only applies to a system-generated CA key.

Sign Node Manager certificate with user-provided CA key

When `host1.com` has an Admin Node Manager, to change admin capabilities from Admin Node Manager to Node Manager, run the following example command:

```
managedomain --edit_host --host 10.142.12.31 --sign_with_user_provided
  --ca=/home/keys/test.p12
```

When `host1.com` has a Node Manager, to change admin capabilities from Node Manager to Admin Node Manager, run the following example command:

```
managedomain --edit_host --host host1.com --sign_with_user_provided --ca=/home/keys/test.p12
  --is_admin
```

In both cases, the command performs the same steps described in the section called “Sign Node Manager certificate with system-generated CA key”, and the same conditions apply.

Sign Node Manager certificate with external CA

When `host1.com` has an Admin Node Manager, to change admin capabilities from Admin Node Manager to Node Manager, run the following example command:

```
managedomain --edit_host --host host1.com --sign_with_external_ca
```

When `host1.com` has a Node Manager, to change admin capabilities from Node Manager to Admin Node Manager, run the following example command:

```
managedomain --edit_host --host host1.com --sign_with_external_ca --is_admin
```

In both cases, this command performs similar steps to adding the first Admin Node Manager to the domain and signing with an external CA.

When you have the signed certificate, use the following command to submit the certificate and finish updating the Node Manager configuration:

```
managedomain --submit_cert --cert cert.pem
```

See also the section called “Additional certificate generation options”.

Regenerate all SSL certificates in a domain

This section explains how to regenerate all SSL certificates used for inter-process communication in the domain. You can do this using a system-generated self-signed CA key and certificate, a user-provided CA key and certificate in a P12 file, or an external CA.

You can also change the domain certificate management option. For example, if the domain was first configured to use a system-generated CA key and certificate to sign all SSL certificates, you can change to a user-provided CA key and cer-

tificate, or an external CA. However, you must use the same option on all hosts in the domain.

Sign certificates in domain with system-generated CA key

Important

You must first regenerate on the Admin Node Manager host that will hold the system-generated self-signed CA key. This may or may not be the same host that held the system-generated self-signed CA key before regeneration.

If you select a different host to hold the CA private key and certificate, you must delete the old CA private key and certificate from the other Admin Node Manager host to prevent certificates being signed by two different CA keys. This must be done before regeneration.

To regenerate all certificates on a host and sign with a system-generated key, run the following command on each host in the domain:

```
managedomain --regen_certs
```

This command uses the `--sign_with_generated` option by default. It regenerates the CA private key and self-signed certificate, and private keys and certificates for the Admin Node Manager, Node Manager, and API Gateway instances on the host that the command was run.

After regenerating certificates, you must reboot the Node Manager and API Gateway instances on the local machine.

Tip

When you regenerate certificates on the first Admin Node Manager, a Node Manager or API Gateway does not need to be running. When you regenerate certificates on subsequent hosts, the Admin Node Manager holding the CA key must be running because it is used to sign the certificates. This applies only for system-generated CA keys.

Sign certificates in domain with user-provided CA key

To regenerate all certificates on a host and sign with a user-provided key, run the following example command on each host in the domain:

```
managedomain --regen_certs --sign_with_user_provided --ca=/home/users/qa.p12
   --domain_passphrase foo
```

This command generates the private keys for the Node Manager and API Gateway instances and signs their certificates. After regenerating the certificates, you must then reboot the Node Manager and API Gateway instances on the local machine.

In this case, `managedomain` does not communicate with an Admin Node Manager when regenerating certificates. The order of hosts selected for certificate regeneration does not matter. An Admin Node Manager does not need to be running when certificates are regenerated or submitted.

Sign certificates in domain with external CA

To regenerate all certificates on a host and sign with an external CA, run the following command on each host in the domain:

```
managedomain --regen_certs --sign_with_external_ca
```

This command generates the private keys for the Node Manager and API Gateway instances. You are prompted to take the CSR files to your CA and return with the signed certificates.

When you have the signed certificates, run the following command for each certificate:

```
managedomain --submit_cert --cert cert.pem
```

After regenerating certificates, you must reboot the Node Manager and API Gateway instances on the local machine. In this case, `managedomain` does not talk to an Admin Node Manager when regenerating certificates. The order of hosts selected for certificate regeneration does not matter. An Admin Node Manager does not need to be running when certificates are regenerated or submitted.

Tip

You can use `managedomain` in interactive mode (option `24`) to regenerate a subset of certificates on the host. If you regenerate the CA key and certificate, the Node Manager and all API Gateway certificates are automatically generated. If you do not regenerate the CA key and certificate, you can choose whether to regenerate the Node Manager certificate, and which API Gateway certificates to regenerate. For more details, see *Managedomain command reference*.

Reset passphrase for CA private key

If a system-generated CA private key is used, you can reset the passphrase on the command line as follows:

```
managedomain --change_domain_passphrase --old_passphrase "oldpass" --new_passphrase "newpass"
```

You can only run this command on the host that has the `domain.p12` file. It does not affect any of the SSL certificates that have already been generated. If a new SSL certificate is to be generated (due to a new Node Manager or new API Gateway instance), the new domain passphrase must be provided to unlock the CA private key. This command modifies local files only. It does not communicate with any Node Managers in the domain.

Change domain SSL certificate expiry date

By default, the SSL certificates are valid for 100 years. To change this value, edit the `apigateway/skel/openssl.cnf` file, and change the value of the `default_days` setting in the `CA_default` section.

```
[CA_default]
 .
 .
default_days = 36500
 .
 .
```

This setting applies to signing with a system-generated or user-generated CA key, or if you are using an API Gateway installation as an external CA. If you are using a third-party external CA, the CA decides on the expiry date.

When the `default_days` setting is changed, all newly signed certificates have the new expiry date. You do not need to reboot the Node Manager. However, you must regenerate certificates if you wish to modify the expiry date of existing certificates.

For more details, see the section called “Location of SSL private keys and certificates”.

Admin Node Manager backup and disaster recovery

This section explains how to create a backup Admin Node Manager for signing certificates, and how to set up an Admin Node Manager for signing certificates from a backup `domain.p12` file.

Create backup Admin Node Manager for signing certificates

If you are using a system-generated CA private key and certificate, the CA private key and certificate are only on the first Admin Node Manager registered in the domain. You can manually copy this CA private key and certificate to other Admin Node Managers so that there is always a backup Admin Node Manager available to sign certificates. Alternatively, you can copy the CA private key and certificate only when there is a problem with the first Admin Node Manager (for example, it is no longer running).

To create a backup Admin Node Manager that can sign certificates, copy the following files from Admin Node Manager 1 to Admin Node Manager 2:

```
apigateway/groups/certs/domaincert.pem
apigateway/groups/certs/index.txt
apigateway/groups/certs/serial
apigateway/groups/certs/private/domain.p12
apigateway/groups/certs/private/domainkey.pem
```

You do not need to reboot Admin Node Manager 2 to be able to sign certificates. You can test the Admin Node Manager's ability to sign certificates by taking down Admin Node Manager 1. Connect the client to Admin Node Manager 2, and add an API Gateway instance on any Node Manager that is running.

Important

You can decide when and if to copy the CA private key to limit the copies of this sensitive data. You must always backup the `domain.p12` file after the registration of the first Admin Node Manager in some secure offline location.

Set up Admin Node Manager for signing certificates from a backup .p12

You can set up an Admin Node Manager for signing certificates from a backup `domain.p12` file only. Perform the following steps:

1. Copy the `domain.p12` file to `apigateway/groups/certs/private/`.
2. Create an `apigateway/groups/certs/serial` file with contents of `01`.
3. Create an empty `apigateway/groups/certs/index.txt` file.
4. Parse `domain.p12` into `domainkey.pem` and `domaincert.pem` (if these files were not backed up originally). For example, assuming a passphrase of `fred` for the CA private key:

```
cd apigateway/groups/certs/private
openssl pkcs12 -in domain.p12 -out ../domaincert.pem -nokeys -passin pass:fred
openssl pkcs12 -in domain.p12 -out domainkey.pem -nocerts -passin pass:fred
  -passout pass:fred
```

Tip

For more details, see the OpenSSL user documentation.

Location of SSL private keys and certificates

This section describes the locations of the SSL private keys and certificates for the CA, Node Manager, and API Gateway instances.

Location of CA private key and certificate

The location of the CA private key and certificate depends on how they are signed.

System-generated CA private key and certificate: If SSL certificates are signed by a system-generated CA private key and certificate, the CA private key and certificate are written to disk when the first Admin Node Manager is created by `managedomain`:

- **CA private key**: `apigateway/groups/certs/private/domainkey.pem`
- **CA private key and certificate**: `apigateway/groups/certs/private/domain.p12`
- **CA certificate**: `apigateway/groups/certs/domaincert.pem`

Warning

The files in `apigateway/groups/certs/private` are security sensitive because they contain the private key. This directory must have restricted access. These files are not copied to any other Node Manager, and must be backed up elsewhere. If the CA private key and certificate are lost, you must recreate certificates for the entire domain before adding new Node Managers or API Gateways.

The CA private key files can be encrypted by a passphrase. You must enter this with `managedomain` when creating a new Node Manager or API Gateway. This is used to unlock the CA private key and sign the new SSL certificate.

User-provided CA private key and certificate: If SSL certificates are signed by a user-provided key and certificate, the key and certificate are in a location managed by the user. A `.p12` file containing the private key and certificate must be made available to `managedomain` when a new Node Manager or API Gateway is created. The CA certificate is copied into the API Gateway configuration because it is part of the SSL certificate chain.

External CA private key and certificate: If SSL certificates are signed by an external CA, the CA private key remains with the external third-party CA, or isolated API Gateway installation acting as an external CA. The CA certificate is copied to the API Gateway configuration because it is part of the SSL certificate chain.

Location of Node Manager private key and certificate

The Node Manager private key and certificate chain are stored in the Node Manager configuration in the following directory:

```
apigateway/conf/fed
```

The private key is encrypted using the entity store passphrase.

Location of API Gateway private key and certificate

The API Gateway private key and certificate chain are stored in the following directory:

```
apigateway/groups/group-2/instance-1/conf/certs.xml
```

The private key is encrypted with the group configuration passphrase.

Warning

When Node Managers and API Gateway instances are being registered, private keys are written in `apigateway/groups/certs/private/temp/instance-id` and `nodemanager-id`. Temporary security sensitive files are removed on completion. If an external CA is signing the certificate, sensitive files may exist for a considerable time, so this directory must have restricted access. You can encrypt this sensitive data using the `--key_passphrase` parameter.

Managedomain command reference

Overview

This topic describes how to run `managedomain` in the following modes:

- *Command interpreter mode*—enter commands from a list using tab completion (default mode)
- *Interactive mode*—follow instructions at the command prompt
- *Command mode*—enter specific commands and parameters

The `managedomain` command is available in the following directory:

Windows	`INSTALL_DIR\apigateway\Win32\bin`
UNIX/Linux	`INSTALL_DIR/apigateway/posix/bin`

For an overview of the `managedomain` command, see *Configure an API Gateway domain*.

Managedomain command interpreter mode

To run in default command interpreter mode, enter `managedomain`, and press Tab to view and select options. For example:

```
C:\Oracle-11.1.2.4.0\apigateway\Win32\bin>managedomain
Running in command interpreter mode. Enter 'quit' to exit.
Enter 'help' to view help topics. Enter 'help topic' to view help for a topic.
Press TAB to view and complete commands.
>

add                         add_service_only          add_tag
change_domain_passphrase    change_passphrase         create_archive
create_instance             delete_group              delete_host
delete_instance             delete_tag                deploy
download_archive            edit_group                edit_host
edit_instance               help                      initialize
list_deployments            print_topology            quit
regen_certs                 remove_service            reset
sign_csr                    submit_cert               topology_check
topology_compare            topology_merge            topology_synch
>
```

View help for a command

You can view detailed help for each command and its parameters by entering `help` followed by the command name. The following example shows the help for the `change_passphrase` command:

```
> help change_passphrase

*** Command Description ***

Change group configuration passphrase.

*** Usage ***
```

```
      change_passphrase old_passphrase "" new_passphrase "12345"


*** Command Parameters ***

group : name of API Gateway's group. Required.

old_passphrase : the old passphrase.

new_passphrase : the new passphrase.

anm_host : host name for Admin Node Manager that managedomain connects to in order to execute
commands. This must be specified for the --add command, otherwise a default Admin Node
Manager will be automatically discovered if none is specified. Note that managedomain may
connect to a remote or local Admin Node Manager in order to execute commands.

anm_port : port of Admin Node Manager that managedomain connects to in order to execute
commands. Defaults to 8090.

username : username for authenticating to Admin Node Manager. Defaults to 'admin'.

password : password for authenticating to Admin Node Manager.

truststore_file : the filename for the truststore which holds all trusted Admin Node Manager
SSL certs. Default to trusting all.

truststore_password : the truststore file password which holds all trusted Admin Node Manager
SSL certs.

debug : enable debugging.
```

Run a command

You can run a command using tab completion to specify parameters. The following example shows the available parameters for the `create_instance` command:

```
> create_instance <press TAB>

name                        group                       host
instance_management_port    instance_services_port      passphrase
sign_with_generated         sign_with_user_provided     sign_with_external_ca
domain_passphrase           sign_alg                    ca
key_passphrase              add_service                 service_user
debug                       anm_host                    anm_port
username                    password                    truststore_file
truststore_password
```

The following example creates a new API Gateway instance with a specific name and group:

```
> create_instance name APIServer1 group Group1

Requesting CSR from Admin Node Manager...

CSR received from Admin Node Manager.

Requesting signed certificate from Admin Node Manager...

Signed certificate received from Admin Node Manager.

Requesting Admin Node Manager to create new API Gateway...

New API Gateway SSL certificate details:
```

```
   dname: CN=instance-1,OU=group-2,DC=host-1

   expires: Fri Apr 27 16:36:37 BST 2114

   thumbprint: E6:C5:B5:6D:52:1D:74:3E:CD:3E:8B:B5:82:24:3A:78:1B:C9:27:F9

   issuer dname: CN=Domain

   issuer thumbprint: B5:68:73:7A:7B:ED:6D:B0:04:40:CF:E3:BC:36:6D:84:F7:49:29:12
The new API Gateway 'APIServer1' in group 'Group1' has been successfully created and
installed


Start the new API Gateway by executing the following command:

C:\Oracle-11.1.2.4.0\apigateway\Win32\bin\startinstance -g "Group1 " -n "APIServer1"

You can alternatively add C:\Oracle-11.1.2.4.0\apigateway\Win32\bin\ to your path and
use "startinstance -g "Group1" -n "APIServer1"".

You can test the connection by visiting the URL:

http://roadrunner:8080/healthcheck
```

Tab completion is also available for some parameter values (instance names, group names and host names) The following example shows available instances for the `delete_instance` command:

```
> delete_instance name <press TAB>

APIServer1   APIServer2
> delete_instance name APIServer
```

Managedomain interactive mode

To run in interactive mode, enter `managedomain --menu`, and follow the instructions at the command prompt. The following options are available:

Host Management

The `managedomain --menu` options for host management are as follows:

Option	Description	Why Use this Option
`1`	`Register host`	Add a new host that runs an API Gateway to a domain topology. You must ensure that the host is registered in order to create and run API Gateways. For example, you can specify the following: • If the host is an Admin Node Manager • Host name • Node Manager name • Node Manager port • Node Manager passphrase • Windows/UNIX service for Node Manager • Trust store details

Option	Description	Why Use this Option
		When registering a Node Manager (Admin or not) in an existing domain, you must specify host details of a running Admin Node Manager in the domain. This is not required when registering the first Admin Node Manager in a new domain because this is always an Admin Node Manager.
2	`Edit a host`	Edit the details for a host registered in a domain topology (used occasionally). You can update the following: • Host name • Node Manager name • Node Manager port • Node Manager passphrase • Windows/UNIX service for Node Manager • Change admin capabilities • Enable metrics Changing admin capabilities enables you to change a Node Manager to Admin Node Manager, or an Admin Node Manager to Node Manager. You can only change admin capabilities of the Node Manager running on the same machine. You cannot remove admin capabilities of the last Admin Node Manager in a domain, or from an Admin Node Manager that has the domain key and certificate used to sign CSRs. When you get a license for an evaluation mode API Gateway, you must use this option to change the host from `127.0.0.1` to a network reachable address or host name. You must also restart the Node Manager to pick up any changes.
3	`Delete a host`	Delete a registered host from a domain topology (used occasionally). You must first stop and delete all API Gateways running on the host. You can use this option to delete an Admin Node Manager or Node Manager. The Admin Node Manager that services this request is not allowed to delete itself from the domain, ensuring the domain always has at least one Admin Node Manager.
4	`Change Admin Node Manager and/or credentials, currently connecting as: user admin with truststore None`	By default, you connect to an Node Manager using `managedomain` with the credentials specified at installation time. You can override these at startup by passing the `--username --password` command line parameters, or reset while running `managedomain` with this option. This username/password refers to an `admin` user configured in Policy Studio. You can also use this option to select which Admin Node Manager `managedomain` connects to. `managedomain` must talk to an Admin Node Manager, which can be local or remote. By default, `managedomain` connects to the local running Admin Node Manager, otherwise it searches the topology and uses the first running Admin Node Manager that it finds.

Option	Description	Why Use this Option

API Gateway Management

The `managedomain --menu` options for API Gateway management are as follows:

Option	Description	Why Use this Option
5	`Create API Gateway instance`	Create a new API Gateway instance. You can also do this in Policy Studio and API Gateway Manager. You can create API Gateway instances locally or on any host configured in the topology.
6	`Edit API Gateway (rename, change management port)`	Rename the API Gateway instance, or change the management port. This functionality is not available in Policy Studio and API Gateway Manager.
7	`Delete API Gateway instance`	Delete an API Gateway instance from the topology, and optionally delete the files on disk. You can also do this in Policy Studio and API Gateway Manager. You must ensure that the API Gateway instance has stopped.
8	`Add a tag to API Gateway`	Add a name-value tag to the API Gateway. The **Topology** view on the API Gateway Manager **Dashboard** displays tags and enables you to filter for API Gateway instances by tag.
9	`Delete a tag from API Gateway`	Delete a name-value tag from the API Gateway. The tag will no longer be displayed in the API Gateway Manager **Dashboard**.
10	`Add or remove a Windows/UNIX service for existing local API Gateway`	Must be run by a user with permission to create a service on the host operating system (`root` on Linux, or `Administrator` on Windows). When run on Linux, adds an `init.d` script.

Group Management

The `managedomain --menu` options for group management are as follows:

Option	Description	Why Use this Option
11	`Edit group (rename it)`	Rename an API Gateway group. This functionality is not available in Policy Studio and API Gateway Manager.
12	`Delete a group`	Delete all API Gateways in the group and the group itself. You must ensure that all API Gateways in the group have been stopped first.

Topology Management

The `managedomain --menu` options for topology management are as follows:

Option	Description	Why Use this Option
13	`Print topology`	Output the contents of the deployed domain topology. This includes the following: • Topology version • Hosts • Admin Node Managers • Node Managers • Groups • API Gateway instances (tags)
14	`Check topologies are in sync`	For advanced users. Check that all Node Managers are running the same topology version. Useful only in multi-host environment. Topologies should be in sync if everything is running correctly.
15	`Check the Admin Node Manager topology against another topology`	For advanced users. Compare the two topologies and highlights differences. There should be no differences if everything is running correctly.
16	`Sync all topologies`	For advanced users. Forces a synchronization of all topologies.
17	`Reset the local topology`	For advanced users. Delete the contents of the `apigateway/groups` directory. This means that you must re-register the host and recreate a local API Gateway instance. Alternatively, you can manually delete the contents of this directory to prevent issues if the host has been registered with other node managers.

Deployment

The `managedomain --menu` options for deployment are as follows:

Option	Description	Why Use this Option
18	`Deploy to a group`	Deploy a configuration (`.fed` file) to API Gateways. This functionality is also available in Policy Studio and API Gateway Manager.
19	`List deployment information`	List the deployment information for all API Gateways in a topology. This functionality is also available in Policy Studio and API Gateway Manager.
20	`Create deployment archive`	Create a deployment archive from a directory that contains a federated API Gateway configuration.
21	`Download deployment archive`	Download the `.fed` file deployed to an API Gateway. This functionality is also available in Policy Studio.
22	`Update deployment archive properties`	Update the manifest properties relating to the deployed configuration only. This functionality is also available in Policy Studio. Enables you to update the properties without performing a new deployment.

Option	Description	Why Use this Option
23	`Change group configuration passphrase`	The default passphrase for the API Gateway configuration is “”. Use this option to set a more secure password. This functionality is also available in Policy Studio.

Domain SSL certificates

The `managedomain --menu` options for group management are as follows:

Option	Description	Why Use this Option
24	`Regenerate SSL certificates on localhost`	Regenerate the SSL certificates used to secure API Gateway components in the domain (for example, Node Manager and the API Gateway instances that it manages). You must restart the Node Manager on the localhost after running this option. You must run this option on all hosts in the domain.
25	`Sign CSR`	Specify a Certificate Signing Request (CSR) to send to the Certificate Authority (CA) when applying for an SSL certificate for a Node Manager or API Gateway instance. You can use this option when `managedomain` acts as the CA, and is passed the CSR to create a signed certificate. You will most likely use an external CA in production. However, this option facilitates testing of certificates signed by an external CA. You can install the API Gateway on a locked-down host, and use this feature only (no license required). You would typically only do this when using a system-generated self-signed domain certificate, and do not wish to store the domain private key on an Admin Node Manager host, and do not wish to use an external CA.
26	`Submit externally signed certificate`	Specify an SSL certificate signed by an externally signed Certificate Authority (CA) to be used by a Node Manager or API Gateway instance. Use this option after registering a host or creating an API Gateway using a certificate signed by an external CA. Submitting the certificate with this option completes the host registration or API Gateway creation.

Managedomain command mode

You can also enter `managedomain` commands and parameters directly on the command line. For example, the following command creates an Admin Node Manager on the first host in the domain and signs with a user-provided domain key:

```
managedomain -i --sign_with_user_provided --ca=/home/keys/test.p12
```

For more details and examples of using `managedomain` in command mode, see the following:

- *Configure Admin Node Manager high availability and security*

- *Configure API Gateway for API Gateway Analytics*

Start and stop the API Gateway

Overview

This topic describes how to start and stop the Node Manager and API Gateway instance on the command line, on all platforms. It also describes how to start the Policy Studio graphical tool. For details on API Gateway components and concepts, see the *API Gateway Concepts Guide*.

You can also start and stop API Gateway instances using the API Gateway Manager web console. For more details, see *Configure an API Gateway domain*.

Set passphrases

By default, data is stored unencrypted in the API Gateway configuration store. However, you can encrypt certain sensitive information such as passwords and private keys using a passphrase. When the passphrase has been set, this encrypts the API Gateway configuration data. You must enter the passphrase when connecting to the API Gateway configuration data (for example, using the Policy Studio, or when the API Gateway starts up). For more details on configuring this passphrase, see *Configure an API Gateway encryption passphrase*.

Start the Node Manager

The following instructions describe how to start the Node Manager on the command line on Windows and UNIX:

Windows
Complete the following steps to start the Node Manager on Windows:

1. Open a DOS prompt at the `/Win32/bin` directory of your API Gateway installation.
2. Run the `nodemanager.bat` file.
3. If you are using an encryption passphrase, you are prompted for this passphrase. Enter the correct encryption passphrase, and press Return. For more details, see *Configure an API Gateway encryption passphrase*.

Linux/Solaris
To start the Node Manager on Linux/Solaris, complete the following instructions:

1. Open a shell at the `/posix/bin` directory of your API Gateway installation.
2. Run the `nodemanager.sh` file, for example:

```
prompt# ./nodemanager
```

3. If you are using an encryption passphrase, you are prompted for this passphrase. Enter the correct encryption passphrase and press Return. For more details, see *Configure an API Gateway encryption passphrase*.

Start the API Gateway instance

The following instructions describe how to start the API Gateway instance on the command line on Windows and UNIX:

Windows
Complete the following steps to start the API Gateway on Windows:

1. Open a DOS prompt at the `/Win32/bin` directory of your API Gateway installation.
2. Use the `startinstance` command to start the API Gateway, for example:

```
startinstance -n "my_server" -g "my_group"
```

3. If you are using an encryption passphrase, you are prompted for this passphrase. Enter the correct encryption passphrase, and press Return. For more details, see Setting Passphrases.
4. When the API Gateway has successfully started up, you can run the `policystudio.exe` file from your Policy Studio installation directory.
5. When the Policy Studio is starting up, you are prompted for connection details for the API Gateway. For more details, see Connecting to the API Gateway.

Linux/Solaris
To start the API Gateway instance and the Policy Studio on Linux/Solaris, perform the following steps:

1. Open a shell at the `/posix/bin` directory of your API Gateway installation.
2. Use the `startinstance` command to start the API Gateway, for example:

```
startinstance -n "my_server" -g "my_group"
```

Note

You must ensure that the `startinstance` file has execute permissions.

3. If you are using an encryption passphrase, you are prompted for this passphrase. Enter the correct encryption passphrase and press Return. For more details, see Setting Passphrases.
4. When the API Gateway has successfully started up, run the `policystudio.sh` file in your Policy Studio installation directory. For example:

```
prompt# cd /usr/home/policystudio
prompt# ./policystudio
```

5. When the Policy Studio is starting up, you are prompted for connection details for the API Gateway.

Tip

You can enter the `startinstance` command without any arguments to display the servers registered on the machine. For example:

```
INSTALL_DIR\Win32\bin>startinstance

usage: "startinstance [[-n instance-name -g group-name [instance-args]] |
[directory-location [instance-args]]]"

The API Gateway instances listed below are available to run on this machine
as follows:
   startinstance -n "server1" -g "group1"
   startinstance -n "server2" -g "group2"
```

If you have a single API Gateway instance on the host on which you run `startinstance`, that instance starts when you specify no arguments.

Startup options

You can specify the following optional instance arguments to the `startinstance` command:

Option	Description
`-b <file>`	Specifies the boot-time trace destination.

Option	Description
`-d`	Runs as daemon/service (UNIX only).
`-h <directory>`	Specifies the service instance directory.
`-k`	Kills the instance (UNIX only).
`-P`	Prompts for a passphrase at startup.
`-q`	Runs in quiet mode (equivalent to `-Dquiet`).
`-v`	Changes to the installation instance directory on startup.
`-s`	Tests if the service is running (UNIX only).
`-e`	Specifies the end of arguments for `vshell`.
`-D prop[=value]`	Sets a configuration file property.

Linux/Solaris
The `-d`, `-s` and `-k` options on UNIX are designed for use with the service controller (for example, traditional SVR4 init, systemd, upstart, and so on). The `-d` option waits until the service is running before returning, and `-k` waits for the process to terminate. This means that when used in a script, the completion of the command indicates that the operation requested has completed.

If the service is running, `-s` will exit with a `0` status code, and with `1` otherwise. For example, you can use the following to print a message if the service is running:

```
startinstance -s -n InstanceName -g GroupName && echo Running
```

Windows
The Service Control Manager (SCM) on Windows requires more direct interaction with the process than on UNIX. The product will run as a service, but needs to be started by the SCM.

The `-d` and `-k` options can be used internally as part of the SCM interaction, but are not for interactive use on Windows. These are used to differentiate between a command-line startup and a startup from the SCM. `-s` is not implemented on Windows. When installed as a service, you can control and query the API Gateway using the SCM (for example, using the applet, or the `net` or `sc` commands).

Connect to the API Gateway in Policy Studio

When starting the Policy Studio, you are prompted for details on how to connect to the Admin Node Manager (for example, the server session, host, port, user name, and password). The default connection URL is:

```
https://HOST:8090/api
```

where `HOST` points to the IP address or hostname of the machine on which the API Gateway runs. For more details, see the *API Gateway Policy Developer Guide*.

Stop the API Gateway instance

Linux/Solaris
To stop the API Gateway instance, you must specify the group and instance name to the `startinstance` command along with the `-k` option. For example:

```
./startinstance -g "my_group" -n "my_server" -k
```

Windows
To stop the API Gateway instance, you can enter **Ctrl-C** on the command prompt. If the API Gateway instance is installed as a Windows service, you should use the Windows Services tool.

Stop the Node Manager

Linux/Solaris
To stop the Node Manager, you must specify the `nodemanager` command along with the `-k` option. For example:

```
./nodemanager -k
```

Windows
To stop the Node Manager, you can enter **Ctrl-C** on the command prompt. If the Node Manager is installed as a Windows service, you should use the Windows Services tool.

Start the API Gateway tools

Overview

This topic describes the prerequisites and preliminary steps. It explains how to start the API Gateway Manager administrator tool and the Policy Studio developer tool.

Before you begin

Before you start the API Gateway tools, do the following:

Install the API Gateway and Policy Studio
If you have not already done so, see the *API Gateway Installation and Configuration Guide*.

Configure a managed domain
If you have not already created a domain, you can use the `managedomain` script to configure a domain. You should also ensure that the Admin Node Manager and an API Gateway instance are running.

Launch API Gateway Manager

To access the web-based API Gateway Manager administration tools, perform the following steps:

Note

You must ensure that the Admin Node Manager is running before you can access the web-based API Gateway Manager tools.

1. Enter the following URL:

   ```
   https://HOST:8090/
   ```

   `HOST` refers to the host name or IP address of the machine on which the API Gateway is running (for example, `https://localhost:8090/`).
2. Enter the administrator user name and password configured at installation time.
3. Click the appropriate button in the API Gateway Manager page in the browser. The **Dashboard** view is displayed by default.

The API Gateway Manager includes the following main views:

- **Dashboard**: System health, traffic summary, and topology (domain, hosts, API Gateways, and groups).
- **Traffic**: Message log and performance statistics on the traffic processed by the API Gateway. For example, all HTTP, HTTPS, JMS, File Transfer, and Directory messages processed by the API Gateway.
- **Monitoring**: Real-time monitoring of all the traffic processed by the API Gateway. Includes statistics at the system level and for services invoked and remote hosts connected to.
- **Logs**: API Gateway trace log, transaction log, and access log files.
- **Events**: API Gateway transaction log points, alerts, and SLA alerts.
- **Settings**: Enables you to configure dynamic API Gateway logging, user roles, and credentials.

Start Policy Studio

To start the Policy Studio tool used to create and manage policies, perform the following steps:

1. In your Policy Studio installation directory, enter the `policystudio` command.
2. In the **Policy Studio**, click a link to connect to the Admin Node Manager session.
3. In the **Open Connection** dialog, enter the administrator user name and password and click **OK**.
4. The **API Gateway** instance is displayed in the **Topology** view.
5. In the **Topology** view, double-click the **API Gateway** instance to load the configuration for the active API Gateway.
6. If Node Manager credential checking is enabled, enter the administrator user name and password in the Node Manager credentials dialog, and click **OK**. To disable credential checking for future deployment or topology operations, deselect the **Always prompt for user credentials** check box.
7. If a passphrase has been set, enter it in the **Enter Passphrase** dialog, and click **OK**. Alternatively, if no passphrase has been set, click **OK**. For more details, see *Configure an API Gateway encryption passphrase*.

Policy Studio enables you to perform the full range of API Gateway configuration and management tasks. This includes tasks such as develop and assign policies, import services, optimize API Gateway configuration settings, and manage API Gateway deployments. For more details on using the Policy Studio to manage API Gateway processes and configurations, see *Manage API Gateway deployments*.

Run API Gateway as non-root on UNIX/Linux

Overview

In a typical deployment on Linux or Solaris, the process for the API Gateway instance runs as `root`. This is typically used to enable the process to listen on Internet domain privileged ports (port numbers less than `1024`). However, this is a potential security issue because the root user has read/write access to all files on the system. Therefore, if the API Gateway process ever becomes compromised, it could be used to return the contents of, or overwrite sensitive files, and the operating system would not prevent this. A solution to this problem is to run the API Gateway as a non-root user, but still allow the API Gateway process to bind to privileged ports.

Important

The steps in this topic apply to the following API Gateway processes:

- API Gateway instance
- Admin Node Manager
- API Gateway Analytics

Linux capabilities

Traditional UNIX implementations distinguish between the following categories of processes:

- *privileged* processes whose effective user ID is `0` (`superuser` or `root`)
- *unprivileged* processes whose effective user ID is non-zero

Privileged processes bypass all kernel permission checks where unprivileged processes are subject to full permissions checking based on the processes credentials, usually from the effective user ID and effective group ID. More recent versions of the kernel divide up the privileges traditionally associated with the `superuser` into a set of capabilities that can be independently enabled or disabled. This allows a more fine-grained control of permissions for a process.

The capability to use with the API Gateway is `CAP_NET_BIND`, which specifically allows binding to privileged ports. The method by which this capability is set on the API Gateway is supported from kernel 2.6.24 onwards. If the kernel version is before 2.6.33, you must enable `CONFIG_SECURITY_FILE_CAPABILITIES`.

Before you begin

The sections that follow describe the steps that you must perform to run the API Gateway as an unprivileged user. This topic assumes that the new unprivileged user is named `admin`, and that the location of your API Gateway installation is `/home/admin/apigateway`.

Modify API Gateway file ownership

The unprivileged user must have ownership of the API Gateway files. If there is a pre-existing API Gateway installation, you must change the ownership of the API Gateway files to the new user that you intend to run the API Gateway with.

You can use the following command to change the user and group ownership of all files under the installation directory:

```
# chown -R user:usergroup /path/to/apigateway/installation
```

For example:

```
# chown -R admin:admin /home/admin/apigateway
```

SSL accelerators for HSM

When using a Hardware Security Module (HSM), the unprivileged user must have access to the device corresponding to the crypto accelerator or HSM card. For HSMs such as Cavium and Ultimaco, this means that you must have access to the following directories:

- **Cavium**: `/dev/pkp_nle_drv`
- **Utimaco**: `/dev/cs2a`

For example, when using Cavium, an `admin` user using `/dev/pkp_nle_drv` should have the following permissions:

```
crw-rw-r-- 1 root admin 126, 0 /dev/pkp_nle_dev
```

If the unprivileged user is different from `admin`, run the following command:

```
# sed -i /admin/d /lib/udev/load_{cavium,utimaco}
```

Set the CAP_NET_BIND capability on vshell

You must add the Linux capability to allow the API Gateway to listen on ports less than `1024` to the `vshell` file.

You can use the following command:

```
# setcap 'cap_net_bind_service=+ep' /path/to/apigateway/installation//platform/bin/vshell
```

For example:

```
# setcap 'cap_net_bind_service=+ep' /home/admin/apigateway/platform/bin/vshell
```

To verify that the permission has been set, run the following command:

```
# getcap /home/admin/apigateway/platform/bin/vshell
```

For example, the output of this command should be as follows:

```
/home/admin/apigateway/platform/bin/vshell = cap_net_bind_service+ep
```

Note

For API Gateway versions prior to 6.3.0 (for example, 6.1.2), this path should be `platform/libexec/vshell` instead of `platform/bin/vshell`.

Install the libcap2 package if required

Depending on your Linux distribution, this may involve installing an additional software package. If the `setcap` command cannot be found, try installing the `libcap2` package.

For yum-based systems (for example, Redhat Enterprise Linux, CentOS, Oracle Enterprise Linux), you can use the following command:

```
# yum install libcap2
```

For Debian or Ubuntu, you can use the following command:

```
# apt-get install libcap2-bin
```

API Gateway appliance version 7.1.0 or later

You must set an additional privilege for API Gateway appliance version 7.1.0 or later. This step applies if you have an appliance and wish to run the `vshell` processes as a user other than the default `admin` user.

Run the following command (all on one line):

```
# setcap 'cap_net_bind_service=+ep cap_sys_rawio=+ep'
   /path/to/apigateway/installation/platform/bin/vshell
```

For example:

```
# setcap 'cap_net_bind_service=+ep cap_sys_rawio=+ep'
   /home/admin/apigateway/platform/bin/vshell
```

To verify that the permissions have been set, run the following command:

```
# getcap /home/admin/apigateway/platform/bin/vshell
```

For example, the output of this command should be as follows:

```
/home/admin/apigateway/platform/bin/vshell = cap_net_bind_service,cap_sys_rawio+ep
```

Add API Gateway library locations

When executing a file with elevated permissions, certain environment variables are disabled as a security precaution. For this reason, you must make the locations of the API Gateway library files available to the operating system. You can do this using the steps described in this section:

1. Create an `ld.so.conf` file with the API Gateway shared object locations
2. Run `ldconfig` to reload the `ld.so.cache` file

Create the ld.so.conf file

For a 64-bit API Gateway installation, create the following file:

```
/etc/ld.so.conf.d/gateway-libs.conf
```

And add the following lines:

```
/home/admin//apigateway/platform/jre/lib/amd64/server
/home/admin/apigateway/platform/jre/lib/amd64
/home/admin/apigateway/platform/lib/engines
/home/admin/apigateway/platform/lib
/home/admin/apigateway/ext/lib
```

For a 32-bit API Gateway installation, create the following file:

```
/etc/ld.so.conf.d/gateway-libs.conf
```

And add the following lines:

```
/home/admin/gateway/platform/jre/lib/i386/server
/home/admin/gateway/platform/jre/lib/i386
/home/admin/gateway/platform/lib/engines
/home/admin/gateway/platform/lib
/home/admin/gateway/ext/lib
```

Note

You should replace `/home/admin/gateway` with the root of your API Gateway installation.

Run ldconfig

After creating the `/etc/ld.so.conf.d/gateway-libs.conf` file, run the following command to reload the library cache file:

```
# ldconfig
```

Modify the init.d script to use sudo

If there is an `init.d` script that runs the API Gateway as a service, you must change this to execute as the unprivileged user. The easiest way is by using `sudo -u admin` in the start command of the script.

For API Gateway version 6.1, this means changing the following line in the `/etc/init.d/vordelgateway` script:

```
daemon --pidfile $PIDFILE $DAEMON $DAEMON_ARGS >> $LOGFILE 2>&1
```

to the following:

```
daemon --pidfile $PIDFILE --user=admin $DAEMON $DAEMON_ARGS >> $LOGFILE 2>&1
```

Modify the jvm.xml file

To modify your `jvm.xml` file, perform the following steps:

1. Open the `system/conf/jvm.xml` file in your API Gateway installation.
2. Near the top of the file, add an extra line after the following line:

   ```
   <JVMSettings classloader="com.vordel.boot.ServiceClassLoader">
   ```

3. For a 64-bit API Gateway installation, add the following:

   ```
   <VMArg name="-Djava.library.path=$VDISTDIR/$DISTRIBUTION/jre/lib/amd64/server:
      $VDISTDIR/$DISTRIBUTION/jre/lib/amd64:$VDISTDIR/$DISTRIBUTION/lib/engines:
      $VDISTDIR/ext/$DISTRIBUTION/lib:$VDISTDIR/ext/lib:$VDISTDIR/$DISTRIBUTION/jre/lib:
      system/lib:$VDISTDIR/$DISTRIBUTION/lib"/>
   ```

4. For a 32-bit API Gateway installation, add the following:

   ```
   <VMArg name="-Djava.library.path=$VDISTDIR/$DISTRIBUTION/jre/lib/i386/server:
      $VDISTDIR/$DISTRIBUTION/jre/lib/i386:$VDISTDIR/$DISTRIBUTION/lib/engines:
      $VDISTDIR/ext/$DISTRIBUTION/lib:$VDISTDIR/ext/lib:$VDISTDIR/$DISTRIBUTION/jre/lib:
      system/lib:$VDISTDIR/$DISTRIBUTION/lib"/>
   ```

Restart the API Gateway

When you have completed the steps described in this topic, start the API Gateway with the unprivileged user.

Run API Gateway as non-root on Solaris

On Solaris 10, a similar system is provided by adding the `net_privaddr` privilege to a user.

Note

On Solaris, you must modify the user instead of the API Gateway process (as already described for Linux in this topic).

In this case, you can modify an existing user (for example, `gwadmin`) as follows:

```
# usermod -K defaultpriv=basic,net_privaddr gwadmin
```

Configure API Gateway high availability

Overview

System administrators can configure High Availability (HA) in an API Gateway environment to ensure that there is no single of point of failure in the system. This helps to eliminate any potential system downtime in production environments. Typically, the API Gateway platform is deployed in the Demilitarized Zone (DMZ) to provide an additional layer of security for your backend systems.

This topic describes the recommended API Gateway architecture in an HA production environment. It includes recommendations on topics such as load balancing, commonly used transport protocols, caching, embedded persistence, and connections to external systems.

Tip

For details on how to configure and secure multiple Admin Node Managers in a domain, see *Configure Admin Node Manager high availability and security*.

HA in production environments

The following diagram shows an overview of an API Gateway platform running in an HA production environment:

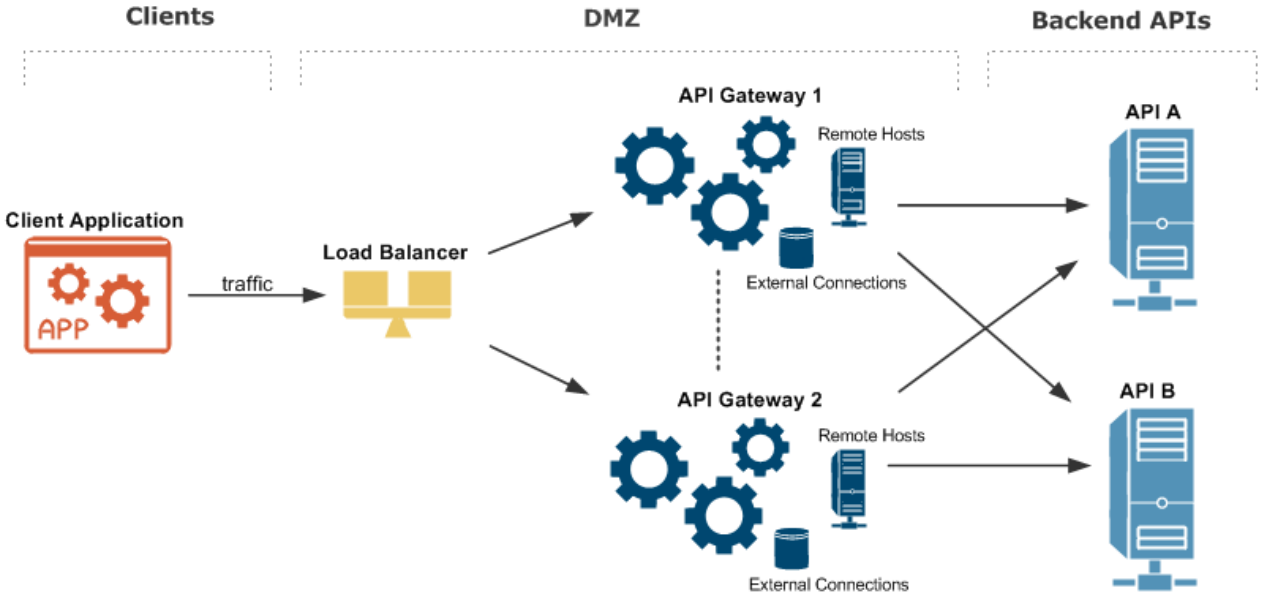

The architecture shown in the diagram is described as follows:

- An external client application makes inbound calls sending business traffic over a specific message transport protocol (for example, HTTP, JMS, or FTP) to a load balancer.
- A standard third-party load balancer performs a health check on each API Gateway instance, and distributes the message load to the listening port on each API Gateway instance (default is `8080`).
- Each API Gateway instance has External Connections to third-party systems. For example, these include databases such as Oracle and MySQL, and Authentication Repositories such as CA SiteMinder, Oracle Access Manager, Local Directory Access Protocol (LDAP) servers, and so on.
- Caching is replicated between each API Gateway instance using a distributed caching system based on Ehcache.

- Each API Gateway instance has Remote Host interfaces that specify outbound connections to backend API systems, and which can balance the message load based on specified priorities for Remote Hosts.
- Each API Gateway instance contains an embedded Apache Cassandra database used by certain features for persistent data storage, and which has its own HA capabilities.
- Each API Gateway instance contains an embedded Apache ActiveMQ messaging system, which can be configured for HA in a shared filesystem.
- Each backend API is also replicated to ensure there is no single point of failure at the server level.
- Management traffic used by the Admin Node Manager, API Gateway Manager, and Policy Studio is handled separately on different port (default is `8090`).

Important

For simplicity, the diagram shows two API Gateway instances only. However, for a resilient HA system, a minimum of at least two active API Gateway instances at any time, with a third and fourth in passive mode, is recommended.

Load Balancing

In this HA production environment, the load balancer performs a health check and distributes message load between API Gateway instances. The API Gateway supports a wide range of standard third-party load balancing products (for example, F5) without any special requirements. Multiple API Gateway instances can be deployed active/active (hot standby) or active/passive (cold standby) modes as required.

The load balancer polls each API Gateway instance at regular intervals to perform a health check on the message traffic port (default `8080`). The load balancer calls the API Gateway Health Check policy, available on the following default URL:

```
http://GATEWAY_HOST:8080/healthcheck
```

The health check returns a simple `<status>ok</status>` message when successful. In this way, if one API Gateway instance becomes unavailable, the load balancer can route traffic to an alternative API Gateway instance.

Both transparent load balancing and non-transparent load balancing are supported. For example, in transparent load balancing, the API Gateway can see that incoming messages are sent from specific client and load balancer IP addresses. The API Gateway can also extract specific client details from the HTTP header as required (for example, the SSL certificate, user credentials, or IP address for Mutual or 2-Way SSL Authentication). In non-transparent load balancing, the API Gateway sees only the virtual service address of the load balancer.

In addition, the API Gateway can also act as load balancer on the outbound connection to the backend APIs. For more details, see the section called “Remote Hosts”.

Java Message System

The API Gateway supports integration with a wide range of third-party Java Message System (JMS) products. For example, these include the following:

- SonicMQ
- Tibco Rendezvous
- Fiorano
- OpenJMS
- JBoss Messaging

The API Gateway can act as a JMS client (for example, polling messages from third-party JMS products or sending mes-

sage to them). For details on configuring API Gateway client connections to JMS systems, see the *API Gateway Policy Developer Guide*. For details on configuring HA in supported third-party JMS systems, see the user documentation available from your JMS provider.

The API Gateway also provides an embedded Apached ActiveMQ server in each API Gateway instance. For more details, see the section called “Embedded Apache ActiveMQ”.

File Transfer Protocol

The API Gateway supports the following protocols:

- File Transfer Protocol (FTP)
- FTP over SSL (FTPS)
- Secure Shell FTP (SFTP)

When using FTP protocols, the API Gateway writes to a specified directory in your filesystem. In HA environments, when the uploaded data is required for subsequent processing, you should ensure that the filesystem is shared across your API Gateway instances—for example, using Storage Area Network (SAN) or Network File System (NFSv4). For more details on configuring FTP connections, see the *API Gateway Policy Developer Guide*.

Remote Hosts

You can use the **Remote Host Settings** in Policy Studio to configure how the API Gateway connects to a specific external server or routing destination. For example, typical use cases for configuring Remote Hosts with the API Gateway are as follows:

- Mapping a hostname to a specific IP address or addresses (for example, if a DNS server is unreliable or unavailable), and ranking hosts in order of priority.
- Specify load balancing settings (for example, whether load balancing is performed on a simple round-robin basis or weighted by response time).
- Allowing the API Gateway to send HTTP 1.1 requests to a destination server when that server supports HTTP 1.1, or resolving inconsistencies in how the destination server supports HTTP.
- Setting timeouts, session cache size, input/output buffer size, and other connection-specific settings for a destination server. For example, if the destination server is particularly slow, you can set a longer timeout.

For details on how to configure **Remote Host Settings** in Policy Studio, see *API Gateway Policy Developer Guide*.

Distributed caching

In an HA production environment, caching is replicated between each API Gateway instance using a distributed caching system. In this scenario, each API Gateway has its own local copy of the cache, but registers a cache event listener that replicates messages to the caches on other API Gateway instances. This enables the put, remove, expiry, and delete events on a single cache to be replicated across all other caches.

In the distributed cache, there is no master cache controlling all caches in the group. Instead, each cache is a peer in the group that needs to know where all the other peers in the group are located. Peer discovery and peer listeners are two essential parts of any distributed cache system.

For more details on configuring distributed cache settings, see the topic on Global Caches in the *API Gateway Policy Developer Guide*. The API Gateway distributed caching system is based on Ehcache. For more details, see http://ehcache.org/.

External Connections

You can use **External Connections** settings in Policy Studio to configure how the API Gateway connects to specific external third-party systems. For example, this includes connections such as the following:

- Authentication Repositories (LDAP, CA SiteMinder, Oracle Access Manager, and so on)
- Databases (Oracle, DB2, MySQL, and MS SQL Server)
- JMS Services
- SMTP Servers
- ICAP Servers
- Kerberos
- Tibco
- Tivoli
- Radius Clients

For details on how to configure **External Connections** in Policy Studio, see *API Gateway Policy Developer Guide*. For details on how to configure HA for any external third-party systems, see the product documentation provided by your third-party vendor.

Important

When configuring connections to server hosts, it is recommended that you specify server hostnames instead of IP addresses, which are less stable and are subject to change in a Domain Name System (DNS) environment.

Embedded Apache ActiveMQ

The API Gateway provides an embedded Apache ActiveMQ broker in each API Gateway instance. In a HA production environment, multiple ActiveMQ broker can work together as a network of brokers in a group of API Gateways. This requires setting up a shared directory that is accessible from all API Gateway instances—for example, using Storage Area Network (SAN) or Network File System (NFSv4).

In this shared network of ActiveMQ brokers, each API Gateway can start a local embedded ActiveMQ broker, which listens on a configured TCP port (61616 by default). This port is accessible from both the local API Gateway and remote clients using the load balancer.

For details on how to configure and manage embedded Apache ActiveMQ, see the following:

- *Embedded ActiveMQ settings*
- *Manage embedded ActiveMQ messaging*

For more details on Apache ActiveMQ, see http://activemq.apache.org/.

Embedded Apache Cassandra database

Each API Gateway instance contains an embedded Apache Cassandra database for default persistent data storage. This Cassandra database is used by the following API Gateway features:

- OAuth tokens and codes
- Client Application Registry
- API Keys
- KPS collections that you create

If your API Gateway system uses any of these features, you must configure the embedded Apache Cassandra database for HA. The embedded Cassandra database is API Gateway group-aware, which means that all API Gateways in a group can share the same Cassandra data source. In a production environment, you must configure the Cassandra data source to provide HA for the API Gateway group.

Note

All nodes in a Cassandra cluster must run on the same operating system (for example, Windows, Linux, or Solaris).

For more details on configuring the embedded Cassandra database for HA, see the *API Gateway Key Property Store User Guide*.

API Gateway backup and disaster recovery

Overview

You must ensure that your API Gateway system can recover from any natural disasters (for example, floods, hurricanes, or earthquakes) and human-induced disasters (for example, failures, fires, or explosions).

Many organizations have a mirrored backup and disaster recovery site with full capacity to recover from any major incidents. These systems are typically kept in a separate physical location on cold stand-by until they need to be brought into action. In this case, the backup and disaster site must be a mirrored image of production environment that replicates all resources and assets (for example, LDAP and databases), and with the same number of API Gateway instances. You should ensure that any required third-party systems are included in your backup and recovery solution.

Example approaches to keeping production and backup environments in sync include making backups to disk or tape, and sending these off-site at regular intervals, or cloud-based solutions that replicate on-site systems, and back up to off-site data centers.

Components that must be backed up

Whichever backup strategy you choose, in a production environment you must ensure that API Gateway installations on all API Gateway host nodes are backed up on a regular basis. For example, this includes hosts that run the following components:

- API Gateway instance
- Admin Node Manager
- Node Manager
- API Gateway Analytics

You must also back up all databases and third-party systems used with the API Gateway. For example, this includes the following:

- Databases used by API Gateway Analytics to store reports (for example, MySQL, Oracle, MS SQL, or DB2)
- Shared disks used by embedded ActiveMQ to store JMS messages
- Any databases or third-party systems that the API Gateway connects to in External Connections

Note

You do not need to back up Policy Studio, Configuration Studio, or API Gateway Explorer because these tools run in a temporary workspace when required. However, if you have modified any third-party dependencies on the **Preferences** page (for example, to connect to a specific database), you must also add the relevant `.jar` on the **Runtime Dependencies** page in your disaster recovery environment.

Back up API Gateway

To back up an API Gateway installation, you must back up files that have changed in the following directory:

```
<install-dir>\apigateway
```

This backs up all Admin Node Manager, Node Manager, , and API Gateway instances in that installation.

For example, the following directories include API Gateway configuration, and will typically need to be backed up on a regular basis:

```
<install-dir>\apigateway\conf
<install-dir>\apigateway\groups
<install-dir>\apigateway\ext
<install-dir>\apigateway\system\conf\nodemanager.xml
```

Note

Before backing up, you can remove the contents of the `apigateway\conf\opsdb.d` directory. This contains transient monitoring data, which can be quite large in some cases, and does not need to be backed up.

Back up API Gateway Analytics

Similarly, to back up an API Gateway Analytics installation, you must back up files that have changed in the following directory:

```
<install-dir>\oaganalytics
```

For example, the following directories include API Gateway Analytics configuration, and will typically need to be backed up on a regular basis:

```
<install-dir>\oaganalytics\conf
<install-dir>\oaganalytics\ext
```

Note

This backs up the API Gateway Analytics installation only. You must also back up the metrics database separately. For more details, see the next section.

Back up databases and third-party systems

You must back up all databases and third-party systems used with the API Gateway. For example, you can back a MySQL database by creating a dump file of the tables in use:

```
mysqldump -u root temp_backup > db_tables.dump
```

For more details, see the user documentation for your third-party system.

Disaster recovery plan and tests

To ensure that your backup and disaster recovery processes are successful, you should conduct full restoration tests on a regular basis. You must ensure that you can restore all required files as quickly and easily as possible.

To ensure this, your backup and disaster recovery plan should include key metrics for recovery points and recovery times for your real-world business processes (for example, creating a purchase order, booking a reservation, and so on).

Example of creating an API Gateway disaster recovery site

This simple example shows how to create a disaster recovery site from a backup of an API Gateway production deployment. It assumes that both the disaster recovery site and the primary production site have the same version of API Gateway installed. In this scenario, the disaster recovery site is a cold standby, and the configuration from production is replicated using a backup of production configuration.

Back up the production environment

To back up the production environment, perform the following steps.

1. Browse to the directory where the API Gateway is installed (for example, `/opt/apigateway` or `c:\apigateway`).
2. Tar or zip the following:
   - `apigateway/groups`
   - `apigateway/conf`
   - `apigateway/system/conf/nodemanager.xml`
   - `apigateway/ext`

   On Linx and UNIX, if you want the API Gateway and Node Manager to start up automatically on the new host, you should also include `/etc/init.d/vshell-*`.

Note

This includes separate startup scripts files for the Node Manager and API Gateway instances if an `init.d` script was created using `managedomain` during initial setup. You can create these at any time using `managedomain`, and choosing option 2, Edit a host, for a Node Manager, or option 10, Add script or service for an existing local API Gateway. For more details, see *Managedomain command reference*.

For example, the following command creates a `.tar` file running from the root directory:

```
>tar -cvf apigateway_backup.tar /opt/apigateway/conf /opt/apigateway/groups
  /opt/apigateway/system/conf/nodemanager.xml
```

The following example creates a `.tar` file containing the startup scripts running from the root directory:

```
> tar -cvf startup_scripts.tar /etc/init.d/vshell-*
```

Copy to the disaster recovery site

To replicate to the disaster recovery site:

1. Copy the created `.tar` file(s) from the primary production machine to the cold standby machine.

   Note

   Ensure the files are tarred before copying because this preserves the permissions and ownership of the files.
2. Untar or unzip the files so that they are extracted in the same directories, overwriting existing files if necessary.

   The following example untars the files from the root directory:

   ```
   >tar -C / -xvf apigateway_backup.tar
   ```

   The following is optional:

   ```
   >tar -C / -xvf startup_scripts.tar
   ```
3. When all the files are copied over and extracted successfully, you should be able to start the Admin Node Manager and API Gateway instances the same way as in primary production site running the same topology and configurations.

Note

This example assumes that backups are collected on regular basis from the master node in the production site.

Further Information

For details on how to back up and restore an Admin Node Manager for signing SSL certificates in a API Gateway domain, see the section called “Admin Node Manager backup and disaster recovery”.

Manage API Gateway settings

Overview

You can configure the underlying settings for the API Gateway using the **Server Settings** node in the Policy Studio tree. Alternatively, select the **Tasks** > **Manage Gateway Settings** menu option in the Policy Studio main menu. This topic describes the settings available in the **Server Settings** screen. Click or expand a tree node on this screen to configure the appropriate settings. You can confirm changes to these settings by clicking the **Apply Changes** button at the bottom right of each screen.

General settings

The top-level **General** settings are applied to all instances of the API Gateway that use this configuration. For example, you can change the tracing level, various timeouts and cache sizes, and other such global information. For more details, see *General settings*.

In addition, you can also configure the following settings under the **General** node:

Cache

If you have deployed several API Gateways throughout your network, you should configure a distributed cache. In a distributed cache, each cache is a peer in a group and needs to know where all the other peers in the group are located. The **Cache** settings enable you to configure settings for peer listeners and peer discovery. For more details, see the *API Gateway Policy Developer Guide*.

MIME/DIME

The API Gateway can filter Multipurpose Internet Mail Extensions (MIME) messages based on the content types (or MIME types) of the individual parts of the message. This also applies to Direct Internet Message Encapsulation (DIME), which is a streamlined version of MIME. The MIME/DIME settings list the default MIME/DIME types that the API Gateway can filter on. These types are then used by the **Content Types** filter to determine which MIME/DIME types to block or allow through to the back end Web service. For more details, see *MIME/DIME settings*.

Namespaces

The **Namespaces** settings are used to determine the versions of SOAP, Web Services Security (WSSE) and Web Services Utility (WSU) that the API Gateway supports. For more details, see *Namespace settings*.

HTTP session

The **HTTP Session** settings enable you to configure HTTP session management settings for the selected cache. For example, you can configure the period of time before expired sessions are cleared from the default `HTTP Sessions` cache. For more details, see *HTTP Session settings*.

Logging settings

The **Logging** settings enable you to configure the following:

Transaction Audit Log

The **Transaction Audit Log** settings enable you to configure the default message transaction logging behavior of the API Gateway. For example, you can configure the API Gateway to log to a database, text or XML file, local or remote UNIX syslog, or the system console. For more details, see the topic on *Transaction audit log settings*.

Transaction Access Log

The **Transaction Access Log** records a summary of all request and response messages that pass through the API Gateway. For example, this includes details such as the remote hostname, username, date and time, first line of the request message, HTTP status code, and number of bytes. For details on configuring these settings per API Gateway, see *Transaction access log settings*.

Transaction Event Log

The **Transaction Event Log** provides a summary of each API Gateway message transaction, which is written to a log file, and used to generate metrics for API Gateway monitoring (for example, in the **Monitoring** view in API Gateway Analytics). For details on configuring these settings per API Gateway, see *Transaction event log settings*.

Messaging settings

The **Messaging** settings enable you to configure settings for the embedded Apache ActiveMQ server that is available in each API Gateway instance. For example, these include the listening interface, port, shared directory, and so on. For more details, see *Embedded ActiveMQ settings*.

Monitoring settings

The **Monitoring** settings enable you to configure the following:

Real Time Monitoring

The **Real Time Monitoring** settings enable you to configure statistics about messages that API Gateway instances store in a database. The API Gateway Analytics monitoring tool can then poll this database, and produce charts and graphs showing how the API Gateway is performing. For more details, see *Real-time monitoring metrics*.

Traffic Monitor

The **Traffic Monitor** settings enable you to configure the web-based Traffic Monitor tool and its message traffic log. For example, you can configure where the data is stored and what message transaction details are recorded in the log. For more details, see *Traffic monitoring settings*.

Security settings

The **Security** settings enable you to configure the following:

Security Service Module

You can configure the API Gateway to act as an Oracle Security Service Module (SSM) to enable integration with Oracle Entitlements Server 10g. The API Gateway acts as a Java SSM, which delegates to Oracle Entitlements Server 10g. For example, you can authenticate and authorize a user for a particular resource against an Oracle Entitlements Server 10g repository. For more details, see the *API Gateway Policy Developer Guide*.

Important

Oracle SSM is required for integration with Oracle OES 10g only. Oracle SSM is not required for integration with Oracle OES 11g.

Kerberos

You can configure Kerberos settings such as the Kerberos configuration file to the API Gateway, which contains information about the location of the Kerberos Key Distribution Center (KDC), encryption algorithms and keys, and domain realms. You can also configure options for APIs used by the Kerberos system, such as the Generic Security Services (GSS) and Simple and Protected GSSAPI Negotiation (SPNEGO) APIs. For more details, see the *API Gateway Policy*

Developer Guide.

Tivoli

You can configure how the API Gateway instance connects to an instance of an IBM Tivoli Access Manager server. Each API Gateway instance can connect to a single Tivoli server. For more details, see the *API Gateway Policy Developer Guide*.

Configure an API Gateway encryption passphrase

Overview

By default, the API Gateway configuration data is stored unencrypted. However, you can encrypt sensitive information such as passwords and private keys using a passphrase. When the passphrase has been set (and the data has been encrypted with it), you must then enter the passphrase when connecting to the API Gateway with the Policy Studio, or when the API Gateway is starting up, so that the encrypted data can be decrypted. You can enter the encryption passphrase at the level of an API Gateway group.

All sensitive information in the API Gateway configuration data is encrypted when you set an encryption passphrase. For example, this sensitive information includes passwords that the API Gateway requires for connecting to external systems (databases, LDAP, and so on), or private keys that are not stored on a Hardware Security Module (HSM). All sensitive information is encrypted using the Password-Based Encryption (PBE) system specified by the Public-Key Cryptography Standard (PKCS#12). For more details, see Appendix B of the PKCS#12 documentation.

This topic describes how to specify the encryption passphrase when connecting to the API Gateway with the Policy Studio, in your API Gateway configuration file, or when the API Gateway is starting up. It also describes how to change the API Gateway group passphrase when it has been set initially.

Warning

It is *crucial* that you remember the passphrase when you change it. Failure to remember the passphrase results in the loss of encrypted data, and may prevent the API Gateway from functioning correctly.

Configure the passphrase in Policy Studio

You can use the Policy Studio topology view to set the passphrase used to encrypt the API Gateway group data. This is the table displayed when you connect to the Admin Node Manager. To change the passphrase, right-click the API Gateway group name in the table (for example, **QuickStart Group**), and select **Change Passphrase**.

Complete the following fields on the **Change Group Passphrase** dialog:

Old Passphrase:
Enter the old passphrase that you wish to change in this field. Alternatively, you can leave this field blank if you are setting the passphrase for the first time.

New Passphrase:
Enter the new passphrase.

Confirm New Passphrase:
Re-enter the new passphrase to confirm it.

Important

You must also re-encrypt the Key Property Store tables after the passphrase has been changed. You can do this using the `kpsadmin` tool. For more details, see the *Key Property Store User Guide*, which is available from Axway Support.

Enter the passphrase when you edit configuration in Policy Studio

If you have set an encryption passphrase for the API Gateway group configuration data, you must enter this passphrase every time you open a configuration for editing in Policy Studio. You can specify this in the **Enter Passphrase** dialog, which is displayed before editing an active server configuration.

When you first open a connection to an API Gateway in Policy Studio, you specify a **Password**. The different roles of the **Passphrase** and the **Password** fields are as follows:

Passphrase	Used to decrypt sensitive data that has already been encrypted (for example, private keys and passwords) . Not required by default, and only needed if you have already set the encryption passphrase in Policy Studio.
Password	Used to authenticate to the API Gateway's management interface using HTTP basic authentication when opening a connection to a server. Required by default.

Provide the passphrase in a configuration file or at startup

For the API Gateway to read (decrypt) encrypted data from its configuration, it must be primed with the passphrase key. You can enter the passphrase directly in a configuration file, prompt for it at startup, or obtain it automatically with a script.

Note

Typically, the passphrase is only entered directly in the file if the server must be started as a Windows service or UNIX daemon. In this case, the administrator cannot enter the passphrase manually when the server is starting. To avoid this, you must enter the passphrase in the configuration file.

Enter the Node Manager passphrase in a configuration file

You can enter a passphrase directly in the Node Manager configuration file. Open the following file in your API Gateway installation:

```
INSTALL_DIR/system/conf/nodemanager.xml
```

This file contains values for general system settings, such as the server name and trace level, and also (if required) the passphrase key that the Node Manager uses to decrypt its own configuration data.

You should specify the passphrase as the value of the `secret` attribute as follows:

```
secret="myPassphrase"
```

Enter the API Gateway passphrase in a configuration file

You can also enter the passphrase for API Gateway instances created using the `managedomain` script. To do this, enter the `secret` attribute in the `group.xml` file for your API Gateway instance. For example:

```
INSTALL_DIR/groups/group-id/conf/group.xml
```

Prompt for the passphrase at server startup

If you do not wish to enter the passphrase directly in the Node Manager or API Gateway instance configuration file, and do not need to start as a Windows service or UNIX daemon, you can configure the Node Manager or API Gateway to prompt the administrator for the passphrase on the command line when starting up. To do this, enter the `"(prompt)"` special value for the `secret` attribute as follows:

```
secret="(prompt)"
```

To configure this for the Node Manager, update your `nodemanager.xml` file. To configure for an API Gateway group, update the relevant `group.xml` file.

Provide the passphrase automatically at startup using a script

Alternatively, you can use a script to automatically provide the passphrase when the API Gateway server starts up. Perform the following steps:

1. Open the following file in your API Gateway installation:

   ```
   INSTALL_DIR/system/conf/service.xml
   ```

2. Add the following to your `service.xml` file:

   ```
   <SystemSettings tracelevel="INFO"
          passphraseExec="&amp;quot;$VINSTDIR/passwd.bat&amp;quot;"
          serviceID="${serviceID}" groupID="${groupID}"
          serviceName="${serviceName}" groupName="${groupName}"
          domainID="${domainID}" title="API Gateway"/>
   ```

   Note

   The `passphraseExec` option is only used if it is present and the `secret` option (described in the previous sections) is not used.

3. Create the passphrase script in the specified location. For example, on Windows, the contents of the `passwd.bat` script is as follows:

   ```
   @echo off
   echo secret
   ```

   For example, on Linux/UNIX, the contents of the `passwd.sh` script is as follows:

   ```
   #!/bin/sh
   echo secret
   ```

Note

The script must be secured by the operating system file permissions so that it is only accessible by, and can only be invoked by the API Gateway. The command should write the password to standard output.

Promotion between environments

When you promote and deploy passphrase protected configuration between environments (for example, from testing to production), the passphrase for the target configuration (production) must be the same as the passphrase in the source configuration (testing).

If you are using a different passphrase in each environment, before the deployment takes place, you must make a copy of the configuration (for example, `.fed` file), load it in Policy Studio, and set it with the correct passphrase for the target configuration. For more details, see the section called “Configure the passphrase in Policy Studio”.

For general details on promoting configuration between environments, see the *API Gateway Deployment and Promotion Guide*.

Further information

For more details on supported security features, see the *API Gateway Security Guide*.

Manage certificates and keys

Overview

For API Gateway to trust X.509 certificates issued by a specific Certificate Authority (CA), you must import that CA's certificate into the API Gateway's trusted certificate store. For example, if API Gateway is to trust secure communications (SSL connections or XML Signature) from an external SAML Policy Decision Point (PDP), you must import the PDP's certificate, or the issuing CA's certificate into the API Gateway certificate store.

In addition to importing CA certificates, you can import and create server certificates and private keys in the certificate store. These can be stored locally or on an external Hardware Security Module (HSM). You can also import and create public-private key pairs. For example, these can be used with the Secure Shell (SSH) File Transfer Protocol (SFTP) or with Pretty Good Privacy (PGP).

View certificates and keys

To view the certificates and keys stored in the certificate store, select **Certificates and Keys** > **Certificates** in the Policy Studio tree. Certificates and keys are listed on the following tabs in the **Certificates** window:

- **Certificates with Keys**: Server certificates with associated private keys
- **Certificates**: Server certificates without any associated private keys
- **CA**: Certificate Authority certificates with associated public keys

You can search for a specific certificate or key by entering a search string in the text box at the top of each tab, which automatically filters the tree.

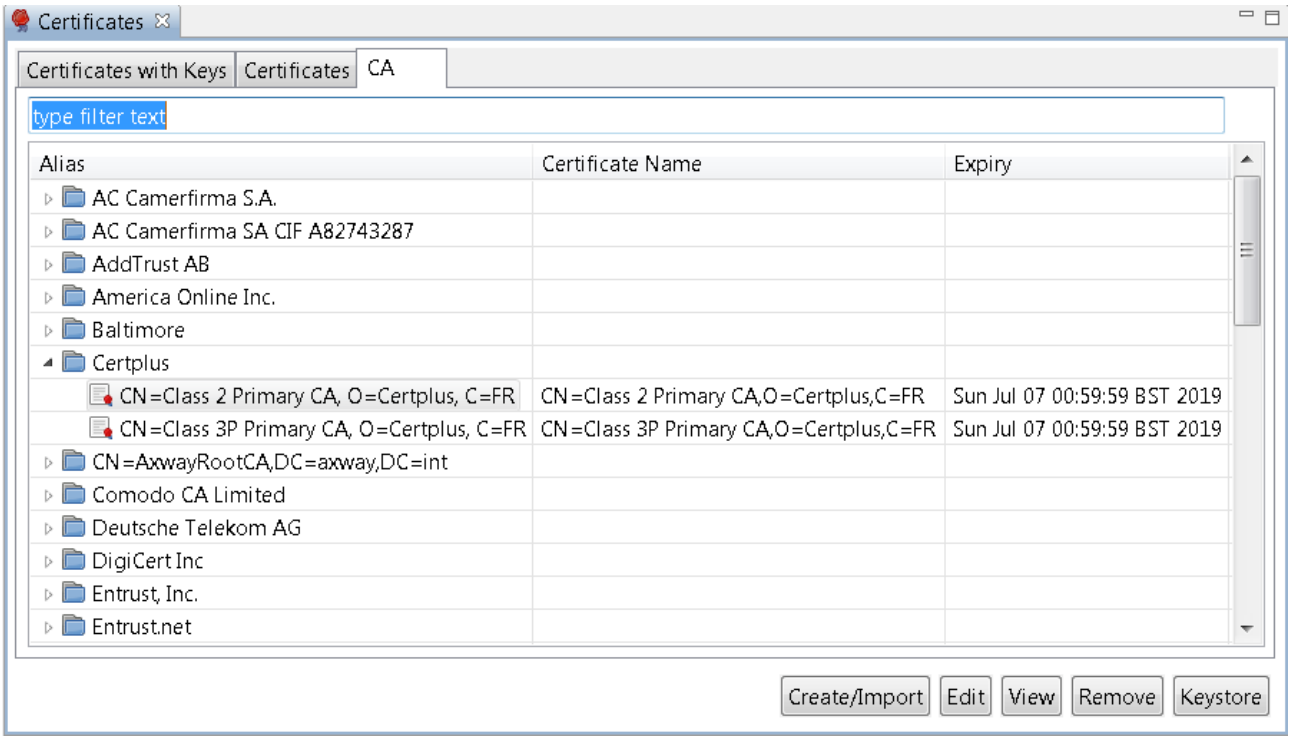

Certificate management options

The following options are available at the bottom right of the window:

- **Create/Import**: Click to create or import a new certificate and private key. For details, see the section called "Configure an X.509 certificate".
- **Edit**: Select a certificate, and click to edit its existing settings.
- **View**: Select a certificate, and click to view more detailed information.
- **Remove**: Select a certificate, and click to remove the certificate from the certificate store.
- **Keystore**: Click this to export or import certificates to or from a Java keystore. For details, see the section called "Manage certificates in Java keystores".

Configure an X.509 certificate

To create a certificate and private key, click **Create/Import**. The **Configure Certificate and Private Key** dialog is displayed. This section explains how to use the **X.509 Certificate** tab on this dialog.

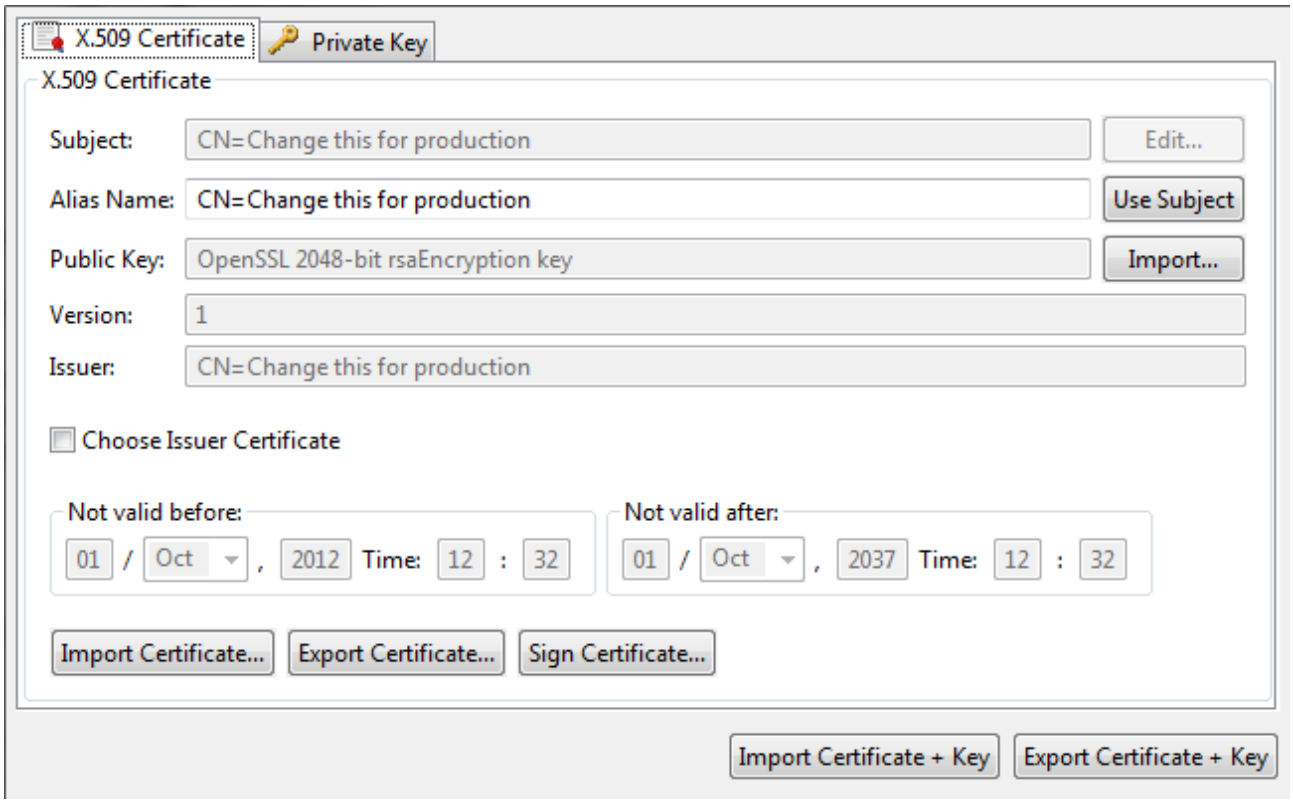

Create a certificate

Configure the following settings to create a certificate:

- **Subject**:
  Click **Edit** to configure the *Distinguished Name* (DName) of the subject.
- **Alias Name**:
  This mandatory field enables you to specify a friendly name (or alias) for the certificate. Alternatively, you can click

Use Subject to add the DName of the certificate in the text box instead of a certificate alias.

- **Public Key**:
  Click **Import** to import the subject's public key (usually from a PEM or DER-encoded file).
- **Version**:
  This read-only field displays the X.509 version of the certificate.
- **Issuer**:
  This read-only field displays the distinguished name of the CA that issued the certificate.
- **Choose Issuer Certificate**:
  Select to explicitly specify an issuer certificate for this certificate (for example, to avoid a potential clash or expiry issue with another certificate using the same intermediary certificate). You can then click the browse button on the right to select an issuer certificate. This setting is not selected by default.
- **Not valid before**:
  Select a date to define the start of the validity period of the certificate.
- **Not valid after**:
  Select a date to define the end of the validity period of the certificate.
- **Sign Certificate**:
  You must click this button to sign the certificate. The certificate can be self-signed, or signed by the private key belonging to a trusted CA whose key pair is stored in the certificate store.

Import certificates

You can use the following buttons to import or export certificates into the certificate store:

- **Import Certificate**:
  Click to import a certificate (for example, from a `.pem` or `.der` file).
- **Export Certificate**:
  Click to export the certificate (for example, to a `.pem` or `.der` file).

Configure a private key

Use the **Private Key** tab to configure details of the private key. By default, private keys are stored locally (for example, in the API Gateway certificate store). They can also be provided by an OpenSSL engine, or stored on a Hardware Security Module (HSM) if required.

API Gateway supports PKCS#11-compatible HSM devices. For example, this includes Thales nShield Solo , SafeNet Luna SA, and so on.

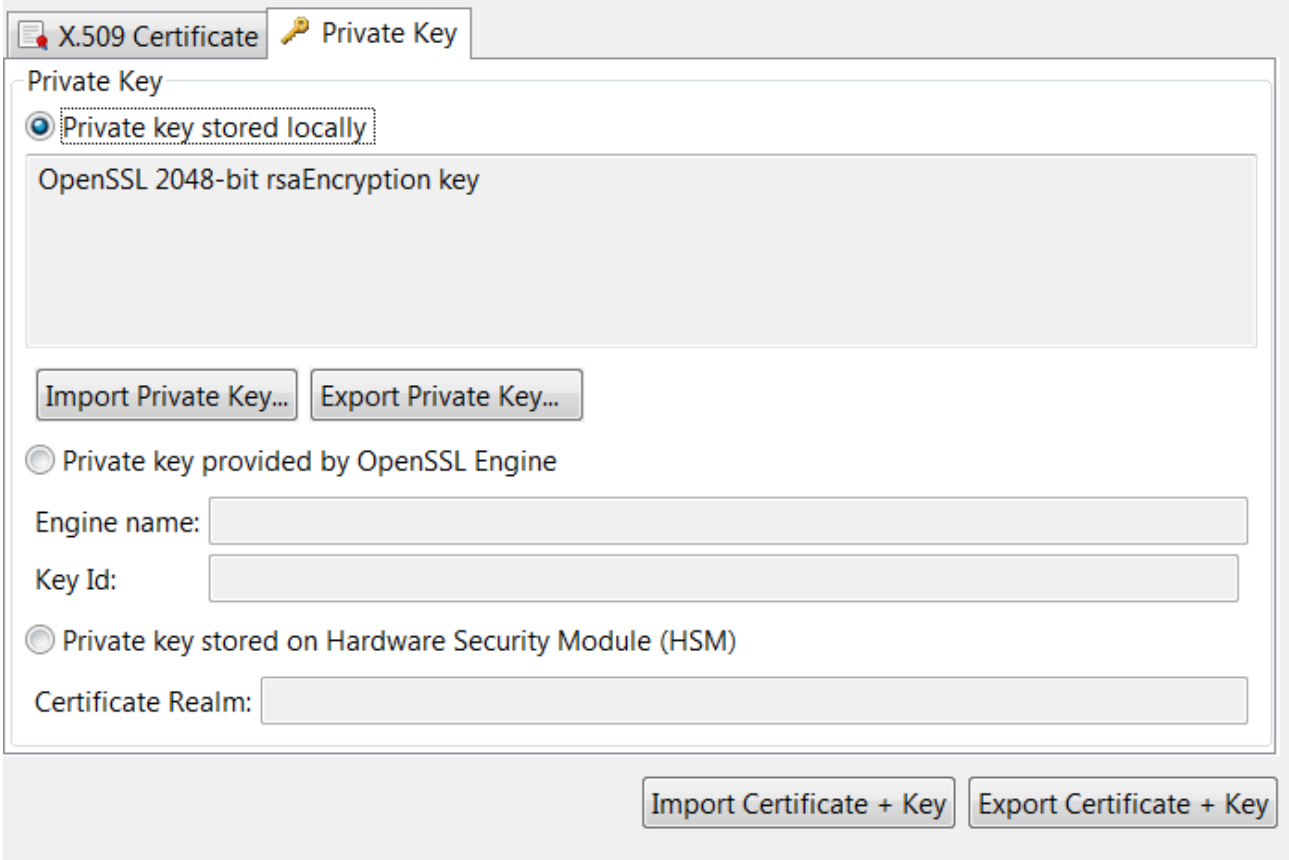

Private key stored locally

If the private key is stored in the API Gateway certificate store, , select **Private key stored locally**. The following options are available for keys stored locally:

- **Private key stored locally**:
  This read-only field displays details of the private key.
- **Import Private Key**:
  Click to import the subject's private key (usually from a PEM or DER-encoded file).
- **Export Private Key**:
  Click to export the subject's private key to a PEM or DER-encoded file.

Private key provided by OpenSSL engine

If the private key that corresponds to the public key in the certificate is provided by an OpenSSL engine, select **Private key provided by OpenSSL Engine**.

Configure the following fields to associate a key provided by the OpenSSL engine with the current certificate:

- **Engine name**:
  Enter the name of the OpenSSL engine to use to interface to an HSM. All vendor implementations of the OpenSSL Engine API are identified by a unique name. See your vendor's OpenSSL engine implementation or HSM documentation to find out the name of the engine.

- **Key Id**:
  Enter the key ID used to uniquely identify a specific private key from all others stored on an HSM. When you complete this dialog, the private key is associated with the certificate that you are currently editing. Private keys are identified by their key ID by default.

Private key stored on external HSM

If the private key that corresponds to the public key stored in the certificate resides on an external HSM, select **Private key stored on Hardware Security Module (HSM)**, and enter the name of the **Certificate Realm**.

Note

To use the API Gateway's PKCS#11 engine to access objects in an external HSM, the corresponding HSM provider and certificate realms must also be configured. For more details, see the section called “Configure HSMs and certificate realms”.

Configure HSMs and certificate realms

Certificate realms are abstractions of private keys and public key certificates, which mean that policy developers do not need to enter HSM-specific configuration such as slots and key labels. Instead, if a private key exists on an HSM, the developer can configure the certificate to show that its private key uses a specific certificate realm, which is simply an alias for a private key (for example, `JMS Keys`).

For example, on the host machine, an administrator could configure the `JMS Keys` certificate realm, and create a *keystore* for the realm, which requires specific knowledge about the HSM (for example, PIN, slot, and private key label). The certificate realm is the alias name, while the keystore is the actual private keystore for the realm.

Manage HSMs with keystoreadmin

The `keystoreadmin` script enables you to perform the following tasks:

- Register an HSM provider
- List registered HSM providers
- Create a certificate realm
- List certificate realms

For example, if a policy developer is using JMS, and wants to indicate that private keys exist on an HSM, they could indicate that the certificate is using the `JMS Keys` certificate realm. On each instance using the configuration, it is the responsibility of the administrator to create the `JMS Keys` certificate realm.

For more details, enter `keystoreadmin` in the following directory, and perform the instructions at the command prompt:

Windows	`INSTALL_DIR\apigateway\Win32\bin`
UNIX/Linux	`INSTALL_DIR/apigateway/posix/bin`

Use keystoreadmin in interactive mode

When you enter `keystoreadmin` without arguments, this displays an interactive menu with the following options:

Option	Description	When to use
`1`	`Change group or instance`	When registering HSMs or configuring certificate realms,

Option	Description	When to use
		you must choose the local group and instance to configure.
2	`List registered HSM providers`	Display the HSMs that are currently registered.
3	`Register an HSM provider`	Before creating certificate realms, you must first register the HSM. This option guides you through the steps. The HSM must be installed, configured, and active, and you must know the full path to the HSM device driver (PKCS#11). You give the HSM an alias (for example, LunaSA), which you use later when registering certificate realms.
4	`List Certificate Realms`	List configured certificate realms and associated keystores.
5	`Create a Certificate Realm`	Create a keystore and assign it to a certificate realm.

Step 1—Register an HSM provider

You must first register an HSM provider as follows:

1. Open a command prompt in the API Gateway `bin` directory (for example, `apigateway\Win32\bin`).
2. Enter the `keystoreadmin` command.
3. Select option `3) Register an HSM provider`.
4. If prompted, select the appropriate API Gateway group or instance.
5. You are prompted for a provider alias name. The alias is local only. For example, if registering a LunaSA HSM, you might enter the `LunaSA` alias.
6. For convenience, `keystoreadmin` searches for supported HSM drivers. If found, it shows the list of supported drivers. If none are found, this does not mean the driver does not exist. You must see your HSM documentation for the location of the drivers. For example:

```
Choose from one of the following:

   1) C:\LunaSA\cryptoki.dll

   o) Other
   q) Quit
```

7. If successful, `keystoreadmin` loads the driver and displays its details. For example:

```
Registering HSM provider...
 Initializing HSM...
 Crypto Version: 2.20
 Manufacter Id: SafeNet, Inc.
 Library Description: Chrystoki
 Library Version: 5.1
 Device registered.
```

Step 2—Create a certificate realm and associated keystore

To create a certificate realm and associated keystore, perform the following steps:

1. Open a command prompt in the API Gateway `bin` directory (for example, `apigateway\Win32\bin`).

2. Enter the `keystoreadmin` command.
3. Select option `5) Create a Certificate Realm`.
4. You are prompted to enter a certificate realm name. This certificate realm name will be used in Policy Studio when configuring the private key of the corresponding X.509 certificate. The realm name is case sensitive (for example, `JMS Keys`).
5. The registered HSMs are listed. For example, select option `1) HSM`.
6. The command connects to the selected HSM, and a list of available slots is displayed. Select the slot containing the private key to use for the certificate realm (for example, select slot `1`).
7. You are prompted to input the PIN passphrase for the slot. The passphrase will not echo any output.
8. When you enter the correct PIN passphrase for the slot, this displays a list of private keys. Choose the key to use for the certificate realm. For example:

```
Choose from one of the following:

   1) server1_priv
   2) jms_priv
   q) Quit

Select option: 2
```

9. You are prompted for a file name for the keystore. For example:

```
Certificate realm filename [jms keys.ks]:
Successfully created the certificate realm: JMS Keys
Press any key to continue...
```

10. The keystore is output to the API Gateway instance directory. For example:

```
apigateway/groups/group-2/instance-1/conf/certrealms/jms keys.ks
```

Note

Each API Gateway instance must have its certificate realm configured. When finished creating certificate realms, you must restart the API Gateway instance for the changes to take effect.

Step 3—Start the API Gateway when using an HSM

When the API Gateway is configured to use certificate realms, these realms are initialized on startup, and a connection to the corresponding HSM is established. This requires the PIN passphrase for the specific HSM slots. At startup, you can manually enter the required HSM slot PIN passphrase, or you can automate this instead.

Start API Gateway with manually entered PIN passphrase

When the API Gateway is configured to use an HSM, the API Gateway stops all processing, prompts for the HSM slot PIN passphrase, and waits indefinitely for input. For example:

```
INFO    07/Jan/2015:16:31:54 Initializing certificate realm 'JMS Keys'...
Enter passphrase for Certificate Realm, "JMS Keys":
```

The API Gateway does not reprompt if the PIN passphrase is incorrect. It logs the error and continues, while any services that use the certificate realm cannot use the HSM.

Start API Gateway with automatic PIN passphrase

You can configure the API Gateway to start and initialize the HSM by invoking a command script on the operating system to obtain the HSM slot PIN passphrase. This enables the API Gateway for automatic startup without manually entering

the PIN passphrase.

To configure an automatic PIN passphrase, perform the following steps:

1. Edit the API Gateway instance's `vpkcs11.xml` configuration file. For example:

```
apigateway/groups/group-2/instance-1/conf/vpkcs11.xml
```

2. Add a `PASSPHRASE_EXEC` command that contains the full path to the script that executes and obtains the passphrase. The script should write the passphrase to stdout, and should have the necessary operating system file and execute protection settings to prevent unauthorized access to the PIN passphrase. The following example shows a `vpkcs11.xml` file that invokes the `hsmpin.sh` to echo the passphrase:

```
<?xml version="1.0" encoding="utf-8"?>
<ConfigurationFragment provider="cryptov">
    <Engine name="vpkcs11" defaultFor="">
        <EngineCommand when="preInit" name="REALMS_DIR"
            value="$VINSTDIR/conf/certrealms" />
        <EngineCommand when="preInit" name="PASSPHRASE_EXEC"
            value=""$VDISTDIR/hsmpin.sh"" />
    </Engine>
</ConfigurationFragment>
```

3. The API Gateway provides the certificate realm as an argument to the script, so you can use the same script to initialize multiple realms. The following examples show scripts that write a PIN of `1234` to stdout when initializing the `JMS Keys` certificate realm:

Example hsmpin.bat file on Windows

```
@echo off
IF [%1]==[] GOTO _END

:: Strip out the double quotes around arg
SET REALM=%1
SET REALM=%REALM:"=%

IF "%REALM%"=="JMS Keys" ECHO 1234
```

Example hsmpin.sh file on Linux/UNIX

```
#!/bin/sh
case $1 in
"JMS Keys")
    echo 1234
    ;;
esac
```

Configure SSH key pairs

To configure public-private key pairs in the certificate store, select **Certificates and Keys** > **Key Pairs**. The **Key Pairs** window enables you to add, edit, or delete OpenSSH public-private key pairs, which are required for the Secure Shell (SSH) File Transfer Protocol (SFTP).

Add a key pair

To add a public-private key pair, click **Add** on the right, and configure the following settings in the dialog:

- **Alias**:

Enter a unique name for the key pair.

- **Algorithm**:
  Enter the algorithm used to generate the key pair. Defaults to RSA.
- **Load**:
  Click to select the public key or private key files to use. The **Fingerprint** field is auto-populated when you load a public key.

Note

The keys must be OpenSSH keys. RSA keys are supported, but DSA keys are not supported. The keys must not be passphrase protected.

Manage OpenSSH keys

You can use the `ssh-keygen` command provided on UNIX to manage OpenSSH keys. For example:

- The following command creates an OpenSSH key:
  `ssh-keygen -t rsa`
- The following command converts an `ssh.com` key to an OpenSSH key:
  `ssh-keygen -i -f ssh.com.key > open.ssh.key`
- The following command removes a passphrase (enter the old passphrase, and enter nothing for the new passphrase):
  `ssh-keygen -p`
- The following command outputs the key fingerprint:
  `ssh-keygen -lf ssh_host_rsa_key.pub`

Edit a key pair
To edit a public-private key pair, select a key pair alias in the table, and click **Edit** on the right. For example, you can load a different public key and private key. Alternatively, double-click a key pair alias in the table to edit it.

Delete key pairs
You can delete a selected key pair from the certificate store by clicking **Remove** on the right. Alternatively, click **Remove All**.

Configure PGP key pairs

To configure Pretty Good Privacy (PGP) key pairs in the certificate store, select **Certificates and Keys** > **PGP Key Pairs**. The **PGP Key Pairs** window enables you to add, edit, or delete PGP public-private key pairs.

Add a PGP key pair

To add a PGP public-private key pair, click the **Add** on the right, and configure the following settings in the dialog:

- **Alias**:
  Enter a unique name for the PGP key pair.
- **Load**:
  Click **Load** to select the public key and private key files to use.

Note

The PGP keys added must not be passphrase protected.

Manage PGP keys

You can use the freely available GNU Privacy Guard (GnuPG) tool to manage PGP key files (available from http://www.gnupg.org/). For example:

- The following command creates a PGP key:
  `gpg --gen-key`
  For more details, see http://www.seas.upenn.edu/cets/answers/pgp_keys.html [http://www.seas.upenn.edu/cets/answers/pgp_keys.html]
- The following command enables you to view the PGP key:
  `gpg -a --export`
- The following command exports a public key to a file:
  `gpg --export -u 'UserName' -a -o public.key`
- The following command exports a private key to a file:
  `gpg --export-secret-keys -u 'UserName' -a -o private.key`
- The following command lists the private keys:
  `gpg --list-secret-keys`

Edit a PGP key pair
To edit a PGP key pair, select a key pair alias in the table, and click **Edit** on the right. For example, you can load a different public key and private key. Alternatively, double-click a key pair alias in the table to edit it.

Delete PGP key pairs
You can delete a selected PGP key pair from the certificate store by clicking **Remove** on the right. Alternatively, click **Remove All**.

Global import and export options

This section desribes global import and export options available when managing certificates and keys.

Import and export certificates and keys

The following global configuration options apply to both the **X.509 Certificate** and **Private Key** tabs:

- **Import Certificate + Key**:
  Use this option to import a certificate and a key (for example, from a `.p12` file).
- **Export Certificate + Key**:
  Use this option to export a certificate and a key (for example, to a `.p12` file).

Click **OK** when you have finished configuring the certificate and private key.

Manage certificates in Java keystores

You can also export a certificate to a Java keystore. You can do this by clicking **Keystore** on the main **Certificates** window. Click the browse button at beside the **Keystore** field at the top right to open an existing keystore, or click **New Keystore** to create a new keystore. Choose the name and location of the keystore file, and enter a passphrase for this keystore when prompted. Click **Export to Keystore**, and select a certificate to export.

Similarly, you can import certificates and keys from a Java keystore into the certificate store. To do this, click **Keystore** on the main **Certificates** window. On the **Keystore** window, browse to the location of the keystore by clicking the browse button beside the **Keystore** field. The certificates/keys in the keystore are listed in the table. To import any of these keys to the certificate store, select the box next to the certificate or key to import, and click **Import to Trusted certificate store**. If the key is protected by a password, you are prompted for this password.

You can also use the **Keystore** window to view and remove existing entries in the keystore. You can also add keys to

the keystore and to create a new keystore. Use the appropriate button to perform any of these tasks.

Further information

For more details on supported security features, see the *API Gateway Security Guide*.

Manage API firewalling

Overview

API Gateway provides API firewalling capabilities by embedding ModSecurity, which is a popular open-source Web Application Firewall (WAF) engine. API Gateway administrators can configure the embedded ModSecurity engine to protect API Gateway HTTP traffic against threats and monitor reported exceptions.

This chapter explains how to enable API firewalling on an API Gateway interface in Policy Studio, and how to monitor it in the API Gateway Manager web console.

For more details on ModSecurity, see www.modsecurity.org.

Configure API firewalling

This section explains how to configure API firewalling by enabling threat protection and configuring threat protection rules.

Enable threat protection

The embedded ModSecurity engine is not enabled by default. To enable ModSecurity on an API Gateway interface, perform the following steps:

1. In the Policy Studio tree on the left, select the interface port that you wish to enable (for example, under **Listeners** > **API Gateway** > **Default Services** > **Ports**).
2. Right-click the HTTP or HTTPS interface in the window on the right, and select **Edit**.
3. Click the **Advanced** tab, and select **Protect this interface with Threat Protection rules**.

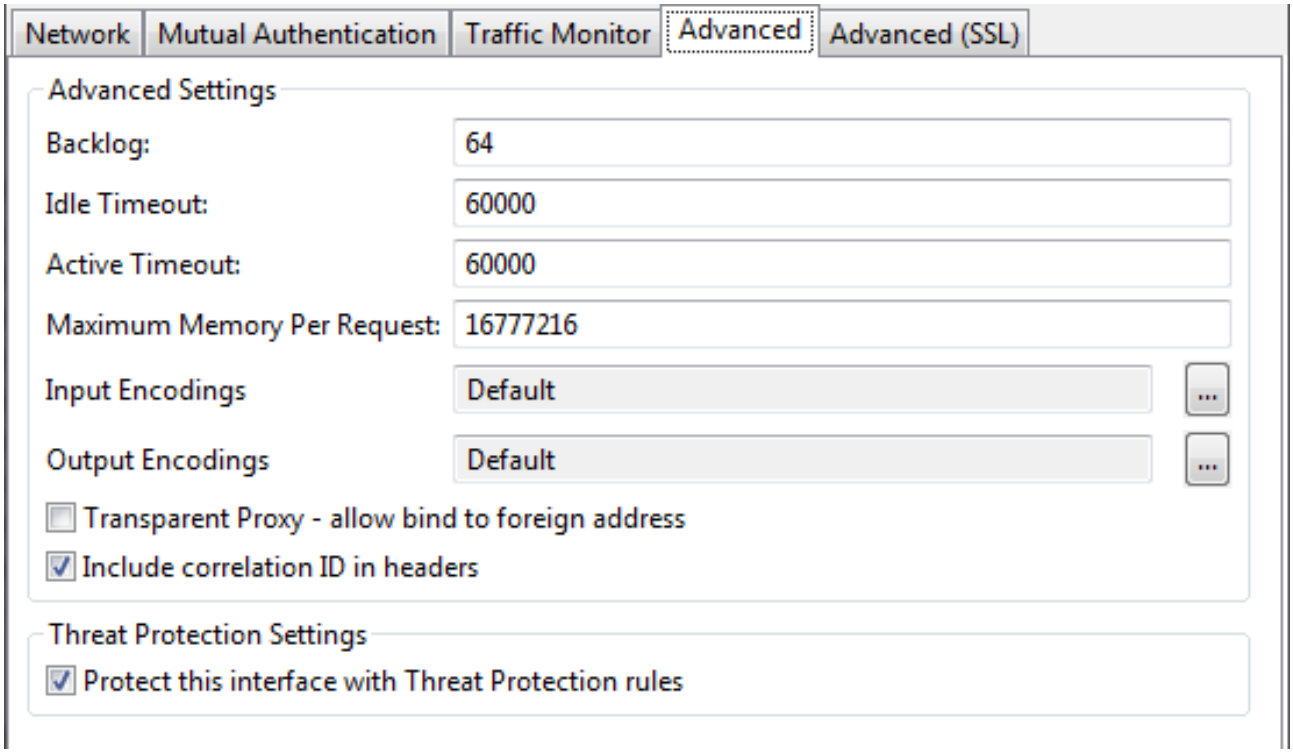

Configure threat protection rules

ModSecurity provides very little protection on its own. However, you can configure the required protection by configuring its rules engine. Protecting against specific threats requires specific rules, and different vendors provide rules for specific threat protection capabilities. The Open Web Application Security Project (OWASP) ModSecurity Core Rule Set (CRS) project provides a popular rule set.

When threat protection has been enabled, the embedded ModSecurity engine looks for threat protection rules configuration in the following directory:

```
$VDISTDIR/system/conf/threat-protection/default
```

The API Gateway uses the OWASP ModSecurity CRS folder structure, and expects to find the following inside the `threat-protection/default` directory:

- `modsecurity.conf` file: Contains the engine global settings. For details on the file format and recommended settings, see the "Recommended Base Configuration" in the ModSecurity documentation.
- `activated_rules` folder: Containing the threat protection rules files. All files with a `.conf` extension in this folder are loaded by the embedded ModSecurity engine.

Note

The recommended ModSecurity default configuration sets the engine mode to `SecRuleEngine DetectionOnly`, which applies the activated rules. But this does not interfere with the traffic (does not block any requests).

You must restart the API Gateway instance after changing any ModSecurity configuration.

Monitor API firewalling

An API Gateway administrator or operator can use the **Traffic** > **HTTP** tab in the API Gateway Manager web console to monitor API firewalling. You can use this tab to show how threat protection affects the HTTP traffic served by the API Gateway.

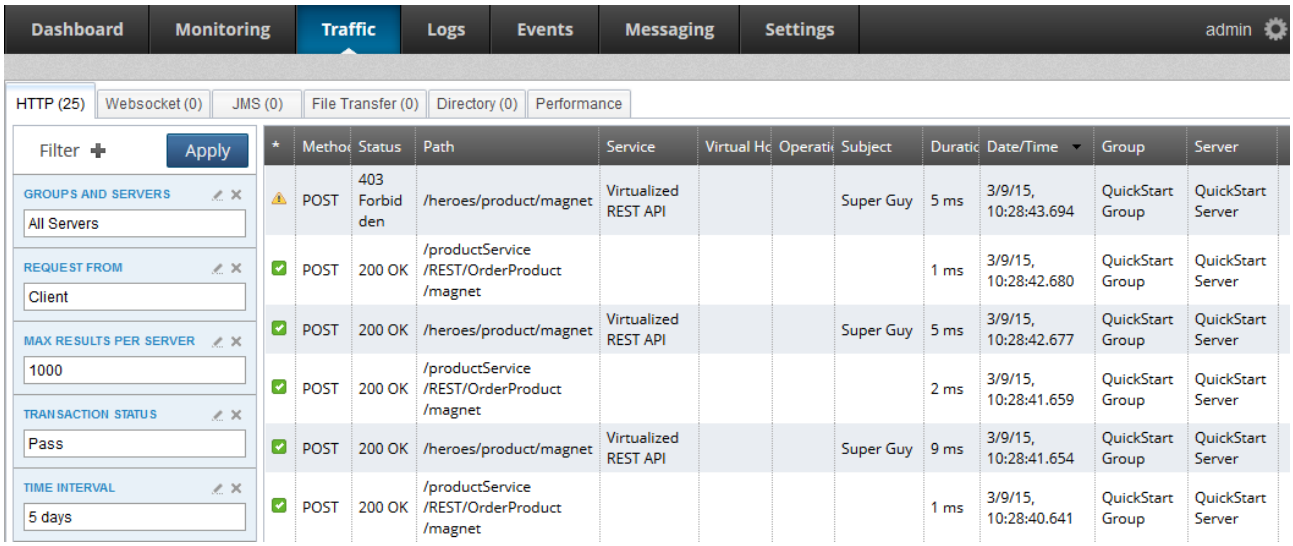

You can filter this tab to display by **Threat Protection** to quickly locate all passed or failed transactions.

For example, perform the following steps:

1. Click **Filter +** in the search pane on the left.
2. Select **Threat Protection** in the list.
3. Select a threat protection status in the dialog:
   - **Pass**: The ModSecurity engine marks all transactions that pass its rules with this status.
   - **Fail**: Transactions that violate any active ModSecurity engine rules are marked with this status. These transactions should be monitored because they represent a false positive (the protection rules might need to be adjusted), or malicious client traffic. You can view more details about the failure reason and specific rule violation by drilling down a specific transaction and looking at the trace details.
   - **Exception**: Transactions that cause a rule processing or other unknown error are marked with this status. These should not occur and probably indicate some rule configuration problem.
4. Click **Apply**.

For example, the following shows detailed trace output from drilling down a failed transaction:

```
Message: Access denied with code 403 (phase 2). Pattern match
"(?i:(?:\\b(?:(?:s(?:ys\\.(?:user_(?:(?:t(?:ab(?:_column|le)|rigger)|object|view)
s|c(?:onstraints|atalog))|all_tables|tab)|elect\\b.{0,40}\\b(?:substring|users?|ascii))
|m(?:sys(?:(?:queri|ac)e|relationship|column|object)s|ysql\\.(db|user))|c(?:onstraint ..."
at ARGS:q. [file "C:\Axway\APIGateway\apigateway\system\conf\threat-protection\default\
activated_rules\modsecurity_crs_41_sql_injection_attacks.conf"] [line "116"] [id "950007"]
[rev "2"] [msg "Blind SQL Injection Attack"] [data "Matched Data: SELECT * FROM USERS
found within ARGS:q: SELECT * FROM USERS"] [severity "CRITICAL"] [ver "OWASP_CRS/2.2.9"]
[maturity "9"] [accuracy "8"] [tag "OWASP_CRS/WEB_ATTACK/SQL_INJECTION"]
[tag "WASCTC/WASC-19"] [tag "OWASP_TOP_10/A1"] [tag "OWASP_AppSensor/CIE1"] [tag "PCI/6.5.2"]
```

Further information

For more details on OWASP, see https://www.owasp.org.

For more details on supported security features, see the *API Gateway Security Guide*.

Manage API Gateway deployments

Overview

When connected to the Admin Node Manager server in Policy Studio, you can deploy configurations to API Gateway instances running in groups in a domain. In Policy Studio, the **Group / API Gateway** topology view enables you to edit the configuration of currently running API Gateway instances. You can update the downloaded configuration, and deploy it to the server, where it can be reloaded later. You can deploy modified configuration to multiple API Gateway instances managed by Policy Studio. You can also create groups and API Gateway instances.

The web-based API Gateway Manager enables you to deploy configurations to API Gateway instances running in groups in a domain, to create groups and API Gateway instances, and to manage Admin users. In this way, Policy Studio and the API Gateway Manager enable policy developers and administrators to centrally manage the policies that are enforced at all nodes throughout the network.

In addition, Policy Studio enables you to compare and merge differences between versions of the same policy. Policies can be merged, and deployed to any running instance that is managed by Policy Studio. One of the most powerful uses of this centralized management capability is in transitioning from a staging environment to a production environment. For example, policies can be developed and tested on the staging environment, and when ready, they can be deployed to all instances deployed in the production environment.

Connect to a server in Policy Studio

Before starting Policy Studio, you should first ensure that the Admin Node Manager and the server instance that you wish to connect to have been started.

When Policy Studio starts up, click a link to a server to display the **Open Connection** dialog. You can use this dialog to specify **Connection Details** (for example, host, port, user name, and password) or to specify **Saved Sessions**. If you wish to connect to the server using a non-default URL, click **Advanced**, and enter the **URL**. The default Admin Node Manager URL is:

```
https://localhost:8090/api
```

Alternatively, you can connect to a server configuration file by clicking the **Open File** button. For more details on connecting using a server URL, configuration file, or deployment archive, see the *API Gateway Policy Developer Guide*.

Note

You must connect to the Admin Node Manager server to deploy API Gateway configuration or manage multiple API Gateway instances in your network.

When the connection to the server has been made, the **Group / API Gateway** topology view is displayed. This displays the list of server instances currently managed by the Admin Node Manager in Policy Studio, and enables you to manage the configuration for server instances.

Edit a server configuration in Policy Studio

The **Group / API Gateway** topology view lists all available instances in each group. Double-click an instance name in the list to load its active configuration. Alternatively, right-click an instance name, and select **Edit Configuration**. The active server configuration is loaded and displayed in the following format: **InstanceName [HostName:Port]** (for example, **test_server [roadrunner.acme.com:8085]**).

When an active server configuration is loaded, its services are displayed under the **Listeners** node in the Policy Studio tree on the left. Expand one of the top-level nodes in the tree to display additional details (for example, **Business Ser-**

vices, **External Connections**, **Resources**, **Libraries**, or **Settings**).

When editing an active server configuration, you can deploy updates using the **Deploy** button in the toolbar (alternatively, press **F6**). You can also deploy configuration packages in the **Group / API Gateway** topology view. For more details, see *Deploy API Gateway configuration*.

Manage deployments in API Gateway Manager

In the web-based API Gateway Manager tool, the **TOPOLOGY** section on the **Dashboard** tab enables you to create groups and API Gateway instances, and to deploy configuration. For details on how to access the API Gateway Manager, see *Start the API Gateway tools*.

Compare and merge configurations in Policy Studio

You can compare and merge differences between the currently loaded API Gateway configuration with a configuration stored in a deployment package (.fed file). Click the **Compare** button on the Policy Studio toolbar to select a `.fed` file to compare the current configuration against. The results are displayed in the **Compare/Merge** tab.

For example, you can view the differences made to particular fields in an Authentication filter that occurs in both configurations. When a difference is located, you can merge the differences, and thereby update the fields in the Authentication filter in the current configuration with the field values for the same Authentication filter in the deployment package.

For more details, see the *API Gateway Policy Developer Guide*.

Manage Admin users in API Gateway Manager

You can add new Admin Users to enable role-based access to the API Gateway configuration managed by Policy Studio and API Gateway Manager. The default `admin` user has access to all API Gateway features in Policy Studio and API Gateway Manager, and can view and modify all API Gateway configurations.

To add or remove Admin Users, click the **Settings** > **Admin Users** tab in the API Gateway Manager. For more details, see *Manage Admin users*.

For more details on role-based access, see *Configure Role-Based Access Control (RBAC)*.

Configure policies in Policy Studio

You can use Policy Studio to manage the configuration of your policies, which can then be deployed to running instances of Oracle API Gateways.

For details on configuring the full range of message filters (for example, for Authentication, Authorization, or Content Filtering), see the *API Gateway Policy Developer Guide*.

Deploy API Gateway configuration

Overview

You can edit API Gateway configuration offline, and then deploy later to a specified API Gateway instance using API Gateway configuration packages. A *deployment package* is a `.fed` file that contains all API Gateway configuration. This includes policies, listeners, external connections, users, certificates, and environment settings. A *policy package* is a `.pol` file that contains policies, listeners, external connections, and environment settings. While an *environment package* is an `.env` file that contains users, certificates, and environment settings. The content of the `.fed` file is equivalent to the combined contents of the `.pol` and `.env` files.

A *package property* is a name-value pair that applies to a specific configuration package (`.fed`, `.pol`, or `.env`). Specifying a property associates metadata with the configuration in that package. For example, the **Name** property with a value of `Default Factory Configuration` is associated with a default installation. For more details on configuration packages and properties, see the *API Gateway Deployment and Promotion Guide*.

You can use Policy Studio to create packages (`.fed`, `.pol`, or `.env`). You can also use Policy Studio to deploy an existing package or factory configuration on selected API Gateway instances, or to deploy a currently loaded configuration. You can use the API Gateway Manager console to deploy a package in a browser. Alternatively, you can use the `managedomain` script to create and deploy deployment packages (`.fed` files) on the command line.

Create a package in Policy Studio

You can create an API Gateway configuration package for a currently loaded configuration, or for a selected server instance in the **Group / API Gateway** topology view.

Currently loaded configuration
To create a package (`.fed`, `.pol`, or `.env`) for a currently loaded API Gateway configuration, perform the following steps:

1. In the main menu, select **File** > **Save** followed by the appropriate option:
   - **Deployment Package**
   - **Policy Package**
   - **Environment Package**
2. Enter a filename, and click **Save**.

Group / API Gateway view
To create a deployment package (`.fed`) in the API Gateway **Group / API Gateway** topology view, perform the following steps:

1. Right-click a server instance in the tree, and select **Save Deployment Package**.
2. Browse to a directory in the dialog.
3. Click **OK**. The file is saved to the specified directory (for example, `c:\temp\5c3b2a3c-23a5-4261-87cb-eca150f0a037.fed`.
4. Click **OK**.

Configure package properties in Policy Studio

You can view or modify API Gateway configuration package properties for a currently loaded configuration, or for a selected server instance in the **Group / API Gateway** topology view.

Currently loaded configuration
To view and modify configuration properties for a currently loaded API Gateway configuration, perform the following

steps:

1. In the Policy Studio tree, and select **Package Properties** > **Policy** or **Environment**.
2. Enter values for the appropriate configuration properties (for example, **Name**, **Description**, or **Version**).
3. If you wish to create any additional properties (for example, **Department**), click the green (+) button on the right, and enter a property value (for example, `Engineering`).
4. If you wish to remove a property, click the red (x) button on the right of the property.
5. Click **Save** on the top right of the screen.

Group / API Gateway view
To view and modify configuration properties for a selected server instance in the API Gateway **Group / API Gateway** topology view, perform the following steps:

1. Right-click a server in the tree, and select **View/Modify Properties**.
2. Select the **Policy Properties** or **Environment Properties** tab.
3. Enter values for the appropriate configuration properties (**Name**, **Description**, and **Version**).
4. If you wish to create any additional properties (for example, **Department**), click the green (+) button on the right, and enter a property value (for example, `Engineering`).
5. If you wish to remove a property, click the red (x) button on the right of the property.
6. Click **Update Configuration Properties**.
7. Click **OK**.

For details on customizing the default package properties displayed, see the **Topology Screen** settings in the *API Gateway Policy Developer Guide*.

Deploy packages in Policy Studio

You can use the Policy Studio to deploy configuration packages to selected API Gateway instances in the **Group / API Gateway** topology view.

Deploy a deployment package
To deploy an existing deployment package (`.fed` file) in the **Group / API Gateway** view, perform the following steps:

1. Click the **Deploy** button in the toolbar.
2. In the **Select the servers(s) you wish to deploy to** section, select a server group from the **Group** list, and select the server instance(s) in the box below.
3. In the **Select the configuration you wish to deploy** section, select **I wish to deploy configuration contained in a single Deployment Package**.
4. In the **Deployment Package** field, click **Browse for .fed**, and select the `.fed` file.
5. Click **Deploy** to upload the package to the Admin Node Manager and deploy to the selected server(s).
6. When the package has deployed, click **Finish**.

Deploy policy and environment packages
To deploy an existing policy package (`.pol` file) and environment package (`.env` file) in the **Group / API Gateway** view, perform the following steps:

1. Click the **Deploy** button in the toolbar.
2. In the **Select the servers(s) you wish to deploy to** section, select a server group from the **Group** list, and select the server instance(s) in the box below.
3. In the **Select the configuration you wish to deploy** section, select **I wish to deploy configuration contained in Policy Package and Environment Package**.
4. In the **Policy Package** field, click **Browse for .pol**, and select the `.pol` file.

5. In the **Environment Package** field, click **Browse for .env**, and select the `.env` file.
6. Click **Deploy** to upload these packages to the Admin Node Manager and deploy to the selected server(s).
7. When the package has deployed, click **Finish**.

Deploy a factory configuration in Policy Studio

To deploy a default factory configuration in the **Group / API Gateway** view, perform the following steps:

1. Click the **Deploy** button in the toolbar.
2. In the **Select the servers(s) you wish to deploy to** section, select a server group from the **Group** list, and select the server instance(s) in the box below.
3. In the **Select the configuration you wish to deploy** section, select **I wish to deploy a factory configuration**.
4. Click **Deploy** to deploy the configuration to the selected server(s).

Deploy currently loaded configuration in Policy Studio

You can also deploy updates to a currently loaded configuration in Policy Studio when editing the configuration. To deploy a currently loaded configuration, perform the following steps:

1. Click the **Deploy** button on the right in the toolbar.
2. In the **Select the servers(s) you wish to deploy to** section, select a server group from the **Group** list, and select the server instance(s) in the box below.
3. Click **Deploy**, and wait for the deployment to complete.
4. Click **Finish**.

Push configuration to a group in Policy Studio

When there is more than one API Gateway instance in a group, and configuration becomes out of sync between instances, you can select which configuration to push to the group. Perform the following steps in Policy Studio:

1. In the **Group / API Gateway** topology view, right-click the API Gateway instance configuration that you want to deploy to other instances in the group.
2. Select **Push this API Gateway's configuration to the group**.
3. In the wizard, select the API Gateway instances in the group that you wish to deploy to.
4. Click **Deploy** to deploy to the selected instances in the group.
5. Click **Finish**.

View deployment results in Policy Studio

When you click **Deploy**, the **Deployment Results** screen is displayed, and deployment to each server occurs sequentially. Feedback is provided using icons in the **Task** column, and text in the **Status** column. When the configuration has deployed, click **Finish**.

Cancel deployments
You can cancel deployments by clicking the **Cancel** button. Feedback is provided in the **Status** column. You cannot cancel a deployment when it has started. The wizard performs the cancellation at the end of the current deployment, with all remaining deployments being cancelled.

Deployment errors
Client-side and server-side errors can occur. Client-side errors are displayed in the **System Trace** in the **Console** view. If any server-side deployment errors occur during the deployment process, you can review these in the **Deployment Er-**

ror Log view. This is displayed at the bottom of the screen when you click **Finish**, and lists any errors that occur for each instance. The corresponding Console **Deployment Log** is also available in the **Console** view.

Redeploy
When you have deployed a configuration to one or more instances, you can click back through the wizard to change your selections and redeploy, without needing to exit and relaunch the wizard.

Deploy on the command line

You can create and deploy a deployment package (`.fed`) using the `managedomain --menu` command in the following directory:

Windows	`INSTALL_DIR\Win32\bin`
UNIX/Linux	`INSTALL_DIR/posix/bin`

The deployment options for the `managedomain --menu` command are as follows:

```
18) Deploy to a group
  19) List deployment information
  20) Create deployment archive
  21) Download deployment archive
  22) Update deployment archive properties
```

Deploy packages in API Gateway Manager

You can also use the API Gateway Manager web console to deploy configuration packages to a group of API Gateway instances. This functionality is available on the default **Dashboard** tab. For more details, see *Manage domain topology in API Gateway Manager*.

Monitor services in API Gateway Manager

Overview

This topic explains how to monitor example services using the API Gateway Manager monitoring tools. For example, these include real-time monitoring metrics, message traffic monitoring, and performance statistics.

Ensure monitoring is enabled

Real-time monitoring and traffic monitoring are enabled by default. If necessary, you can enable these settings as follows:

1. In the Policy Studio tree, select the **Server Settings** node, and select **Monitoring** > **Traffic Monitor**.
2. In the **Traffic Monitor** settings, ensure **Enable Traffic Monitor** is selected.
3. Select **Monitoring** > **Real Time Monitoring**, and ensure **Enable Real Time Monitoring** is selected.
4. Click the **Deploy** icon in the Policy Studio toolbar to deploy these settings to the API Gateway. Alternatively, press F6.

Important

Enabling traffic monitoring may have a negative impact on performance. If you wish to maximize performance, you can disable these settings. For more details, see *Traffic monitoring settings*.

View real-time monitoring

You can view a wide range of monitoring data in the API Gateway Manager. For example, this includes message status, message traffic, filter execution path, message content, system, services, and remote hosts. You can view real-time traffic monitoring summary data on the main **Dashboard** tab in the **TRAFFIC** section. The following example shows the number of messages that have been passed by the API Gateway on to a service:

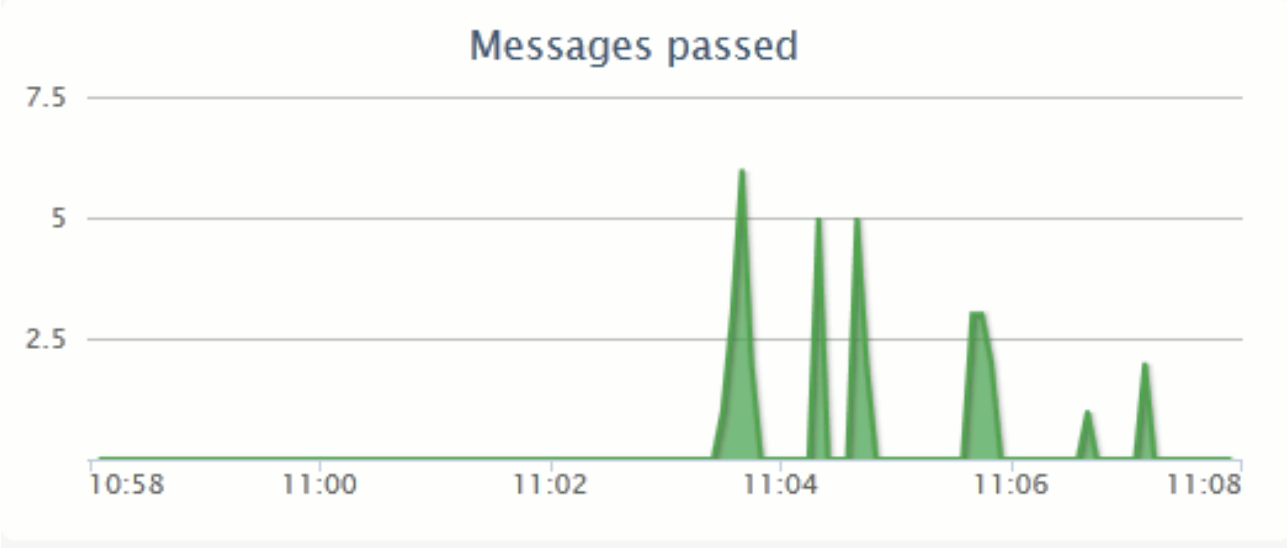

Each time you send test messages through the API Gateway to an example service (for example, using API Gateway Explorer or the Send Request (SR) tool), the message status is displayed in the **TRAFFIC** section.

View traffic monitoring

You can use the traffic monitoring tools in API Gateway Manager for operational diagnostics and root cause analysis. The **Traffic** view provides a web-based message log of the HTTP, HTTPS, JMS, and FTP traffic processed by the API Gateway. You can perform tasks such as the following:

- Filter messages on a range of criteria (for example, transaction ID, service name, or remote host)
- Drill down to view message contents
- View performance statistics (for example, number of requests, average bytes sent, or average duration)

For example, you can click the **Traffic** button in the API Gateway Manager to view summary information for each message sent to the API Gateway. Alternatively, you can click one of the summary charts displayed on the **Dashboard** (for example, **Messages passed** or **Messages failed**). This displays the message traffic automatically filtered according to your selection. The following simple example shows the details displayed on the **Traffic** tab for **Messages passed** by the API Gateway:

*	Methoc	Status	Path	Service	Virtual Hos	Operat	Subject	Duratio	Date/Time	Group	Server
☑	POST	200 OK	/productService /REST/OrderProduct /magnet					1 ms	1/30/15, 15:18:56.823	QuickStart Group	QuickStart Server
☑	POST	200 OK	/heroes/product/magnet	Virtualized REST API			Super Guy	5 ms	1/30/15, 15:18:56.821	QuickStart Group	QuickStart Server
☑	POST	200 OK	/productService /REST/OrderProduct /magnet					1 ms	1/30/15, 15:18:55.495	QuickStart Group	QuickStart Server
☑	POST	200 OK	/productService /REST/OrderProduct /magnet	Virtualized REST API			Super Guy	3 ms	1/30/15, 15:18:55.494	QuickStart Group	QuickStart Server
☑	POST	200 OK	/heroes/product/magnet	Virtualized REST API			Super Guy	7 ms	1/30/15, 15:18:55.491	QuickStart Group	QuickStart Server
☑	POST	200 OK	/productService /REST/OrderProduct /magnet					1 ms	1/30/15, 15:18:54.433	QuickStart Group	QuickStart Server

Filter message traffic

In the **SELECTION** pane on the left of the **Traffic** tab, you can click the **Apply** button to filter the messages displayed based on a range of criteria. For example, the default filters include **REQUEST FROM** (Client or API Gateway), **MAX RESULTS PER SERVER**, **TRANSACTION STATUS**, and **TIME INTERVAL**.

You can click **Add Filter** to search on different criteria (for example, Service Name, Remote Host, Authentication Subject, Transaction ID, and Operation). The API Gateway inserts a transaction ID in all HTTP and HTTPS traffic in a header named `X-CorrelationID`. When you have selected your search criteria, click the **Apply** button.

View message content

When you click a selected message listed on the **Traffic** tab, this displays the message filter execution path and the contents of each request message and response message. The following example displays the message path for a simple Google Search message:

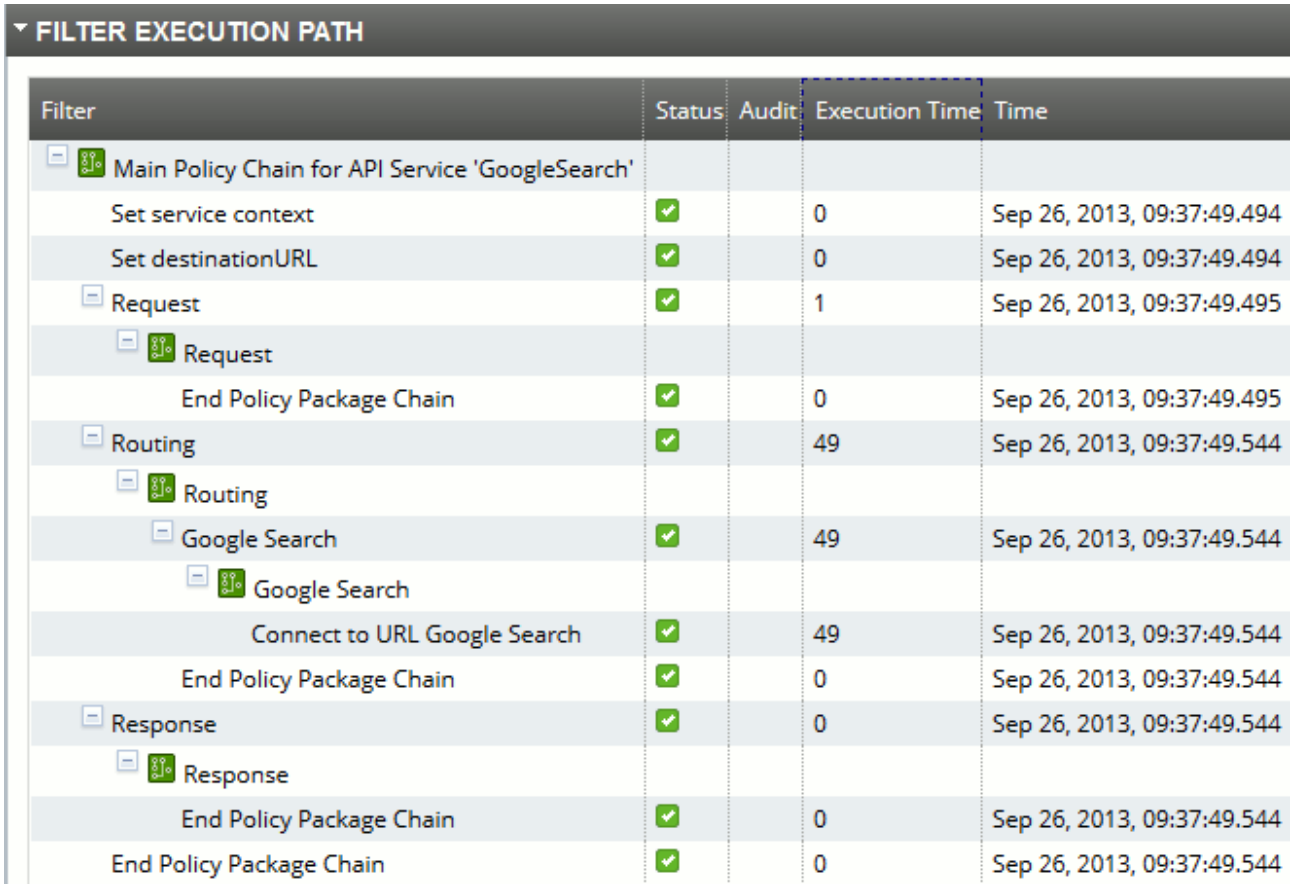

FILTER EXECUTION PATH

Filter	Status	Audit	Execution Time	Time
Main Policy Chain for API Service 'GoogleSearch'				
Set service context	✓		0	Sep 26, 2013, 09:37:49.494
Set destinationURL	✓		0	Sep 26, 2013, 09:37:49.494
Request	✓		1	Sep 26, 2013, 09:37:49.495
Request				
End Policy Package Chain	✓		0	Sep 26, 2013, 09:37:49.495
Routing	✓		49	Sep 26, 2013, 09:37:49.544
Routing				
Google Search	✓		49	Sep 26, 2013, 09:37:49.544
Google Search				
Connect to URL Google Search	✓		49	Sep 26, 2013, 09:37:49.544
End Policy Package Chain	✓		0	Sep 26, 2013, 09:37:49.544
Response	✓		0	Sep 26, 2013, 09:37:49.544
Response				
End Policy Package Chain	✓		0	Sep 26, 2013, 09:37:49.544
End Policy Package Chain	✓		0	Sep 26, 2013, 09:37:49.544

The following example shows the corresponding message content for the selected message displayed below:

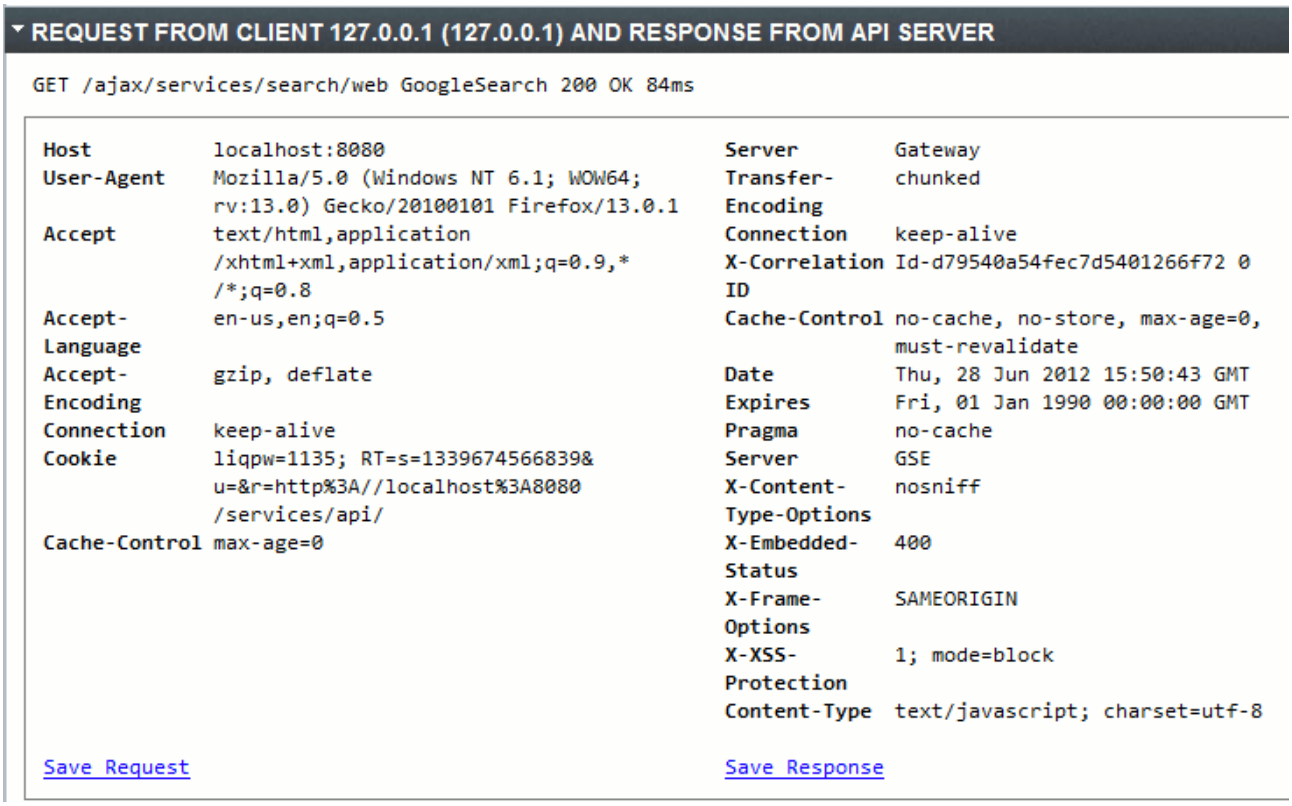

REQUEST FROM CLIENT 127.0.0.1 (127.0.0.1) AND RESPONSE FROM API SERVER

GET /ajax/services/search/web GoogleSearch 200 OK 84ms

Host	localhost:8080	Server	Gateway
User-Agent	Mozilla/5.0 (Windows NT 6.1; WOW64; rv:13.0) Gecko/20100101 Firefox/13.0.1	Transfer-Encoding	chunked
Accept	text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8	Connection	keep-alive
		X-Correlation ID	Id-d79540a54fec7d5401266f72 0
Accept-Language	en-us,en;q=0.5	Cache-Control	no-cache, no-store, max-age=0, must-revalidate
Accept-Encoding	gzip, deflate	Date	Thu, 28 Jun 2012 15:50:43 GMT
Connection	keep-alive	Expires	Fri, 01 Jan 1990 00:00:00 GMT
Cookie	liqpw=1135; RT=s=1339674566839&u=&r=http%3A//localhost%3A8080/services/api/	Pragma	no-cache
		Server	GSE
		X-Content-Type-Options	nosniff
Cache-Control	max-age=0	X-Embedded-Status	400
		X-Frame-Options	SAMEORIGIN
		X-XSS-Protection	1; mode=block
		Content-Type	text/javascript; charset=utf-8

Save Request Save Response

You can click **Save Request** or **Save Response** to download the message contents and save them to a file.

View performance statistics

The **Performance** tab displays performance statistics for the HTTP and HTTPS traffic processed by the API Gateway. For example, these include the number of requests, average bytes sent, and average duration. For example, the **Performance** page is displayed as follows:

Request from	Path	Num. requests	Avg Bytes Received	Avg Bytes Sent	Avg Duration	Min Duration	Max Duration
Client	//ajax/services/search/web	4	436	552	14	3	29
Client	/healthcheck	7	389	531	104	5	620
Client	/services/api/	2	457	602	175	18	333
Gateway	/ajax/services/search/web?v=1.0	5	469	372	255	154	371
Client	/ajax/services/search/web	5	432	578	9064	493	21802

Filter performance statistics

You can click the **Apply** in the left pane to filter the performance statistics displayed based on different criteria. By default, the statistics are grouped by path name, with a time interval of 1 day. You can select different criteria from the **GROUP BY** and **TIME INTERVAL** lists. When you have selected your search criteria, click the **Apply** button.

Detect malformed messages

Messages with malformed content or an incorrect relative path are blocked by the API Gateway and displayed on the

Dashboard tab in the **TRAFFIC** section as follows:

You can click the chart to display the list of failed messages automatically filtered on the **Traffic** tab. Click a message in the list to display the filter execution path and message content. The following example shows the execution path of a malformed message that has been blocked by the API Gateway:

Filter	Status	Audit Trail Message
Return HTTP Error 403: Access Denied (Forbidden)		
Make the "Forbidden" Message	✓	
Pass the "Forbidden" message back	✓	Successfully echoed back the message to IP Address /192.168.0.52:51959
Flag as blocked message	✓	

Monitor real-time metrics

The **Monitoring** view enables you to monitor successes, failures, exceptions, and real-time metrics for the following:

- **System**: Metrics for memory, disk space, and CPU.
- **API Services**: Metrics for messages and processing times.
- **API Methods**: Metrics for messages and processing times.
- **Clients**: Metrics for messages.
- **Remote Hosts**: Metrics for transactions, bytes sent and received, and response times.

For example, on the **System** tab, when you click a panel in the **ALL SYSTEMS** section at the top, a graph for the selected setting is displayed below. The following example shows the graph displayed for the **System CPU Avg (Max)** setting selected on the right:

You can drill down to view metrics for specific components at the bottom (for example, for a specific API Gateway group or instance, service, client, method, or remote host). You can also configure the metrics time window on the right (for example, 10 minutes, 10 hours, or 5 days).

Configure dynamic trace, logging, and monitoring

You can click the **Settings** > **Dynamic** tab to configure trace, logging, and monitoring settings on-the-fly. These are dynamic settings, which means that you do not need to refresh or deploy to the API Gateway. For example, you can specify these settings for an API Gateway system, instance, service, interface, or path. For more details, see *Configure API Gateway logging and events*.

Monitor and report on services with API Gateway Analytics

This topic shows how to monitor API services using the monitoring tools provided with the API Gateway. API Gateway Analytics is a separately installed component that enables you to monitor services and to generate reports on the stored message traffic history in your domain. For more details, see *Monitoring and reporting with API Gateway Analytics*.

Redact sensitive data from API Gateway Manager

Overview

The API Gateway enables you to remove sensitive content from messages monitored in the API Gateway Manager web console and traffic monitoring database. This also applies to API Gateway trace files at the DATA trace level. You can redact sensitive content across a range of message protocols and content types such as HTTP, JSON, XML, HTML form, and plain text.

For example, sensitive data such as user passwords or credit card details can be redacted from both request and response messages. This means that such sensitive data is only ever present in the API Gateway memory during message processing, and is never displayed onscreen or persisted to disk. This is shown in the following architecture diagram:

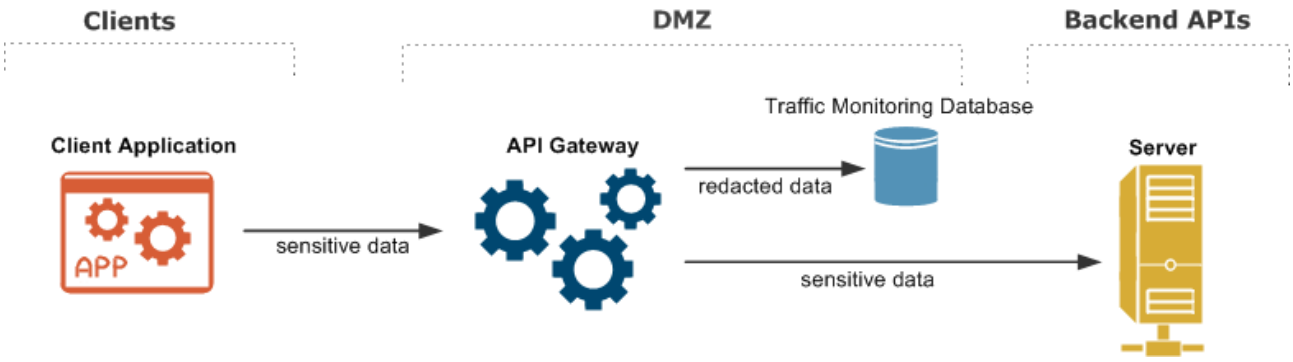

Redaction configuration

In the API Gateway configuration, message redaction rules are configured in the following XML file:

```
apigateway/system/conf/redaction.xml
```

When the API Gateway configuration is loaded, this creates redactors for the specified message protocol and content. This XML-based configuration uses the following model:

```
<Redaction enabled=”true” provider="redactors">
   <JSONRedactor>...</JSONRedactor>
   <RawRedactor>...</RawRedactor>
   <XMLRedactor>...</XMLRedactor>
   <HTTPRedactor>...</HTTPRedactor>
   <FormRedactor>...</FormRedactor>
</Redaction >
```

During the transaction processing, for each traffic monitoring stream, a chain of redactors is created for redacting the received and sent data. Each redactor removes any sensitive data that it finds and passes the data for the next redactor for processing. The redacted content is then written to the traffic monitoring database.

Each redactor defines its supported content types in `RedactMime` child elements. For example, the following shows content types for a JSON redactor:

```
<JSONRedactor>
   <RedactMime mimeType="application/json"/>
   <RedactMime mimeType="text/json"/>
   ...
</JSONRedactor>
```

Enable redaction for an API Gateway

To enable redaction for an API Gateway instance, perform the following steps:

1. Copy the sample redaction configuration file from the following directory:

   ```
   apigateway/samples/redaction/sample_redaction.xml
   ```

   to the following directory:

   ```
   apigateway/groups/<group-name>/<instance-name>/conf/redaction.xml
   ```

2. Ensure that redaction is enabled in `redaction.xml` as follows:

   ```
   <ConfigurationFragment>
     <Redaction enabled="true" provider="redactors">
     ...
   </ConfigurationFragment>
   ```

   You can customize this file to configure redactors for different message payloads (HTTP, JSON, HTML form, and plain text). This is described in the next sections.
3. Edit the following file:

   ```
   apigateway/groups/<group-name>/<instance-name>/conf/service.xml
   ```

   And add the following line at the end of the file:

   ```
   <NetService provider="NetService">
      ...
      <include file="$VDISTDIR/system/conf/redaction.xml"/>
   </NetService>
   ```

4. Restart the API Gateway instance.

Important

For all message content (HTTP, JSON, HTML form, and plain text), you must first ensure that the appropriate URL is defined in an `HTTPRedactor`. For more details, see the section called “Redact HTTP message content”.

Redact HTTP message content

You can redact any HTTP header or parameter value from the API Gateway message stream based on HTTP URLs specified in configuration. This applies to both HTTP requests and responses. The following shows a simple example configured in `redaction.xml`:

```
<HTTPRedactor>
   <HTTPURL value="/payment"/>
   <HTTPParam value="credit_card"/>
   <HTTPParam value="password"/>
   <HTTPHeader value="Authorization"/>
</HTTPRedactor>
```

This example specifies to remove the `credit_card` and `password` query string parameters and `Authorization` header from messages sent to and from the `/payment` URL.

URL path matching

Each HTTPURL value is used to match URL paths, and to determine if the redaction applies to the transaction. You can use the match attribute to specify a match for an exact URL path or for a URL prefix. The following example shows an exact URL path match:

```
<HTTPURL value="/secure_folder" match="exact"/>
```

In this exact match example:

- /secure_folder matches
- /secure_folder/ does not match
- /secure_folder/123 does not match

The following example shows a URL prefix match:

```
<HTTPURL value="/creditcard/" match="prefix"/>
```

In this prefix match example:

- /creditcard/ matches
- /creditcard/charge matches
- /creditcard/charge/1234 matches
- /creditcard does not match

HTTPURL values are also case sensitive. For example:

```
<HTTPURL value="/order/shiptoaddress"/>
```

This is different from:

```
<HTTPURL value="/ORDER/shiptoaddress"/>
```

Supported HTTP features

HTTP features such as the following are supported:

- Chunked transfer encoding
- Multipart body entities (Content-Type:multipart/)

Example: redact an HTTP Basic authorization header

This section shows an end-to-end example of redacting an HTTP Basic authorization header. Given the following HTTP request message:

```
GET /securefiles/ HTTP/1.1
Host: www.httpwatch.com
Authorization: Basic aHR0cHdhdGNoOmY=
```

And the following HTTP redactor configuration:

```
<HTTPRedactor>
   <HTTPURL value="/securefiles/" match="exact"/>
   <HTTPHeader value="Authorization"/>
</HTTPRedactor>
```

The HTTP message is redacted and stored in the traffic monitoring database as follows:

```
GET /securefiles/ HTTP/1.1
Host: www.httpwatch.com
```

Redact JSON message content

You can redact JSON content from a message by configuring a specific path to be removed. You can define a relative or absolute location for elements inside a JSON document. When you configure a specific path in the JSON redactor configuration, all elements found in that element are removed. The following general syntax is used to remove JSON content:

```
rule = path_seg [“.”path_seg]*
path_seg = wildcard/rel_wildcard/identifier/array_elem
array_elem = identifier ”[”wildcard/number“]”
identifier = char*
wildcard = “*”
rel_wildcard = “**”
```

The following simple examples show how this syntax works:

```
a.b.c, a.*.c, a.**.c, **.b.c,**.b[0].c,**.b[*].c, *.b[0].*
```

This results in the following configuration model:

```
<JSONRedactor>
   <RedactMime mimeType="text/json"/>
   ...
   <JSONPath path="a.b.c"/>
   <JSONPath path="**.b[0].c"/>
</JSONRedactor>
```

JSON redactor configuration

The following shows a simple example from `redaction.xml`:

```
<JSONRedactor>
   <RedactMime mimeType="application/json"/>
   <JSONPath path="**.subject[0].id"/>
</JSONRedactor>
```

This example removes JSON content such as the following:

```
authentication.subject[0].id
cert.subject[0].id
attribute.subject[0].id
```

Example: redact OAuth message tokens from a JSON message

This section shows an end-to-end example of redacting an OAuth message token. Given the following JSON request message:

```
{
   "access_token":"2YotnFZFEjr1zCsicMWpAA",
   "token_type":"example",
   "expires_in":3600,
```

```
   "refresh_token":"tGzv3JOkF0XG5Qx2TlKWIA",
   "example_parameter":"example_value"
 }
```

And the following JSON redactor configuration:

```
<JSONRedactor>
    <RedactMime mimeType="application/json"/>
    <JSONPath path="**.access_token"/>
    <JSONPath path="**.refresh_token"/>
</JSONRedactor>
```

The JSON message is redacted and stored in the traffic monitoring database as follows:

```
{
   "access_token":null,
   "token_type":"example",
   "expires_in":3600,
   "refresh_token":null,
   "example_parameter":"example_value"
 }
```

Redact XML message content

You can redact specific XML content from a message by configuring XML elements or attributes to be removed. The XML redactor removes sensitive data based on the document location. You can define the locations to be removed using the fully qualified name of the redacted element.

For example, to redact all the children of an element named `oracle:sensitive_data`, where `xmlns:oracle` is `http://oracle.com/`, you can use the following syntax:

```
<XMLRedactedElement localname=”sensitive_data” namespace=”http://oracle.com”
  redactionDisposition=”redactChildren”/>
```

You can specify the following XML redaction directives:

`redactChildren`	Removes all children of a specified element
`redactElement`	Redacts the specified element and all its descendants
`redactText`	Removes all text nodes from the specified element
`redactDescendents`	Redacts children and text descendants of the specified node

If you need to redact attributes of the specified node, you can configure this using `XMLRedactedAttribute` (child of `XMLRedactedElement`). `XMLRedactedElement` has two mandatory attributes, `localname` and `namespace`, which have the same meaning for `XMLRedactedAttribute` .

Note

An empty XML namespace name is the same as the default document namespace.

XML redactor configuration

The following example from `redaction.xml` removes all children from `a_namespace:a_name`. It also removes the `an_attribute_name` and `another_attribute_name` attributes:

```
<XMLRedactor>
   <RedactMime mimeType="application/xml"/>
   <RedactMime mimeType="text/xml"/>
   <!--Remove children of a_namespace:a_name and some attributtes-->
    <XMLRedactedElement localname="a_name" namespace="a_namespace"
     redactionDisposition="redactChildren">
     <XMLRedactedAttribute localname="an_attribute_name" namespace="an_attribute_namespace"/>
     <XMLRedactedAttribute localname="another_attribute_name" namespace="o"/>
    </XMLRedactedElement>
</XMLRedactor>
```

The following example removes the `b:a` element and all its children:

```
<XMLRedactor>
   <RedactMime mimeType="application/xml"/>
   <RedactMime mimeType="text/xml"/>
   <!---Remove element b:a and all its descendants-->
   <XMLRedactedElement localname="a" namespace="b" redactionDisposition="redactElement"/>
</XMLRedactor>
```

Example: redact a WS-Security username token from a XML message

This section shows an end-to-end example of redacting a WS-Security username token. Given the following XML request message:

```
<?xml version="1.0" encoding="UTF-8"?>
<Envelope xmlns="http://www.w3.org/2003/05/soap-envelope">
 <Header>
  <Security xmlns="http://docs.oasis-open.org/wss/2004/01/
     oasis-200401-wss-wssecurity-secext-1.0.xsd">
    <UsernameToken>
      <Username>root</Username>
      <Password Type="http://docs.oasis-open.org/wss/2004/01/
        oasis-200401-wss-username-token-profile-1.0#PasswordDigest">EVfkjdkljla=</Password>
      <Nonce>tKUH8ab3Rokm4t6IAlgcdg9yaEw=</Nonce>
      <Created xmlns="http://docs.oasis-open.org/wss/2004/01/
        oasis-200401-wss-wssecurity-utility-1.0.xsd">2010-08-10T10:52:42Z</Created>
    </UsernameToken>
  </Security>
 </Header>
 <Body>
   <SomeRequest xmlns="http://example.ns.com/foo/bar" />
 </Body>
</Envelope>
```

And the following XML redactor configuration:

```
<XMLRedactor>
   <RedactMime mimeType="text/xml" />
   <XMLRedactedElement localname="UsernameToken"
      namespace="http://docs.oasis-open.org/wss/2004/01/
        oasis-200401-wss-wssecurity-secext-1.0.xsd"
      redactionDisposition="redactChildren">
   </XMLRedactedElement>
</XMLRedactor>
```

The XML message is redacted and stored in the traffic monitoring database as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<Envelope xmlns="http://www.w3.org/2003/05/soap-envelope">
 <Header>
  <Security xmlns="http://docs.oasis-open.org/wss/2004/01/
    oasis-200401-wss-wssecurity-secext-1.0.xsd">
    <UsernameToken></UsernameToken>
  </Security>
 </Header>
  <Body>
    <SomeRequest xmlns="http://example.ns.com/foo/bar" />
  </Body>
</Envelope>
```

Redact HTML form message content

You can redact the content of specific HTML form fields by configuring the fields to be removed. The following shows an example from `redaction.xml`:

```
<FormRedactor>
   <RedactMime mimeType="application/x-www-form-urlencoded"/>
   <FormField value="credit_card"/>
   <FormField value="phone_number"/>
</FormRedactor>
```

This example removes the contents of the `credit_card` and `phone_number` form fields from the message.

Supported HTML form content types are as follows:

- `application/x-www-form-urlencoded`
- `multipart/formdata`

Redact raw message content

You can redact specific plain text by configuring regular expressions to define content to be removed. The following shows a configuration example:

```
<RawRedactor>
   <RedactMime mimeType="text/plain"/>
   <Regex exp="creditcard\s*=\s*(\d*{16})" redact="1" icase="true"/>
   <Regex exp="source:\b\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}\b" redact="1,2" icase="true"/>
   <Regex exp="\d*{16}" redact="0" icase="false" />
</RawRedactor>
```

In this configuration model, the `Regex` element includes the following attributes to define the redactor behavior:

`exp`	Regular expression used to match the desired content. Possible values are valid regular expressions.
`redact`	Specifies which groups in the match are redacted. Possible values are comma-separated lists of group indexes (for example, `1` or `1,2` or `4,6,7`, and so on). You can specify `0` to redact the entire match.
`icase`	Specifies whether the match is case insensitive. Possible values are `true` (case insensitive) and `false` (case sensitive).

Example: redact credit card details from raw text

This section shows some configured regular expressions and the behavior with specific raw message content. The following expression specifies to redact a defined group:

```
<Regex exp="creditcard\s*=\s*(\d{16})" redact="1" icase="false"/>
```

The following shows example message content and the behavior with this expression:

Message content	Behavior
`&creditcard=1234123412341234`	Content matches expression. Defined group 1 (`\d{16}`) is redacted (in this case, `1234123412341234`).

The following expression specifies to redact multiple defined groups:

```
<Regex exp="ccdigits:(\d{1,4})\.(\d{1,4})\.(\d{1,4})\.(\d{1,4})" redact="1,2,3"
icase="false"/>
```

The following shows example message content and the behavior with this expression:

Message content	Behavior
`ccdigits:1234.2345.3456.4567`	Content matches expression. Defined groups 1 (`\d{1,4}`), 2 (`\d{1,4}`)), and 3 (`\d{1,4}`) are redacted (in this case `1234`, `2345`, and `3456`. Defined group 4 (`\d{1,4}`) is left intact (in this case `4567`).

The following expression specifies to redact content using case insensitivity:

```
<Regex exp="credit\s*card\s*\d{16}" redact="0" icase="true/>
```

The following shows example message content and the behavior with this expression:

Message content	Behavior
`credit card 123456781234567`	Content matches expression. Entire match (`credit card 1234567812345678`) is redacted.
`Credit Card 1234567812345678`	Content matches expression because of the `icase` attribute. Entire match (`Credit Card 1234567812345678`) is redacted.

Redact sensitive data from log files

For details on how to redact sensitive data from domain audit log and access log files, see the following topics:

- *Configure API Gateway logging and events*
- *Transaction access log settings*

Configure API Gateway for API Gateway Analytics

Overview

This topic explains how to configure an API Gateway instance and Node Manager to store metrics in the metrics database used by API Gateway Analytics. The following data streams are used to populate the metrics database:

- Transaction and system data: Transaction data includes clients, services, remote hosts, and protocols. System data includes CPU, memory and disk usage, and SLA breaches. The API Gateway writes this data to a transaction event log, with a new log file automatically created every 5 minutes. The Node Manager parses completed event logs and updates the metrics database. Data from this stream is displayed in API Gateway Analytics on the **System**, **API Services**, **Remote Hosts** and **Protocols** tabs.
- Transaction audit log events: These are written directly to the metrics database by the API Gateway instance, and are displayed in the **Transaction Audit Log** in API Gateway Analytics.

This topic explains how to perform the following tasks:

- Use Policy Studio to configure an API Gateway instance to write audit logging events to the metrics database, and to write metrics data to the transaction event log.
- Use the `managedomain` command to configure the Node Manager to process event logs and update the metrics database.

Important

To view API Gateway metrics in API Gateway Analytics, you must perform the steps in this topic.

Prerequisites

This topic assumes that you have already installed and configured API Gateway Analytics using the steps described in the *API Gateway Installation Guide*.

Connect to the API Gateway in Policy Studio

To connect to the API Gateway in Policy Studio, perform the following steps:

1. Ensure the Admin Node Manager and API Gateway are running. For more details, see *Start and stop the API Gateway*.
2. On the Policy Studio welcome page, click the Admin Node Manager server session to make a connection.
3. Specify your connection details (host, port, user name, and password). The default connection URL is:

   ```
   https://localhost:8090/api
   ```

4. In the API Gateway topology view, double-click the API Gateway instance to load its configuration.

Configure the metrics database connection

To configure the API Gateway connection to the metrics database, perform the following steps:

1. Expand the **External Connections** > **Database Connections** node in the Policy Studio tree.
2. Right-click the **Default Database Connection** tree node, and select **Edit**.
3. Configure the database connection to point to the same metrics database created when API Gateway Analytics was installed. For details on connection settings, see the *API Gateway Policy Developer Guide*.

4. Verify that your database connection is configured correctly by clicking the **Test Connection** button on the **Configure Database Connection** dialog.

Tip

You can troubleshoot your database connection by viewing the contents of your server `.trc` file in the `INSTALL_DIR/apigateway/trace` directory. For more details, see *Configure API Gateway diagnostic trace*.

Configure transaction audit logging to the metrics database

To configure the API Gateway instance to write transaction audit log data to the metrics database, perform the following steps:

Note

This is only required to display data on the **Transaction Audit Log** tab in API Gateway Analytics. It is not required to display data on the **System**, **API Services**, **Remote Hosts** and **Protocols** tabs.

1. In the Policy Studio tree, select the **Server Settings** node, and select **Logging** > **Transaction Audit Log** in the window on the right.
2. Select the **Database** tab, and select **Enable logging to database**.
3. Select the **Default Database Connection** from the drop-down list if appropriate. Alternatively, select a database connection that you have configured. You must ensure that your database connection points to the same metrics database configured when API Gateway Analytics was installed.
4. For more details, see *Transaction audit log settings*.

Tip

To write the content of message transactions to the database, you must also configure the **Log Message Payload** filter in your policies (for example, at the start and end of the policy). For more details, see the *API Gateway Policy Developer Guide*.

Configure the API Gateway to write to the transaction event log

To configure the API Gateway instance to write transaction data to the transaction event log, perform the following steps:

1. In the Policy Studio tree, select the **Server Settings** node, and select **Logging** > **Transaction Event Log** in the window on the right.
2. Ensure **Writing to Transaction Event Log** is selected.
3. To enable monitoring of protocol and remote host metrics, select the **Monitoring** > **Traffic Monitor** node, and ensure the following settings are selected:
   - **Enable Traffic Monitor**
   - **Record inbound transactions**
   - **Record outbound transactions**

For more details, see *Transaction event log settings*.

Deploy the updated configuration to the API Gateway

You must deploy these configuration changes to the API Gateway. Click the **Deploy** button in the toolbar, or press F6.

The API Gateway now sends transaction audit logging to the metrics database, and writes transaction data to the transaction event log. The final step is to configure the Node Manager to read the transaction event logs and write system and transaction metrics to the metrics database.

Configure the Node Manager to process event logs and update the metrics database

If you have not already done so, you must use the `managedomain` tool to enable the Node Manager to process event logs from your API Gateway host, and to write metrics data to the metrics database.

All API Gateway instances running on the host node generate transaction event log files. These files are all written to the same folder, and are collectively processed and aggregated by the Node Manager on the host, and then written to the metrics database. The metrics database provides the data for the graphical charts in the **Monitoring** view in API Gateway Analytics.

Use the managedomain interactive menu

You can enable metrics using the interactive `managedomain --menu` command. The following shows an example:

```
Select option: 2
Select a host:
   1) LinuxMint01
   2) Enter host name
Enter selection from 1-2 [2]: 1
Enter a new host name [LinuxMint01]:
Enter a new Node Manager name [Node Manager on LinuxMint01]:
Enter a new Node Manager port [8090]:
There is only one Node Manager in this domain so it must remain as an Admin Node Manager
Do you want to create an init.d script for this Node Manager [n]:
Do you want to reset the passphrase for the Node Manager on this host ? [n]:
Do you wish to edit metrics configuration (y or n) ? [n]: y
Do you wish to enable metrics (y or n) ? [y]: y
Enter metrics database URL [jdbc:mysql://127.0.0.1:3306/Reporter]:
Enter metrics database username [root]:
Enter metrics database plaintext password [*******]:
Testing Database connectivity for : jdbc:mysql://127.0.0.1:3306/Reporter, user : root
Metrics database connectivity succeeded
Metrics generation enabled. All other specified metrics settings updated.
Metrics settings updated successfully. Please reboot Node Manager on completion of this
program.
Completed successfully.
Hit enter to continue...
```

Use the managedomain command options

Alternatively, you can use `managedomain` command options to enable metrics when initializing a host, adding a host other than the Admin Node Manager, or editing a host.

The following example shows enabling metrics when initializing a host machine:

```
./managedomain --initialize --metrics_enabled y
   --metrics_dburl jdbc:mysql://127.0.0.1:3306/reports
   --metrics_dbuser root --metrics_dbpass MY_DB_PWD
```

The following example shows enabling metrics when adding a host machine other than the Admin Node Manager:

```
./managedomain --add --anm_host MY_HOSTNAME --nm_entitystore_passphrase MY_CONFIG_PWD
   --metrics_enabled y --metrics_dbuser root --metrics_dbpass MY_DB_PWD
   --metrics_dburl jdbc:mysql://1.2.3.4:3306/reports --nm_name MY_NODE_MNGR --port 8055
```

The following example shows enabling metrics when editing a host machine in the domain:

```
./managedomain --edit_host --nm_entitystore_passphrase bonjour
   --metrics_enabled y --metrics_dburl jdbc:mysql://127.0.0.1:3306/reports
   --metrics_dbuser root --metrics_dbpass MY_DB_PWD
```

The `managedomain` metrics options are described as follows:

Option	Description
`--nm_entitystore_pass`	Specifies the encryption passphrase used to access the API Gateway instance configuration. If no passphrase has been set, omit this argument. For more details, see *Configure an API Gateway encryption passphrase*.
`--metrics_enabled`	Specifies whether writing of metrics data is enabled. Enter `y` or `n`.
`--metrics_dburl`	Specifies the JDBC URL for the metrics database (for example, `jdbc:mysql://127.0.0.1:3306/reports`).
`--metrics_dbuser`	Specifies the metrics database user (for example, `root` or `analyticsuser`).
`--metrics_dbpass`	Specifies the password for the metrics database user.

Important

When the `managedomain` command has finished, you must restart the Node Manager.

For more details on `managedomain`, see *Managedomain command reference*.

Configure additional options for event log processing in the Node Manager

The parameters described in this section specify how transaction event logs are processed in the Node Manager. You can configure these optional settings by editing the Node Manager configuration using the `esexplorer` tool.

For example, perform the following steps:

1. Change to the following directory:

Windows	`INSTALL_DIR\apigateway\Win32\bin`
UNIX/Linux	`INSTALL_DIR/apigateway/posix/bin`

2. Enter the `esexplorer` command.
3. Select **Store** > **Connect**.
4. Browse to `INSTALL_DIR/apigateway/conf/fed/configs.xml`
5. Select **System Components** > **Metrics Generation Configuration** in the tree on the left.
6. Configure the appropriate fields in the window on the right:

Option	Description
`sourceEventLogDir`	Specifies the folder in which the Node Manager looks for event log files. This should match the API Gateway transaction event log directory set in Policy Studio (see *Transaction event log settings*). Defaults to `${environment.VDISTDIR}/events`.

Option	Description
`retainProcessedEventLogs`	Specifies whether processed event logs should be deleted or retained in a separate directory. By default, event logs are deleted when their contents are written to the metrics database. Logs can be retained if they are needed for audit purposes or as input to a custom analytics process. Defaults to `false`.
`processedEventLogDir`	When `retainProcessedEventLogs` is `true`, specifies the directory to which event files are moved after being processed by the Node Manager. Defaults to `${environment.VDISTDIR}/events/processed`.
`dirSizeMb`	If `retainProcessedEventLogs` is `true`, specifies the maximum size of the `processedEventLogDir`. When the configured size is reached, the oldest log files in the directory are deleted. Defaults to 1024 MB.
`processCustomMessageAttributes`	Specifies whether message attributes contained in the transaction event log, are written to the database `transaction_data` table. Defaults to `true`. For more details, see *Transaction event log settings*.
`processCustomMetrics`	Specifies whether custom metrics generated by the API Gateway Java Metrics API and written to the transaction event log are written to the database. Defaults to `true`. For more details, see the following: • `../apigateway/docs/javadoc/com/vordel/reporting/rtm/api` • *Transaction event log settings*

7. You must restart the Node Manager after editing its configuration using `esexplorer`.

Tip

When making changes using `esexplorer`, ensure that you open the latest configuration. For example, you could overwrite changes made using `managedomain` if an old version of the configuration was loaded into `esexplorer` and then updated.

Further information

For details on how to use API Gateway Analytics to monitor your domain, see *Monitoring and reporting with API Gateway Analytics*.

Monitoring and reporting with API Gateway Analytics

Overview

API Gateway Analytics monitors, records, and reports on the history of message traffic between API Gateway instances and various services, remote hosts, and clients running in an API Gateway domain. You can use API Gateway Analytics to monitor traffic and perform root cause analysis at the level of the domain, API Gateway instance, service, remote host, and client. You can also filter the display based on any selected time period. For example, this defaults to the last 7 days, but you can specify any date range.

Note

This topic describes how to use the API Gateway Analytics web-based interface to monitor your domain. The Node Manager on each host must be configured to write metrics data to the same database that API Gateway Analytics reads from. The API Gateway can write to the same database for transaction audit logging if required. This topic assumes that you have already performed the steps in *Configure API Gateway for API Gateway Analytics*.

Launch API Gateway Analytics

To launch API Gateway Analytics, perform the following steps:

1. Start the API Gateway Analytics server using the `oaganalytics` script in the `/bin` directory of your API Gateway Analytics installation.
2. Using the default port, connect to the API Gateway Analytics interface in a browser at the following URL:

   ```
   http://HOST:8040/
   ```

   where `HOST` points to the IP address or hostname of the machine on which API Gateway Analytics is installed.
3. Log in using the default `admin` user with a password of `changeme`.

Note

You should change the default credentials by editing your API Gateway Analytics configuration in Policy Studio under the **Users and Groups** tree node.

The default `admin` user is local to API Gateway Analytics. Updating these credentials does not affect access to other Oracle products.

Monitor the API Gateway system

The API Gateway Analytics **System** view includes the following tabs:

- **Systems**
- **System Resources**

Systems

On the **Systems** tab, click a panel in the **SYSTEMS** section at the top to display graph for the selected system-level metrics below. For example, the available metrics include the following:

- **Successes**: The number of successful messages processed in the domain (that generated a success in an API

Gateway policy).

- **Failures**: The number of failed messages processed in the domain (that generated a failure in an API Gateway policy).
- **Exceptions**: The number of messages in the domain that generated an exception in an API Gateway policy.
- **Active**: The uptime of API Gateway instances.
- **SLA breaches**: The number of Service Level Agreement (SLA) breaches in the domain.

The following example shows messages successfully sent displayed in a simple domain with a single API Gateway instance:

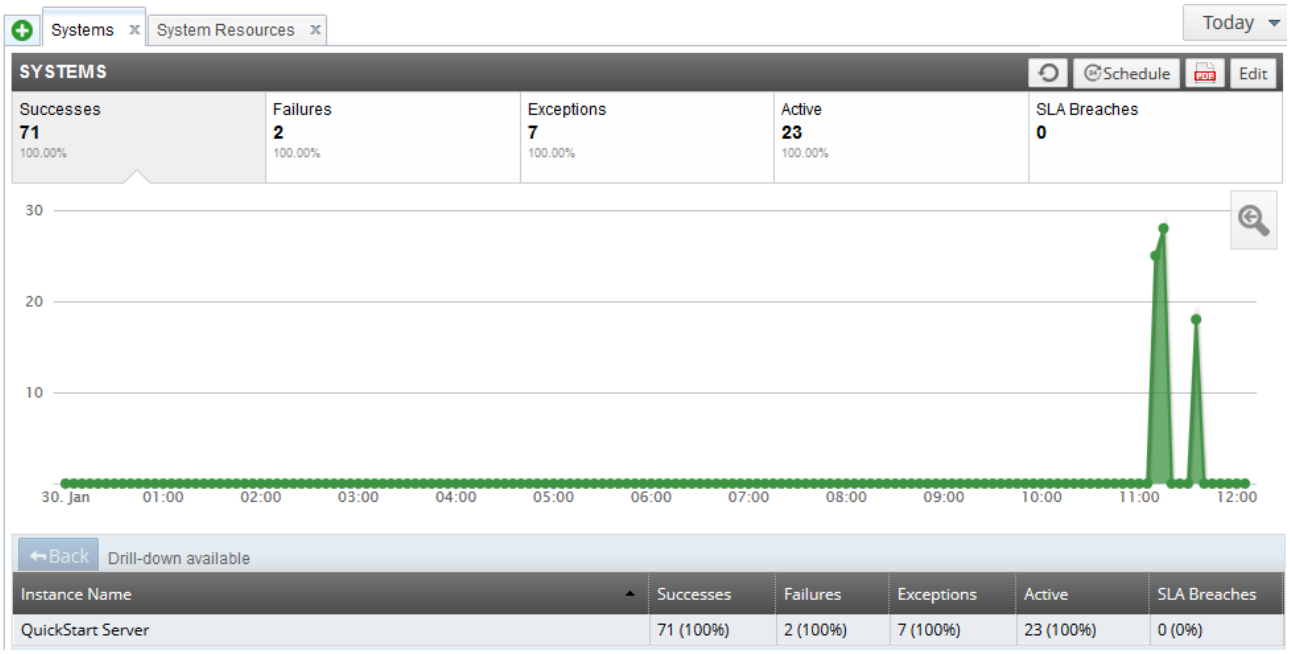

The table at the bottom shows all the API Gateway instances that are sending monitored traffic to protected services, clients, and remote hosts in your domain. You can click an API Gateway instance in the table to drill down and view graphs for the selected instance. Click **Back** on the left to return to the **ALL SYSTEMS** view.

System resources

On the **System Resources** tab, click a panel in the **SYSTEM RESOURCES** section at the top to display graph for the selected system resource metrics below. For example, the available metrics include the following:

- **Instance CPU (Avg %)**: Average amount of CPU used by the API Gateway instance.
- **System CPU (Avg %)**: Average amount of CPU used on the host machine running the API Gateway.
- **Instance Memory Used (Avg MB)**: Average amount of memory used by the API Gateway instance.
- **System Memory (Max MB)**: Maximum amount of memory used on the machine hosting the API Gateway. This includes memory used by the API Gateway and all other processes running on the machine.
- **System Memory Total (Max MB)**: Total amount of available memory on the machine hosting the API Gateway.

The following example shows the average system CPU displayed in a simple domain with one API Gateway instance:

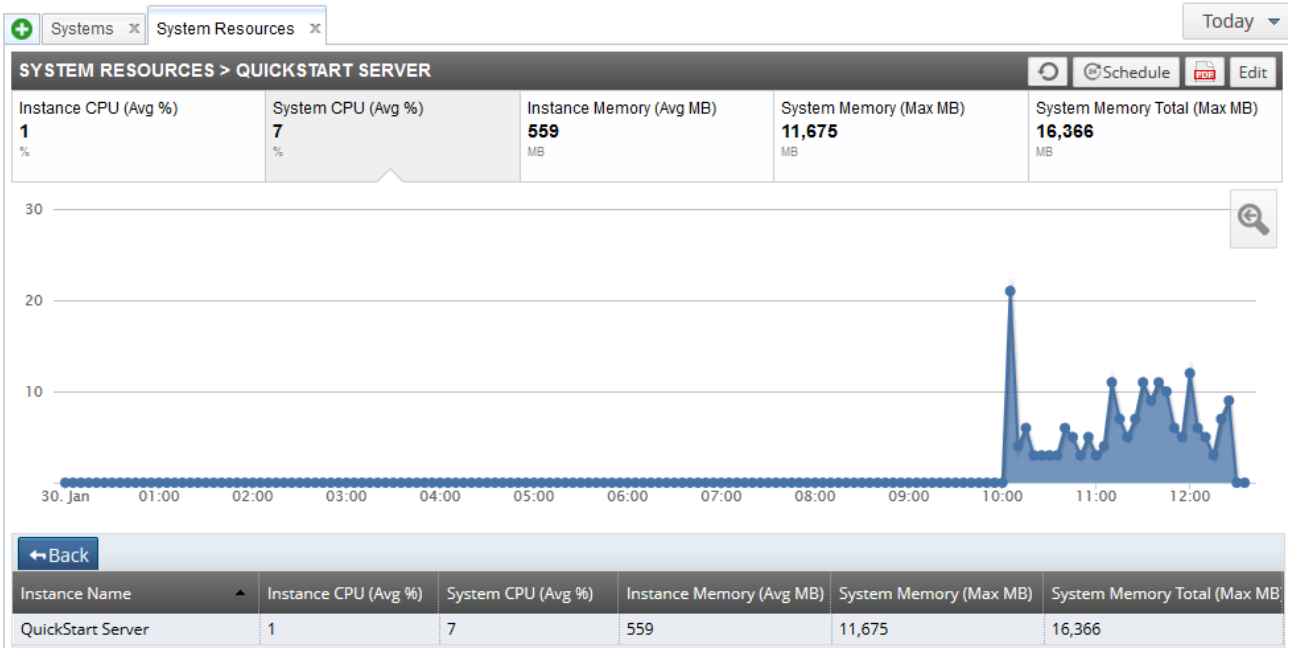

The table at the bottom shows all API Gateway instances that are sending monitored traffic to protected services, clients, and remote hosts in your domain. You can click an API Gateway instance in the table to drill down and view graphs for the selected instance. Click **Back** on the left to return to the **SYSTEM RESOURCES** view.

Monitor API services, methods, and clients

The **API Services** view shows metrics for services that are virtualized by API Gateway instances in your domain. For more details on virtualizing services, see the *API Gateway Policy Developer Guide*.

The **API Services** view includes the following tabs:

- **Load Balance**: Metrics for the number of messages and processing times.
- **Clients (Service of)**: Metrics for the number of messages, successes, and failures.
- **Clients**: Metrics for the number of messages, successes, failures, and exceptions.
- **API Services (Methods of)**: Metrics for the number of messages, successes, failures, and exceptions.
- **API Services (Clients of)**: Metrics for the number of messages, successes, failures, and exceptions.
- **API Services**: Metrics for the number of messages, successes, failures, and exceptions.
- **API Service Performance**: Metrics for number of messages, successes, failures, and exceptions.

Example: API service performance

For example, In the **API Services** > **API Service Performance** tab, click a panel in the **API SERVICE PERFORMANCE** section at the top to display a graph for the selected service-level metric below. For example, the available metrics include the following:

- **Messages**: The number of API service messages processed in the API Gateway domain.
- **Processing Time (Min ms)**: The minimum time taken to process a message, including all calls to remote servers.
- **Processing Time (Max ms)**: The maximum time taken to process a message, including all calls to remote servers.
- **Processing Time (Avg ms)**: The average time taken to process a message, including all calls to remote servers.

The following example shows the maximum processing time in a simple domain with multiple API services:

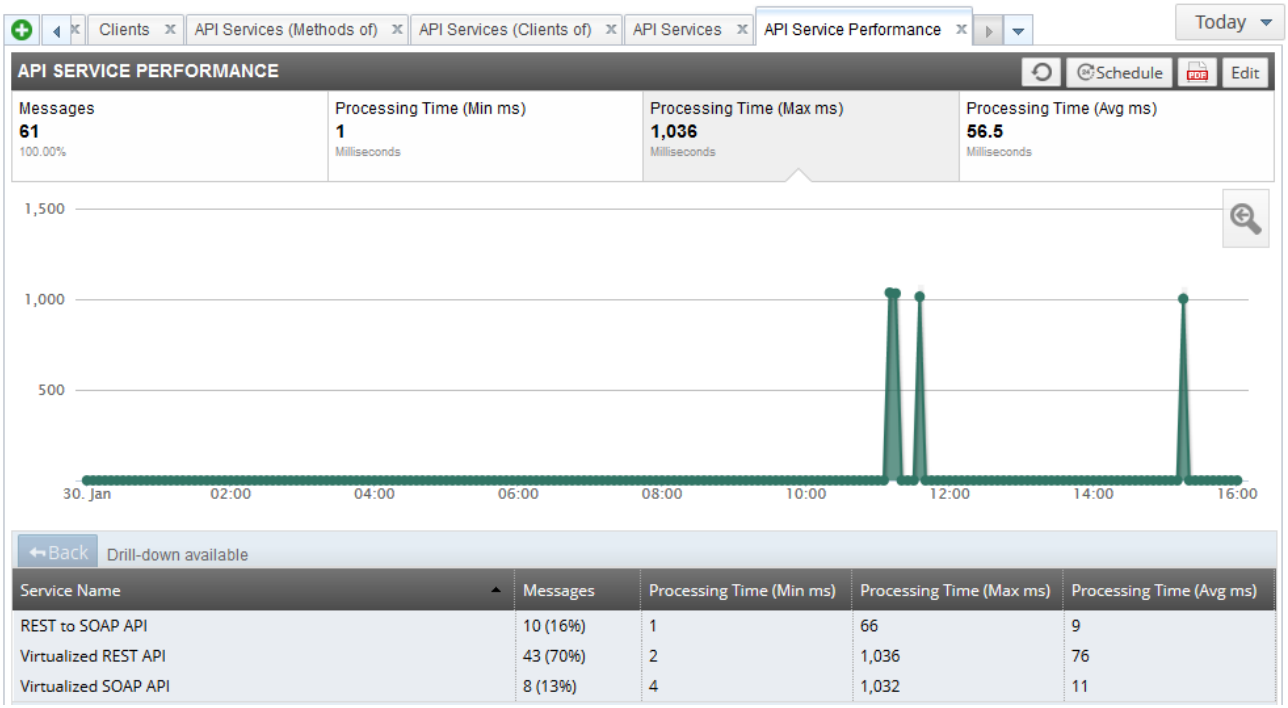

The table at the bottom shows all services protected by API Gateway instances in your domain. You can click a service in the table to drill down and view graphs for the selected service. Click **Back** on the left to return to the **API SERVICE PERFORMANCE** view.

Note

A service must first have been sent a message before it is displayed in the **API Services** view.

Monitor remote hosts

The **Remote Hosts** view displays metrics for all the remote hosts that have been configured in your domain. It shows details such as the number of message transactions that have been sent to this remote host, together with the total number of bytes sent to and received from this host.

In the **Remote Hosts** view, click a panel in the **REMOTE HOSTS** section at the top to display graph for the selected remote host metric below. For example, the available metrics include the following:

- **Transactions**: The number of message transactions to remote hosts.
- **Volume Bytes (In)**: The total number of bytes sent to remote hosts.
- **Volume Bytes (Out)**: The total number of bytes received from remote hosts.
- **Response Time (Avg ms)**: The average response time to remote hosts.

The following example shows the total number of bytes sent to remote hosts:

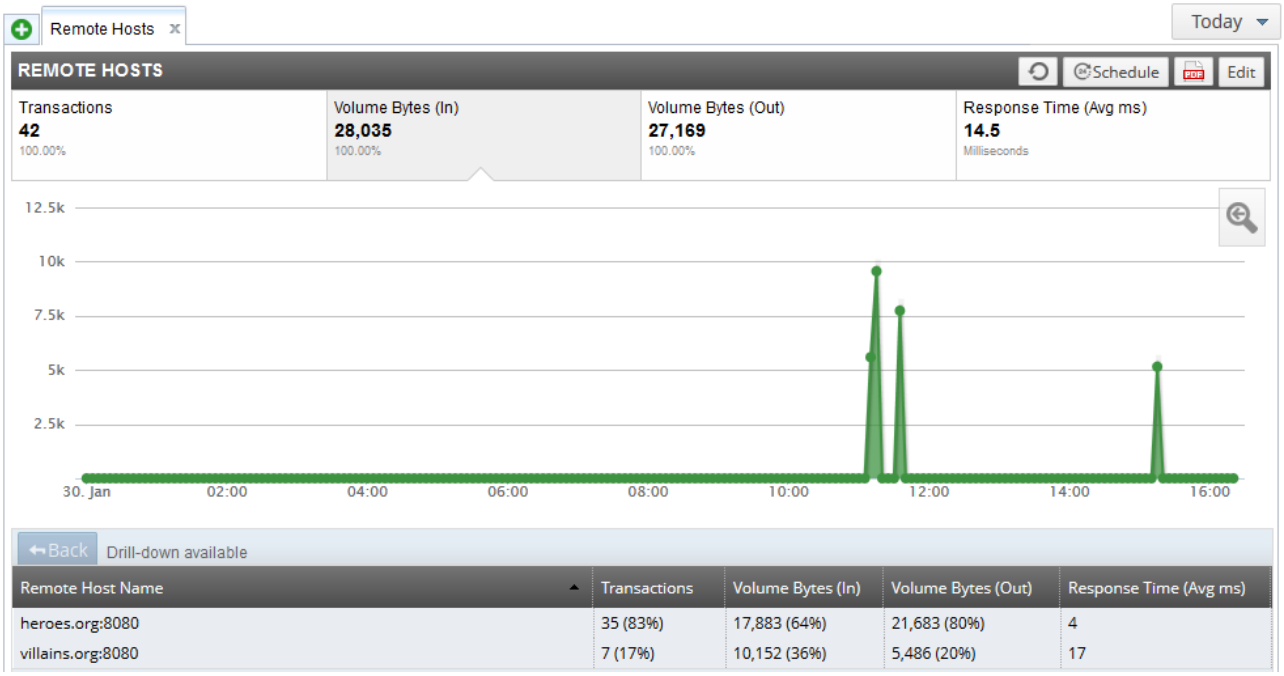

The table at the bottom shows all the remote hosts connected to by API Gateway instances in your domain. You can click a remote host in the table to drill down and view graphs for the selected remote host. Click **Back** on the left to return to the **REMOTE HOSTS** view.

Monitor protocols

The **Protocols** view enables you to monitor the different message protocols in your domain (for example, HTTP, Web-socket, JMS, File Transfer, and so on). Click a panel in the **INBOUND PROTOCOLS** section at the top to display graph for the selected protocol metrics below. For example, the available metrics include the following:

- **Successes**: The number of successful messages that are processed in the domain.
- **Failures**: The number of failed messages that are processed in the domain.
- **Exceptions**: The number of messages that generated an exception in an API Gateway policy.

The following example shows the number of HTTP messages that generated an exception in the domain:

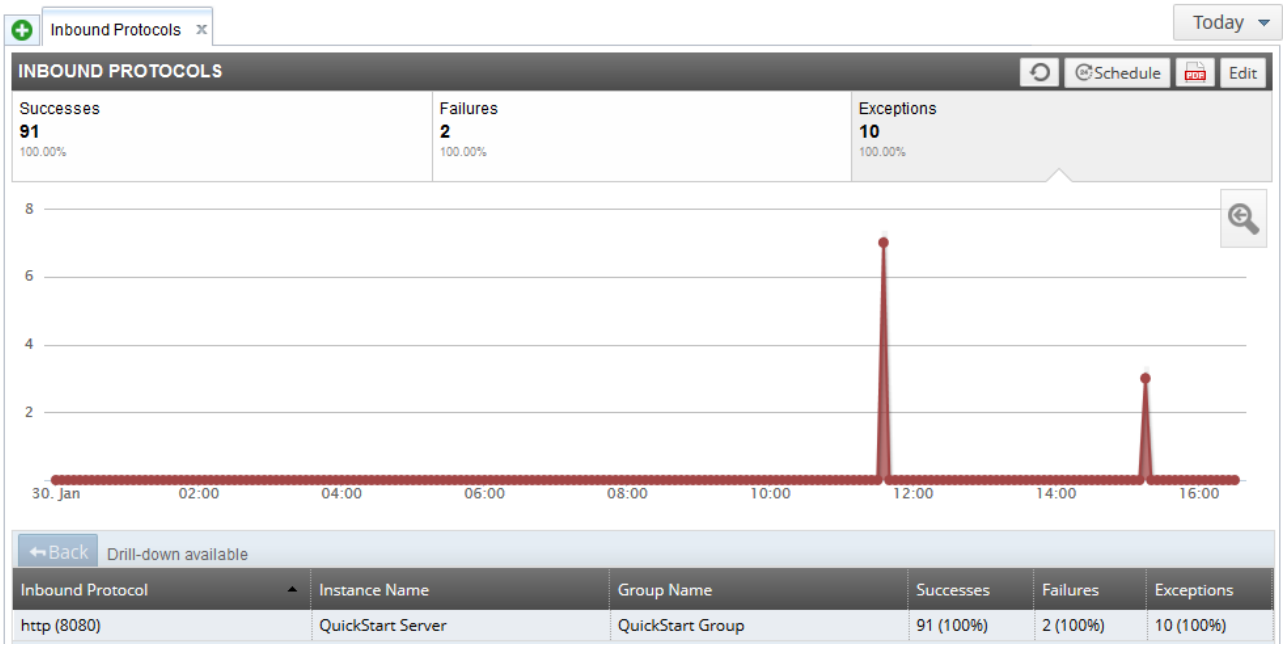

The table at the bottom shows all the protocols used by API Gateway instances in your domain. You can click a protocol in the table to drill down and view graphs for the selected protocol (for example, **http 8080**). Click **Back** on the left to return to the **INBOUND PROTOCOLS** view.

Audit transactions

The **Transaction Audit Log** view enables you to filter the transaction audit log messages messages generated by API Gateway instances in your domain. For example:

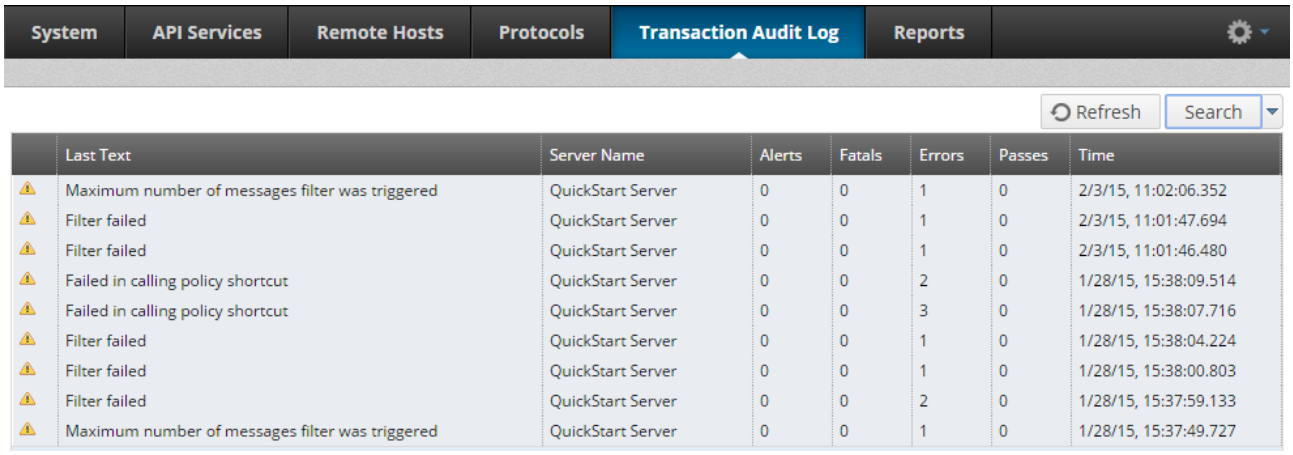

You can filter the log messages by clicking the **Search** button on the right in the toolbar. The **Query Editor** dialog enables you to create a query to filter log messages by details such as time period, severity level, filter type or name, and message text. When you have added your search criteria, click **Search** at the bottom to run the query. You can also save the query for later use.

When you click **Search**, the log messages that match the search criteria specified in the query are displayed in the table.

For example, the details displayed in the table include the log message text, API Gateway name, alerts, and time. You can also double-click an item in the list for more details (for example, transaction ID, filter category, and filter name).

Important

To view data in the **Transaction Audit Log** tab, you must configure the API Gateway to write to the metrics database. In the Policy Studio tree, select **Server Settings** > **Logging** > **Transaction Audit Log** > **Database**. For more details, see *Transaction audit log settings*

Schedule custom reports

API Gateway Analytics uses message metrics stored in a centralized metrics database by the Node Managers running in your domain. The Node Managers store metrics for the virtualized services exposed by the local API Gateway instances, and for the services, clients, and remote host connections that they protect.

API Gateway Analytics can generate usage reports and charts based on the stored metrics data, and enables you to configure custom reports to suit the needs of your environment. This includes viewing available metrics for each target report type, grouping and filtering metrics, and what to display on drill through.

Create reports in a monitoring view

In the **System**, **API Services**, **Remote Hosts**, and **Protocols** views, you can create reports by clicking **Schedule** on the right in the toolbar. For example, to generate reports on API Service clients, perform the following steps:

1. Click the **API Services** > **Clients** view.
2. Click **Edit** on the right in the toolbar to specify the report settings:
   - **Name**: Defaults to the tab name (for example, **Clients**).
   - **Group by**: Select how the report is displayed (for example, by **Client Name**, **Instance Name**, and so on). You can also select the metrics that are displayed (for example, **Messages**, **Successes**, and so on).
   - **Filter by**: Select condition options to filter the display (for example, only show entries with a specific **Client Name**).
   - **Enable drill-down**: Select this to use the value of the **Group by** columns to provide an additional drill-down report. You can choose to group by additional fields. This is enabled by default.

   These filtering and grouping mechanisms enable you to answer questions such as what clients used an API Service, or which API services were used by a client. For example, to show clients that used `Service1`, you can create a custom report that groups by **Client Name** and filters where **Service Name** is `Service1`.

   Note

   The group-by mechanism only applies to the data table below the report chart. The chart remains the same.
3. Click **Save** in the toolbar.
4. Click **Schedule** in the toolbar.
5. You can also click the **PDF** button in the toolbar to generate a PDF version of the report.

You can follow a similar sequence of steps to generate reports in the **System**, **Remote Hosts**, and **Protocols** views.

Using the reports view

Alternatively, you can create a report in the **Reports** view. Perform the following steps:

1. Click **New report** on the left in the toolbar.
2. Enter the report **Name** in the dialog.
3. Select the report **Type**, and click **OK**.

4. Configure the schedule and output options on the right as appropriate. Defaults to PDF output daily at 9am, starting from today.
5. Click **Apply** to save your settings.
6. Click **Create** to create the report.

The following example shows a custom report grouped by **Service Name**:

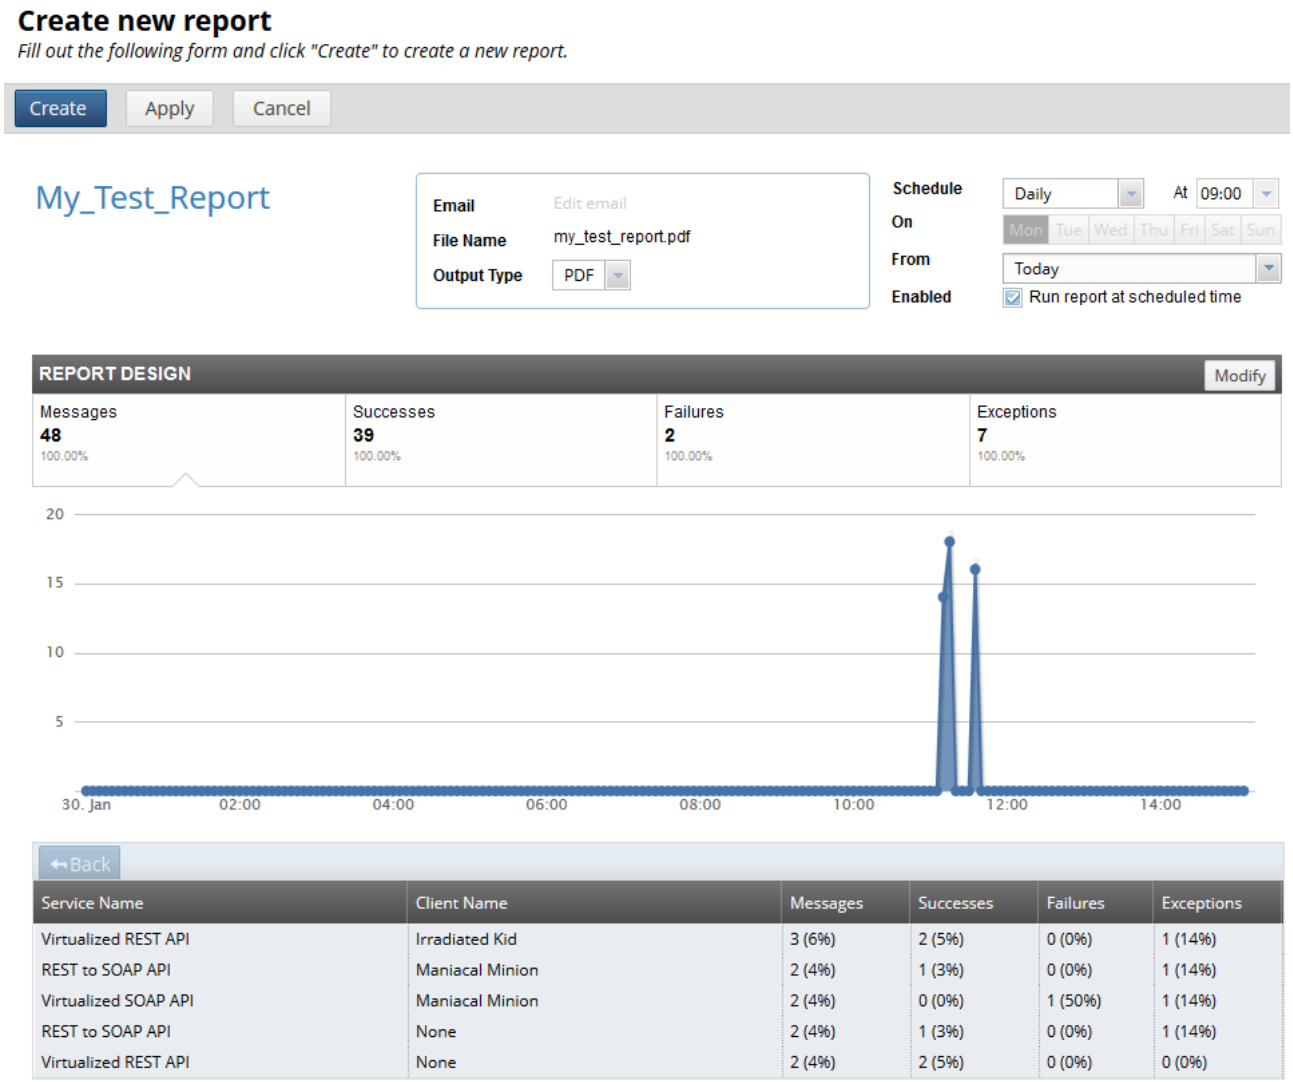

When you have created a report, you can select it in list on the **Managing Reports** page, and click **Generate now**. Alternatively, click **Delete** to remove the report from the list.

Further information

For more details, see the following:

- *API Gateway Installation Guide*
- *Configure API Gateway for API Gateway Analytics*

- *Configure scheduled report settings*
- *Purge the metrics database*

Configure scheduled report settings

Overview

You can schedule reports to run on a regular basis, and have the results emailed to the user in PDF format. These reports include summary values at the top (for example, the number of requests, SLA breaches, alerts triggered, and unique clients in a specified week) followed by a table of services, and their aggregated usage data (for example, the number of requests on each service).

The report data is for the configured *current week of the report*, which is compared to the week before. You can set the configured *current week of the report* to be the actual current calendar week or any prior week (provided there is corresponding data in the database).

To configure scheduled reports, right-click the **Listeners** > **Oracle Analytics** node in the Policy Studio tree, and select **Database Archive**.

Database configuration

Click the browse button the right, and select a pre-configured database connection in the dialog. This setting defaults to the `Default Database Connection`. To add a new database connection, right-click the **Database Connections** node, and select **Add DB connection**. You can also edit or delete existing nodes by right-clicking and selecting the appropriate option. Alternatively, you can add database connections under the **External Connections** node in the Policy Studio tree view. For more details on creating database connections, see the *API Gateway Policy Developer Guide*.

Scheduled reports configuration

You can configure the following settings for scheduled reports:

Enable Report Generation:
Select whether to enable scheduled reports in PDF format. When selected, by default, this runs a scheduled weekly report on Monday morning at 0:01. For details on configuring a different time schedule, see the next setting. This setting is not selected by default.

When **Enable Report Generation** is enabled, you can configure the following settings on the **Report Generator Process** tab:

Connect to API Gateway Analytics as User:
Enter the username and password used to connect to the report generator process. Defaults to the values entered using the `configureserver` script.

Output:
Enter the directory used for the generated report files in the **Output Directory** field, or click **Choose** to browse to the directory. Defaults to the directory entered using the `configureserver` script (for example, `c:\temp\reports`). You can also select to **Do not delete report files after emailing**. This setting is not selected by default.

SMTP configuration

When **Enable Report Generation** is enabled, you can configure the following settings on the **SMTP** tab. These settings default to those entered using the `configureserver` script. For more details, see the *API Gateway Installation Guide*.

Email generated reports:
Select whether to email generated PDF report files. This is not selected by default.

Do not delete report files after emailing:
Select whether to keep generated PDF report files after they have been sent. This is not selected by default.

Email Recipient (To):
Enter the recipient of the automatically generated email (for example, `user@mycorp.com`). Use a semicolon-separated list of email addresses to send reports to multiple recipients.

Email Sender (From):
The generated report emails appear *from* the sender email address specified here (for example, `no-reply@mycorp.com`).

Note

Some mail servers do not allow relaying mail when the sender in the **From** field is not recognized by the server.

SMTP Server Settings:
Specify the following fields:

Outgoing Mail Server (SMTP)	Specify the SMTP server used to relay the report email (for example, `smtp.gmail.com`).
Port	Specify the SMTP server port to connect to. Defaults to port 25.
Connection Security	Select the connection security used to send the report email (`SSL`, `TLS`, or `NONE`). Defaults to `NONE`.

Log on Using:
If you are required to authenticate to the SMTP server, specify the following fields:

User Name:	Enter the user name for authentication.
Password:	Enter the password for the user name specified.

Purge the metrics database

Overview

You can use the `dbpurger` command to connect to your metrics database and to purge old data. This command also enables you to retain a specified amount of data, and to archive all data.

For details on configuring the connection to your metrics database, see the *API Gateway Installation Guide*.

Run the dbpurger command

You can run the `dbpurger` command from the following directory:

Windows	`INSTALL_DIR\oaganalytics\Win32\bin`
Linux/UNIX	`INSTALL_DIR/oaganalytics/posix/bin`

Command options
You can specify the following options to the `dbpurger` command:

Option	Description
`-h, --help`	Displays help message and exits.
`-p PASSPHRASE, --passphrase=PASSPHRASE`	Specifies the configuration passphrase (leave blank for zero length).
`--dbname=DBNAME`	Specifies the database name (mutually exclusive with `dburl`, `dbuser`, and `dbpass` options).
`--dburl=DBURL`	Specifies the database URL.
`--dbuser=DBUSER`	Specifies the database user.
`--dbpass=DBPASS`	Specifies the database passphrase.
`--archive`	Archive all data.
`--out=OUT`	Archive all data in the specified directory.
`--purge`	Purge data from the database. You must also specify the `--retain` option.
`--retain=RETAIN`	Specifies the amount of data to retain (for example, `30days`, `1month`, or `1year`). You must specify this option with the `--retain` option.

Example commands

This section shows examples of running `dbpurger` in default interactive mode and of specifying command-line options.

Running dbpurger in interactive mode
The following example shows the output when running the `dbpurger` command in interactive mode. This example archives all data, retains three months of data, and purges older data from the database:

```
    >dbpurger
```

```
    Choosing: Default Database Connection
    Archive database (Y, N) [N]: y
    Archive path [./archive]:
    Purge an amount of data from the database (Y, N) [N]: y
    Amount of data to retain (e.g. 1year, 3months, 7days) [3months]:
    Wrote archive: .\archive\process_groups.xml
    Wrote archive: .\archive\processes.xml
    Wrote archive: .\archive\metric_types.xml
    Wrote archive: .\archive\audit_log_sign.xml
    Wrote archive: .\archive\time_window_types.xml
    Wrote archive: .\archive\audit_log_points.xml
    Wrote archive: .\archive\audit_message_payload.xml
    Wrote archive: .\archive\transaction_data.xml
    Wrote archive: .\archive\metric_groups.xml
    Wrote archive: .\archive\metric_group_types.xml
    Wrote archive: .\archive\metrics_alerts.xml
    Wrote archive: .\archive\metrics_data.xml
    Purging data older than: Wed Jun 27 15:26:00 BST 2012
    Purging table: audit_log_sign... deleted 0 rows
    Purging table: transaction_data... deleted 0 rows
    Purging table: audit_message_payload... deleted 7 rows
    Purging table: audit_log_points... deleted 16 rows
    Purging table: metrics_alerts... deleted 4 rows
    Purging table: metrics_data... deleted 703 rows
```

Specify command-line options to dbpurger

The following example shows the output when specifying options the `dbpurger` command. This example retains 30 days of data, and purges older data from the database:

```
dbpurger.bat --dburl=jdbc:mysql://127.0.0.1:3306/reports --dbuser=root
--dbpass=fred --purge --retain=30days
```

Note

You can run `dbpurger` without a password by specifying the name of the database connection. For example:

```
dbpurger.bat --dbname="Default Database Connection" --archive --out=archive.dat
```

Configure API Gateway logging and events

API Gateway logs and events

You can use the **Logs** and **Events** view in the API Gateway Manager web console to view and search API Gateway log and event files. You can configure API Gateway log and event settings in the **Server Settings** in Policy Studio. The **Settings** > **Dynamic** tab in API Gateway Manager also enable you to configure log settings at runtime.

This section first provides an overview of the different types of logs and events that are displayed in API Gateway Manager. This topic then explains how to configure the different types of logs and events, and shows where to find more detailed information.

Domain management and diagnostics

The **Logs** view displays the following API Gateway logs for domain management and runtime diagnostics:

- **Domain Audit**: Displays management changes at the API Gateway domain level (for example, updates to API Gateway configuration, topology, login, or deployment). The domain audit log is configured by default. For more details, see the section called “Configure audit logs per domain”.
- **Trace**: Records detailed diagnostic and debugging information on API Gateway instance execution (for example, services starting or stopping, and messages sent through the API Gateway). The trace log is configured by default. You should submit trace log files when raising issues with Oracle Support. For more details, see *Configure API Gateway diagnostic trace*.

Message transactions

The **Logs** view also displays the following API Gateway logs for message transactions:

- **Transaction Audit**: Records policy-level message transaction log entries generated by each filter as the message passes through the filter. You can define the transaction audit log output for each filter in Policy Studio (for example, success, failure, or abort), and configure different log output destinations. For more details, see the section called “Configure transaction audit log destinations”.
- **Transaction Access**: Provides a summary of HTTP request and response message transactions in Apache HTTP Server format. You can configure the access log per API Gateway path. For more details, see the section called “Configure transaction access logs per path”.

Events and alerts

The **Events** view displays in-memory events for transaction audit logs, system alerts, and SLA alerts. For more details, see the section called “Manage API Gateway events and alerts”.

Configure audit logs per domain

The domain audit log captures management changes in the API Gateway domain that are written by the Admin Node Manager and by API Gateway instances. This includes details such as API Gateway configuration changes, log in/log out, deployments, user, or topology changes. For example, user Joe deployed a new configuration, admin user created a new group, or user Jane has read deployment data. The domain audit log is enabled by default. However, you can configure filtering options such as the number of events displayed, time interval, and event type.

For example, the displayed event types include the following:

- Configuration
- Service

- Application
- Topology
- User store
- Key Property Store (KPS)

View in API Gateway Manager

To view domain audit log events in the API Gateway Manager web console, perform the following steps:

1. In the API Gateway Manager, select **Logs** > **Domain Audit**.
2. Configure the number of events displayed in the **Max results per server** field on the left. Defaults to `1000`.
3. Configure the **Time Interval** for events. Defaults to `1 day`.
4. Click the **Filter** button to add more viewing options (**Event Type** or **Groups and Servers**).
5. Click **Apply** when finished.

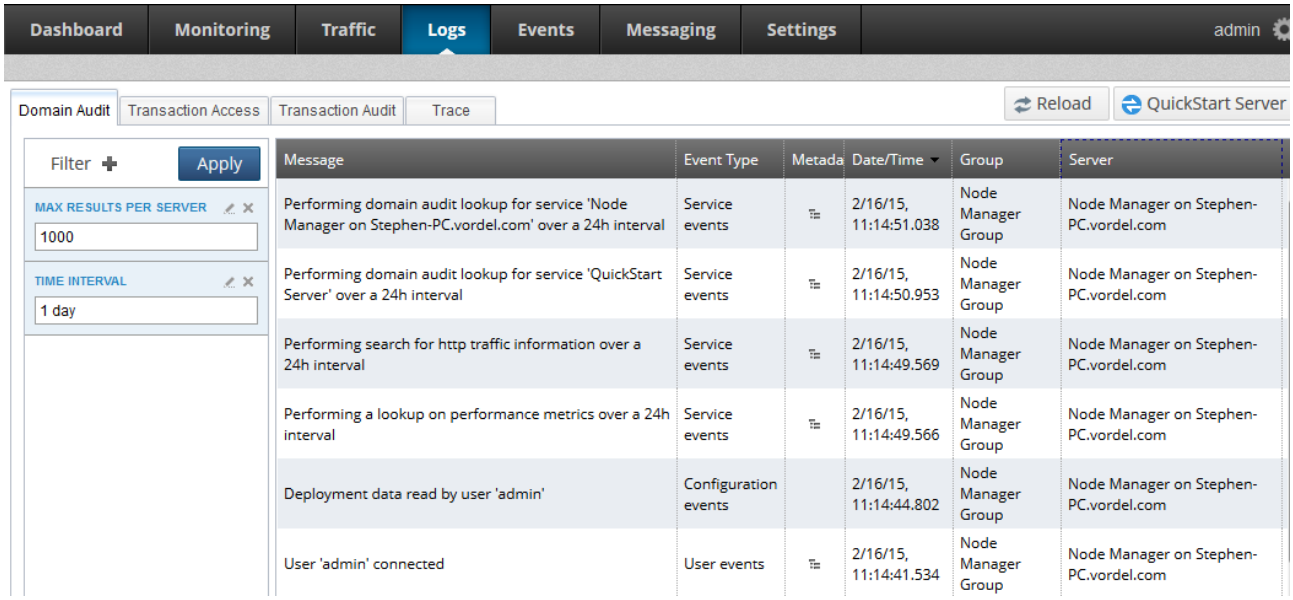

View the domain audit log file

Alternatively, you can view contents of the domain audit log file. For example, the following shows the file for the Admin Node Manager:

```
<install-dir>/apigateway/logs/audit.log
```

For example:

```
{"timestamp":1397724538713,"message":"User 'admin' connected with 3 defined user roles",
   "eventId":107,"metadata":{"userID":"admin"}}
{"timestamp":1397724539638,"message":"Deployment data read by user 'admin'","eventId":1037,
   "metadata":{}}
{"timestamp":1397726232992,"message":"Performing domain audit lookup for service
   'Node Manager on cayote.acme.com' over a 24h interval","eventId":9,"metadata":
   {"userID":"admin","serviceID":"nodemanager-1"}}
{"timestamp":1397726235233,"message":"Performing domain audit lookup for service
   'Node Manager on cayote.acme.com' over a 24h interval","eventId":9,"metadata":
```

```
   {"userID":"admin","serviceID":"nodemanager-1"}}
```

The default maximum size for the audit log file is 1 MB. A new file is created when the server instance restarts. The maximum of files stored in the `logs` directory is `50`. When this maximum number of log files is reached, the files roll over, and the oldest files are deleted.

Redact domain audit log output

You can also customize and redact the contents of the output in the `logs/audit.log` file. For example, for security purposes, you can redact sensitive information, such as specific query parameters that contain customer details, passwords, or credit card information. Alternatively, you can prevent the file from becoming flooded with specific messages, such as `GET` API calls for metrics.

You can use the following file to customize the output of the domain audit log file:

```
<install-dir>/apigateway/conf/apiaudit.xml
```

This file enables you to specify rules to filter out sensitive details or noisy API calls. The default file contains some predefined rules (for example, filtering out metrics). You can use this file to specify whether an entry is made to the domain audit log file, and to specify the contents of the text in the output message.

For example, the following entry specifies an `outputMessage` for all `GET` messages on the `ops/setserviceconfig` path:

```
<apiauditrule>
   <method>GET</method>
   <path>^ops/getserviceconfig$</path>
   <pathMatch>MATCHES</pathMatch>
     <queryArgs>*</queryArgs>
     <outputMessage>Update configuration for service '${serviceName}':
        ${queryArgs}</outputMessage>
</apiauditrule>
```

The following example specifies no `outputMessage` for `GET` messages on the `api/monitoring/metrics` path:

```
<apiauditrule>
   <method>GET</method>
   <path>api/monitoring/metrics</path>
   <pathMatch>BEGINS_WITH</pathMatch>
</apiauditrule>
```

Domain audit rule syntax

The rules in the `apiaudit.xml` file analyze the traffic passing through the API Gateway router service, and control the entries in the domain audit log. These rules are checked in the order specified in the file. The `method`, `path` and `pathMatch` elements determine whether a rule is triggered. If a rule is triggered, all subsequent rules are ignored. You should specify all rules in order of priority (for example, most sensitive or noisy first).

The domain audit rule elements are described as follows:

Element	Description
`method`	Required comma-separated list of HTTP methods (`GET`, `PUT`, and so on). Use the * wildcard to specify all methods.
`path`	Required regular expression that specifies a URL path (for example, `^api/domainaudit/search$`). Use the * wildcard to specify all paths.
`pathMatch`	Required path matching statement (one of the following: `MATCHES`, `BE-`

Element	Description
	GINS_WITH, ENDS_WITH, CONTAINS, DOES_NOT_CONTAIN, IS, IS_NOT, DOES_NOT_MATCH).
queryArgs	Option to specify query string arguments output in the log. To redact certain arguments, you must explicitly list only the arguments you wish to show in a comma-separated list. Leaving this blank or omitting the element specifies that no query arguments are displayed. The * wildcard specifies that all query arguments are available for printing.
outputMessage	Option to specify the message output printed in the log. Leaving this blank or omitting the element means that no entry is made in the domain audit log for this rule.

For more details and example rules, see the contents of the `conf/apiaudit.xml` file.

Configure transaction audit log destinations

The API Gateway provides detailed transaction audit logging for specific message filters (for example, the request, time of the request, where the request was routed to, and response returned to the client). You can configure transaction logging output to a number of different destinations:

- Text file
- XML file
- Database
- Local syslog
- Remote syslog
- System console

Transaction audit logging is not configured by default. To configure where transaction audit log information is sent, perform the following steps:

1. In the Policy Studio tree, select **Server Settings** > **Logging** > **Transaction Audit Log**.
2. Specify the required settings on the appropriate tabs (for example, **Text File**, **Database**, and **XML File**).
3. When finished, click **Apply Changes** at the bottom right.
4. Click the **Deploy** button in the toolbar to deploy your settings to the API Gateway.

When a transaction audit log **Text File** destination is enabled, a text log file is written and displayed in API Gateway Manager in the **Logs** -> **Transaction Audit** view. For details on configuring all available options, see *Transaction audit log settings*.

Configure transaction audit logs per filter

You can configure the transaction audit log level and log message for a specific filter as follows:

1. In the Policy Studio tree, click any policy to display it in the canvas on the right (for example, **QuickStart** > **Virtualized Services** > **REST** > **GetProducts**.
2. Double-click a filter on the canvas to edit (for example, **Connect to Heroes REST Service**).
3. Click **Next** to display the **Transaction Audit Logging Level and Message** window.
4. Select **Override Logging Level for this filter**.

5. Select the log levels required for troubleshooting (for example, **Fatal** and **Failure**).
6. Enter any non-default log messages if required.
7. Click **Finish**.
8. Click the **Deploy** button in the toolbar to deploy your settings to the API Gateway.

Message payload logging

You can also enable logging of the message payload in the transaction using a **Log Message Payload** filter in your policy, or by enabling dynamic payload logging in API Gateway Manager (see the section called "Configure dynamic trace and log settings"). When message payload logging is enabled, the transaction payload is logged to the transaction audit log destinations that are configured.

For more details on transaction audit logging for specific message filters, see the following in the *API Gateway Policy Developer Guide*:

- **Transaction Audit Logging Level and Message** monitoring in each filter
- **Log Message Payload** filter

Configure transaction event logs per API Gateway

The transaction event log provides a summary of each API Gateway transaction. These logs are persisted, and when configured, are used to generate metrics on the **Monitoring** tab in API Gateway Analytics.

Transaction event logging is enabled by default. To configure the transaction event log output, perform the following steps:

1. In the Policy Studio tree, select **Server Settings** > **Logging** > **Transaction Event Log**.
2. Specify the required settings (for example, enable API method or client monitoring).
3. When finished, click **Apply Changes** at the bottom right.
4. Click **Deploy** in the toolbar to deploy your settings to the API Gateway.

For details on configuring all the available options, see *Transaction event log settings*.

Configure transaction access logs per path

The access log provides summary of the HTTP request and response messages that are written to an access log file in the format used by Apache HTTP Server. For example, this includes details such as the remote hostname, user login name, and authenticated user name.

Access logging is not configured by default. To configure the access log output, perform the following steps:

1. In the Policy Studio tree, select the **Server Settings** > **Logging** > **Transaction Access Log**.
2. Specify the required settings (for example, file name, directory name, max files, and so on).
3. When finished, click **Apply Changes** at the bottom right.
4. You must also configure the access log at the service level on a specific relative path. In the Policy Studio tree, select the relative path, right-click it in the **Resolvers** pane, and select **Edit**.
5. Click the **Logging Settings** tab, and select **Include in server access log records**.
6. Click the **Deploy** button in the toolbar to deploy your settings to the API Gateway.

For details on configuring all the available options, see *Transaction access log settings*.

Manage API Gateway events and alerts

The **Events** view in API Gateway Manager enables you to view and search the contents of the following in-memory API Gateway events and alerts:

- **Transaction Audit**: When policy-level transaction audit logging is configured in API Gateway filters, this displays an in-memory list of transaction audit log events. A transaction audit log destination does not need to be enabled for the in-memory list of events to be updated. For more details, see the section called “Configure transaction audit logs per filter”.
- **Alerts**: When system **Alert** filters are configured in your policies, this displays an in-memory list of alert events. An alert destination does not need to be enabled for in-memory list of events to be updated. For more details, see the *API Gateway Policy Developer Guide*.
- **SLA Alerts**: When **SLA Alert** filters are configured in your policies, this displays an in-memory list of Service Level Agreement (SLA) alert events. An alert destination does not need to be enabled for in-memory list of events to be updated. For more details, see the *API Gateway Policy Developer Guide*.

The following example shows some **Transaction Audit** events:

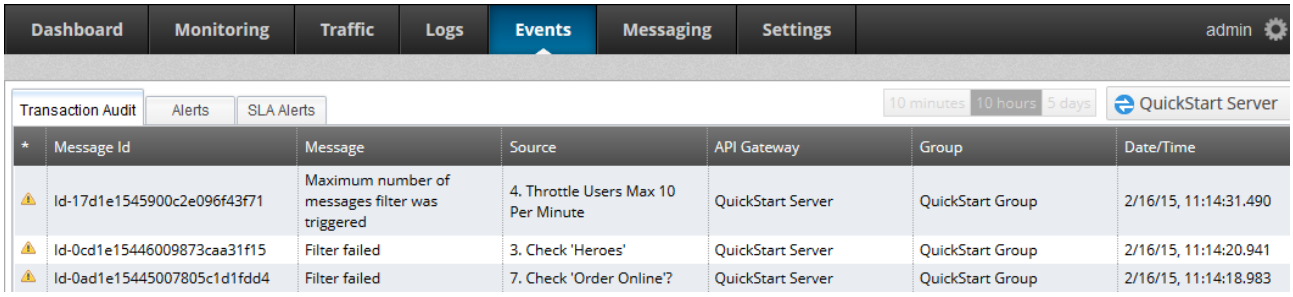

Configure dynamic trace and log settings

You can click the **Settings** > **Dynamic** tab in API Gateway Manager to configure dynamic trace, logging, and monitoring settings on-the-fly at runtime. These settings are dynamic because you do not need to refresh or deploy updates to the API Gateway. You can specify these settings for an API Gateway system, instance, service, interface, or path.

For example, the top-level **SYSTEM SETTINGS** allow you to configure logging of inbound and outbound transactions, policy paths, and message trace. You can select an HTTP interface in the tree on the left to configure the **INTERFACE SETTINGS**, **TRAFFIC MONITORING SETTINGS**, and **Trace level**. You can select the API Gateway instance in the tree to configure its trace level.

You can also select a relative path or service in the tree, and configure the following options:

- **SERVICE SETTINGS**: Select whether the service is enabled.
- **TRANSACTION AUDIT LOGGING LEVEL**: You can select from **Fatal**, **Failure**, and **Success**. See the **Transaction Audit Logging Level and Message** monitoring setting in the *API Gateway Policy Developer Guide*.
- **TRANSACTION AUDIT PAYLOAD LOGGING**: Disabled by default. See the **Log Message Payload** filter in the *API Gateway Policy Developer Guide*.
- **TRANSACTION ACCESS LOGGING**: See *Transaction access log settings*.

When finished, click **Apply Changes** at the bottom of the page.

The following example shows the dynamic settings available for a service:

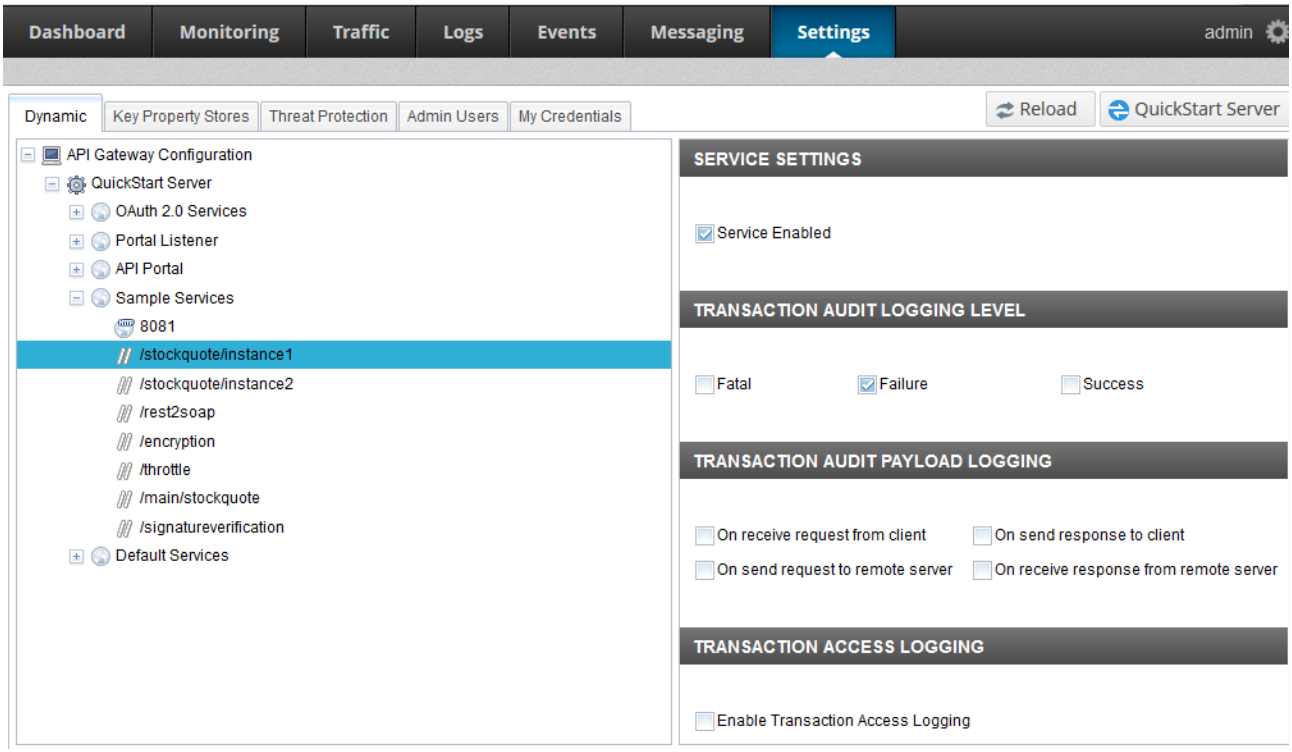

Further information

For more details on real-time monitoring and traffic monitoring, see *Monitor services in API Gateway Manager*.

Configure API Gateway diagnostic trace

Overview

By default, the API Gateway outputs diagnostic trace and debugging information to record details about its runtime execution. For example, this includes services starting or stopping, exceptions, and messages sent through the API Gateway. This information can then be used by API Gateway administrators and developers for diagnostics and debugging purposes, and is useful when contacting Oracle Support. You can view and search the contents of API Gateway tracing in the following locations:

- **Logs** > **Trace** view in API Gateway Manager
- Trace files in the following location:
  - Admin Node Manager: `<install-dir>/trace`
  - API Gateway instance: `<install-dir>/groups/<group-id>/<instance-id>/trace`
  - API Gateway Analytics: `<install-dir>/trace`
- Windows or UNIX console window for the running server

The **Logs** view in API Gateway Manager enables you to view and search the contents of the API Gateway trace log, domain audit log, and transaction logs. This topic explains how to configure the trace log only. For more details, see *Configure API Gateway logging and events*.

For details on how to redact sensitive data from trace files (for example, user passwords or credit card details), see *Redact sensitive data from API Gateway Manager*.

View API Gateway trace files

Each time the API Gateway starts up, by default, it outputs a trace file to the API Gateway `trace` directory (for example, `INSTALL_DIR\groups\group-2\server1\trace`). The following example shows an extract from a default API Gateway trace file:

```
INFO    15/Jun/2012:09:54:01.047 [1b10] Realtime monitoring enabled
INFO    15/Jun/2012:09:54:01.060 [1b10] Storing metrics in database disabled
INFO    15/Jun/2012:09:54:03.229 [1b10] cert store configured
INFO    15/Jun/2012:09:54:03.248 [1b10] keypairs configured
...
```

The trace file output takes the following format:

```
TraceLevel   Timestamp [thread-id] TraceMessage
```

For example, the first line in the above extract is described as follows:

`TraceLevel`	`INFO`
`Timestamp`	`15/Jun/2012:09:54:01.047` (*`day`*`:`*`hours`*`:`*`minutes`*`: `*`seconds`*`:`*`milliseconds`*)
`Thread-id`	`[1b10]`
`TraceMessage`	`Realtime monitoring enabled`

Set API Gateway trace levels

The possible trace levels in order of least to most verbose output are as follows:

- FATAL
- ERROR
- INFO
- DEBUG
- DATA

where FATAL is the least verbose and DATA is the most verbose. The default trace level is INFO.

Set trace levels
You can set the trace level using the following different approaches:

Startup trace	When the Admin Node Manager is starting up, it gets its trace level from the tracelevel attribute of the SystemSettings element in /system/conf/nodemanager.xml. You can set the trace level in this file if you need to diagnose boot up issues.
Default Settings trace	When the API Gateway has started, it reads its trace level from the Default Settings for the API Gateway instance. To set this trace level in the Policy Studio, click the **Server Settings** node in the Policy Studio tree, select a **Trace level** from the drop-down list.
Interface level trace	You can configure an HTTP/HTTPS interface with a different trace level to that specified in the Default Settings. For example, the default API Gateway management port (8090) has a trace level set to ERROR to ensure it is not too verbose. To configure the trace level for an interface in the Policy Studio, right-click the interface under the **Listeners** node, select **Edit**, and select a **Trace level** from the drop-down list.
Dynamic trace	You can also change dynamic API Gateway trace levels on-the-fly in API Gateway Manager. For more details, see the section called "Configure dynamic trace, logging, and monitoring".

Configure API Gateway trace files

By default, trace files are named *servername_timestamp*.trc (for example, server1_20130118160212.trc). You can configure the settings for trace file output in INSTALL_DIR/system/conf/trace.xml, which is included by INSTALL_DIR/system/conf/nodemanager.xml. By default, trace.xml contains the following setting:

```
<FileRolloverTrace maxfiles="500" filename="%s_%Y%m%d%H%M%S.trc"/>
```

This setting means that the API Gateway writes Node Manager trace output to nodemanageron*hostname_timestamp*.trc (for example, nodemanageron127.0.0.1_20130118160212.trc) in the trace directory of the API Gateway installation. And the maximum number of files that the trace directory can contain is 500.

FileRolloverTrace Attributes
The FileRolloverTrace element can contain the following attributes:

filename	File name used for trace output. Defaults to the tracecomponent attribute read from the SystemSettings element.

directory	Directory where the trace file is written. Defaults to INSTALL_DIR/trace when not specified.
maxlen	Maximum size of the trace file in bytes before it rolls over to a new file. Defaults to 16777216 (16 MB).
maxfiles	Maximum number of files that the trace directory contains for this filename. Defaults to the maximum integer value (2147483647).
rollDaily	Whether the trace file is rolled at the start of the day. Defaults to true.

The following setting shows example attributes:

```
<FileRolloverTrace maxfiles="5" maxlen="10485760" rollDaily="true"
    directory="/mydir/log/trace" filename="myserver.trc"/>
```

This setting means that the API Gateway writes trace output to myserver.trc in the /mydir/log/trace directory, and rolls the trace files over at the start of each day. The maximum number of files that this directory can contain is 5, and the maximum trace file size is 10 MB.

Writing Trace Output to Syslog
On UNIX and Linux, you can send API Gateway trace output to syslog. In your INSTALL_DIR/system/conf/trace.xml file, add a SyslogTrace element, and specify a facility. For example:

```
<SyslogTrace facility="local0"/>
```

Run trace at DEBUG level

When troubleshooting, it can be useful to set to the trace level to DEBUG for more verbose output. When running a trace at DEBUG level, the API Gateway outputs the status of every policy and filter that it processes into the trace file.

Debugging a Filter
The trace output for a specific filter takes the following format:

```
Filter name {
    Trace for the filter is indented
    to the following point to make it clear
    to identify output from the filter
} status, in x milliseconds
```

The status is 0, 1, or 2, depending if the filter failed, succeeded, or aborted. For example, the result of an **WS-Security Username Token** filter running successfully is as follows:

```
DEBUG 12:43:59:093 [11a4] run filter [WS-Security Username Token] {
DEBUG 12:43:59:093 [11a4]    WsUsernameTokenFilter.invoke: Verify username and password
DEBUG 12:43:59:093 [11a4]    WsAuthN.getWSUsernameTokenDetails:
                             Get token from actor=current actor
DEBUG 12:43:59:093 [11a4]    Version handler -  creating a new ws block
DEBUG 12:43:59:108 [11a4]    Version handler - adding the ws element to the wsnodelist
DEBUG 12:43:59:108 [11a4]    Version handler -  number of ws blocks found:1
DEBUG 12:43:59:124 [11a4]    No timestamp passed in WS block, no need to check timestamp
DEBUG 12:43:59:139 [11a4]    WsAuthN.getWSUsernameTokenDetails: Check <Created> element
                             in token. Value=2010-08-06T11:43:43Z
DEBUG 12:43:59:139 [11a4]    WS Nonce TimeStamp Max Size is 1000 and wsNonces cache is 4
DEBUG 12:43:59:139 [11a4]    Add WS nonce for key [joe:2010-08-06T11:43:43Z].
                             New cache size [5].
```

```
DEBUG 12:43:59:155 [11a4]    WsBasicAuthN.getUsername: Getting username
DEBUG 12:43:59:171 [11a4]    WS-Security UsernameToken authN via CLEAR password
DEBUG 12:43:59:171 [11a4]    VordelRepository.checkCredentials: username=joe
DEBUG 12:43:59:186 [11a4] } = 1, in 62 milliseconds
```

Debug a policy
The trace output for a policy shows it running with all its contained filters, and takes the following format:

```
policy Name {
    Filter 1{
        Trace for the filter
    } status, in x milliseconds
    Filter 2{
        Trace for the filter
    } status, in x milliseconds
}
```

For example, the following extract shows a policy called when running a simple service:

```
DEBUG ... run circuit "/axis/services/urn:cominfo"...
DEBUG ... run filter [Service Handler for 'ComInfoServiceService'] {
DEBUG ...   Set the service name to be ComInfoServiceService
DEBUG ...   Web Service context already set to ComInfoServiceService
DEBUG ...   close content stream
DEBUG ...   Calling the Operation Processor Chain [1. Request from Client]...
DEBUG ...   run filter [1. Request from Client] {
DEBUG ...      run filter [Before Operation-specific Policy] {
DEBUG ...         run circuit "WS-Security UsernameToken AuthN"...
DEBUG ...           run filter [WS-Security Username Token] {
                     ...
DEBUG ...           } = 1, in 62 milliseconds
DEBUG ...           ..."WS-Security UsernameToken AuthN" complete.
DEBUG ...       } = 1, in 74 milliseconds
...
```

Debug at startup
When running a startup trace with a `DEBUG` level set in the `SystemSettings`, the API Gateway outputs the configuration that it is loading. This can often help to debug any incorrectly configured items at start up, for example:

```
DEBUG 14:38:54:206 [1ee0] configure loadable module type RemoteHost, load order = 500000
DEBUG 14:38:54:206 [1ee0] RemoteHost {
DEBUG 14:38:54:206 [1ee0]     ESPK: 1035
DEBUG 14:38:54:206 [1ee0]     ParentPK: 113
DEBUG 14:38:54:206 [1ee0]     Key Fields:
DEBUG 14:38:54:206 [1ee0]         name: {csdwmp3308.wellsfargo.com}
DEBUG 14:38:54:206 [1ee0]         port: {7010}
DEBUG 14:38:54:221 [1ee0]     Fields:
DEBUG 14:38:54:221 [1ee0]         maxConnections: {128}
DEBUG 14:38:54:268 [1ee0]         turnMode: {off}
DEBUG 14:38:54:268 [1ee0]         inputBufSize: {8192}
DEBUG 14:38:54:268 [1ee0]         includeContentLengthRequest: {0}
DEBUG 14:38:54:268 [1ee0]         idletimeout: {15000}
DEBUG 14:38:54:268 [1ee0]         activetimeout: {30000}
DEBUG 14:38:54:268 [1ee0]         forceHTTP10: {0}
DEBUG 14:38:54:268 [1ee0]         turnProtocol: {http}
DEBUG 14:38:54:268 [1ee0]         includeContentLengthResponse: {0}
DEBUG 14:38:54:268 [1ee0]         addressCacheTime: {300000}
DEBUG 14:38:54:268 [1ee0]         outputBufSize: {8192}
DEBUG 14:38:54:268 [1ee0]         sessionCacheSize: {32}
DEBUG 14:38:54:268 [1ee0]         _version: {1}
DEBUG 14:38:54:268 [1ee0]         loadorder: {500000}
```

```
DEBUG 14:38:54:268 [1ee0]         class: {com.vordel.dwe.NativeModule}
DEBUG 14:38:54:268 [1ee0] }
```

For details on setting trace levels and running a startup trace, see the section called “Set API Gateway trace levels”.

Run trace at DATA level

When the trace level is set to DATA, the API Gateway writes the contents of the messages that it receives and sends to the trace file. This enables you to see what messages the API Gateway has received and sent (for example, to reassemble a received or sent message).

Note

On Windows, you can not rely on the console output because it truncates large messages.

Searching by thread
Every HTTP request handled by the API Gateway is processed in its own thread, and threads can be reused when an HTTP transaction is complete. You can see what has happened to a message in the API Gateway by following the trace by thread ID. Because multiple messages can be processed by the API Gateway at the same time, it is useful to eliminate threads that you are not interested in by searching for items that only match the thread you want.

You can do this using the search feature in the API Gateway Manager **Logs** view. Enter the thread you wish to search for in the **Only show lines with text** textbox, and click **Refresh**. Alternatively, you can do this on the command line using vi by specifying the thread ID as a pattern to search for in the trace file. In this example, the thread ID is 145c:

```
:g!/145c/d
```

The following example shows the DATA trace when a message is sent by the API Gateway (message starts with snd): >

```
DATA 17:45:35:718 [145c]  snd 1495: <POST /axis/services/urn:cominfo HTTP/
     1.1Connection: closeContent-Length: 1295User-Agent: VordelSOAPAction:
     ""Via: 1.0 devsupport2 (Vordel)Host: devsupport2:7070Content-Type:
     text/xml<?xml version="1.0" encoding="UTF-8" standalone="no"?>
     <soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
     xmlns:xsd="http://www.w3.org/2001/XMLSchema"
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
     <soap:Header>
       <wsse:Security xmlns:wsse="http://docs.oasis-open.org/wss/2004/01/
       oasis-200401-wss-wssecurity-secext-1.0.xsd">
       <wsse:UsernameToken xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/
         oasis-200401-wss-wssecurity-utility-1.0.xsd"
         wsu:Id="Id-00000128d05aca81-00000000009d04dc-10">
              <wsse:Username>joeuser</wsse:Username>
              <wsse:Nonce EncodingType="utf-8">
                  gmP9GCjoe+YIuK1einlENA==
              </wsse:Nonce>
              <wsse:Password Type="http://docs.oasis-open.org/wss/2004/01/
              oasis-200401-wss-username-token-profile-1.0#PasswordText">
                    joepwd
              </wsse:Password>
              <wsu:Created>2010-05-25T16:45:30Z</wsu:Created>
         </wsse:UsernameToken>
         </wsse:Security>
      </soap:Header>
      <soap:Body>
          <ns1:getInfo xmlns:ns1="http://stock.samples">
              <symbol xsi:type="xsd:string">CSCO</symbol>
              <info xsi:type="xsd:string">address</info>
```

```
          </ns1:getInfo>
      </soap:Body>
    </soap:Envelope>
>
```

The following example shows the DATA trace when a message is received by the API Gateway (message starts with rcv):

```
DATA 17:45:35:734 [145c]  rcv 557: <HTTP/1.0 200 OKSet-Cookie: 8Set-Cookie2:
     8Content-Type:
     text/xml; charset=utf-8<?xml version="1.0" encoding="UTF-8"?>
     <soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
     xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/
     XMLSchema-instance">
     <soapenv:Body>
        <ns1:getInfoResponse soapenv:encodingStyle="http://schemas.xmlsoap.org/soap/
        encoding/" xmlns:ns1="http://stock.samples">
            <getInfoReturn xsi:type="xsd:string">San Jose, CA</getInfoReturn>
        </ns1:getInfoResponse>
      </soapenv:Body>
     </soapenv:Envelope>
>
```

If you want to see what has been received by the API Gateway on this thread, run the following command:

```
:g!/145c] rcv/d
```

All snd and rcv trace statements start and end with < and > respectively. If you are assembling a message by hand from the DATA trace, remember to remove characters. In addition, the sending and/or receiving of a single message may span multiple trace statements.

Integrate trace output with Apache log4J

Apache log4j is included on the API Gateway classpath. This is because some third-party products that the API Gateway interoperates with require log4j. The configuration for log4j is found in the API Gateway INSTALL_DIR/system/lib directory in the log4j.properties file.

For example, to specify that the log4j appender sends output to the API Gateway trace file, add the following setting to your log4j.properties file:

```
log4j.rootLogger=DEBUG, A1, Vordel
log4j.appender.Vordel=com.vordel.trace.VordelTraceAppender
```

Get help with API Gateway

Context-sensitive help is available from the Policy Studio screens. Click the **Help** button on any screen to display the relevant help page for that screen. If you require further information or assistance, please contact the Oracle Support team.

Details for Oracle Support
It is important to include as much information as possible when contacting the Oracle Support team. This helps to diagnose and solve the problem in a more efficient manner. The following information should be included with any Support query:

- Name and version of the product (for example, Oracle API Gateway 11.1.2.4.0).
- Details of patches that were applied to the product, if any.

- Platform on which the product is running.
- A clear (step-by-step) description of the problem or query.
- If you have encountered an error, the error message should be included in the email. It is also useful to include any relevant trace files from the `/trace` directory of your product installation, preferably with the trace level set to `DEBUG`.

API Gateway performance tuning

Overview

This topic explains how to optimize API Gateway performance using various configuration options. For example, general performance tuning options include tracing, monitoring, and logging. More advanced performance tuning options include database pooling, HTTP keep alive, chunked encoding, client threads, and Java memory.

Note

This topic applies to performance tuning for both API Gateway and API Gateway Analytics.

General performance tuning

You can optimize API Gateway performance by using Policy Studio to configure the general settings described in this section.

Minimize tracing

The **Trace Log** is displayed in the **Logs** view in the API Gateway Manager web console. When tracing is running at a verbose level (for example, DEBUG), this means that the API Gateway is doing more work and is very dependent on disk input/output. You can set a less verbose trace level for an API Gateway instance or API Gateway port interface (for example, ERROR or FATAL).

To set the tracing for an API Gateway instance, select **Server Settings** > **General** in the Policy Studio tree, and select the **Trace Level** (for example, FATAL):

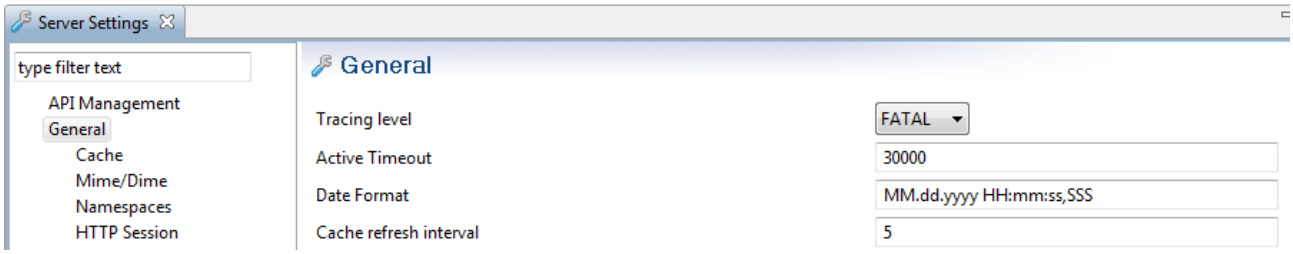

You can also override the trace level for an API Gateway port interface, and set it to a quieter level. For example, in the Policy Studio tree, select **Listeners** > **API Gateway** > **Sample Services** > **Ports**. Right-click an interface in the list on the right, select **Edit**, and set the **Trace level** (for example, FATAL):

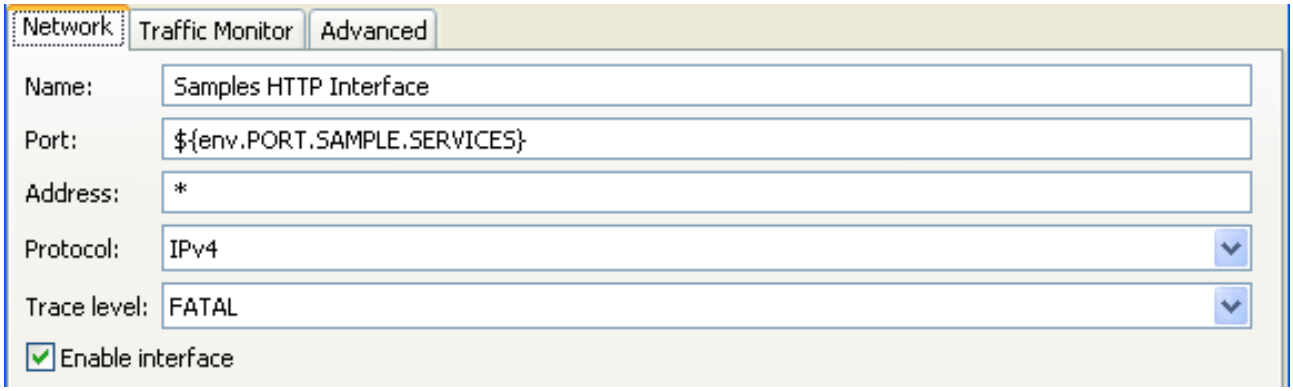

For more details, see *Configure API Gateway diagnostic trace*

Disable real-time monitoring

Real-time monitoring is displayed in the **Monitoring** view in the API Gateway Manager web console. This caches recent message transactions in the memory of the API Gateway. You can remove this overhead by disabling real-time monitoring.

To disable in Policy Studio, select **Server Settings** > **Monitoring** > **Real Time Monitoring**, and deselect **Enable Real Time Monitoring**:

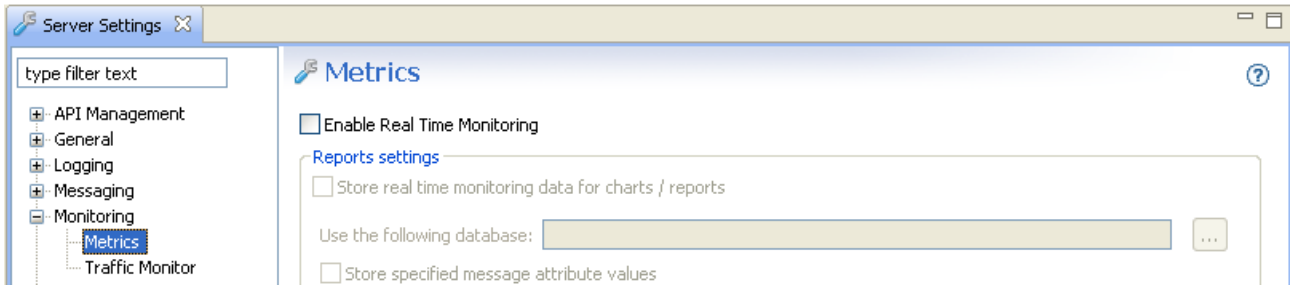

For more details, see *Real-time monitoring metrics*

Disable traffic monitoring

Traffic monitoring is displayed in the **Traffic** view in the API Gateway Manager web console. By default, the API Gateway stores recent HTTP traffic summaries to the API Gateway disk for use in API Gateway Manager. You can remove this overhead by disabling traffic monitoring.

To disable in Policy Studio, select **Server Settings** > **Monitoring** > **Traffic**, and deselect **Enable Traffic Monitor**:

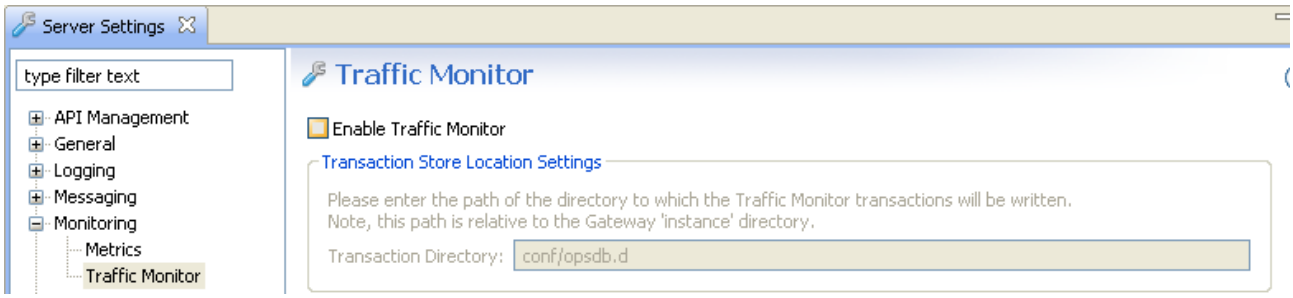

For more details, see *Traffic monitoring settings*

Disable transaction logging

The **Transaction Log** is displayed in the **Logs** view in the API Gateway Manager web console. You should ensure that the API Gateway is not sending transaction log messages or events to transaction log destinations. This is because the performance of the API Gateway will be determined by the log destination.

To disable transaction logging in the API Gateway, you must disable all log destinations in Policy Studio. For example, select **Server Settings** > **Logging** > **Transaction Log**, and deselect **Enable logging to a file**. The following example shows disabling logging to file, you must perform this step in all tabs on this screen:

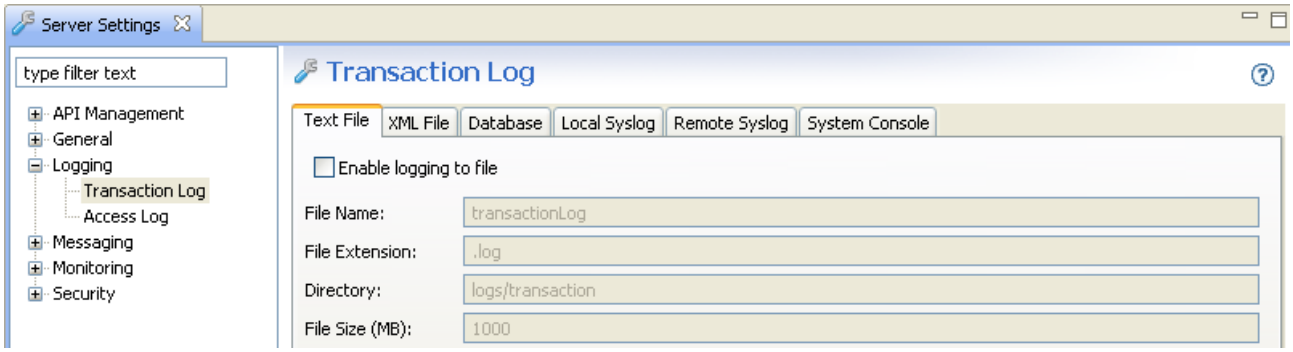

For more details, see *Transaction audit log settings*

Disable Access logging

You should also ensure that the API Gateway is not sending log messages to the access log. To disable access logging in the API Gateway, select **Server Settings** > **Logging** > **Transaction Access Log**, and deselect **Transaction Access Log Enabled**:

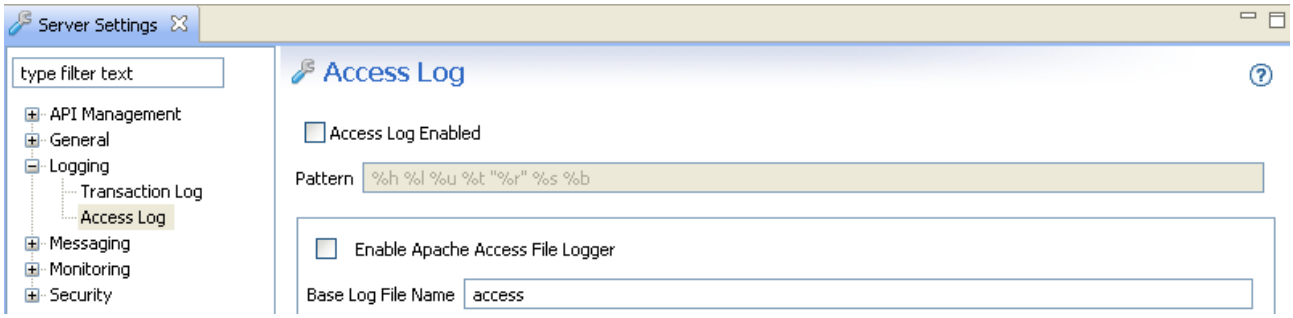

For more details, see *Transaction access log settings*

Advanced performance tuning

You can also use the advanced configuration settings described in this section to optimize API Gateway performance.

Configure spill to disk

When stress testing with large messages (greater than 4 MB), the API Gateway spills data to disk instead of holding it in memory. By default, the `spilltodisk` option is triggered with payload sizes of 4 MB or more. For example, you can configure this in the `service.xml` file in the following directory by adding the `spilltodisk` option configured in bytes:

```
<install-dir>/apigateway/groups/<group-id>/<instance-id>/conf/service.xml
```

For example:

```
<NetService provider="NetService">
    <SystemSettings
            tracelevel="&server.tracelevel;"
            tracecomponent="&server.title;"
            title="&server.title;"
            homedir="$VINSTDIR"
            secret="&server.entitystore.secret;"
            servicename="&server.servicename;"
            spilltodisk="10485760"
            ...
```

This setting specifies the limit of what is considered a reasonably-sized message to hold in memory. After this limit is exceeded, to preserve memory, the system writes the content of the incoming request to disk when it arrives.

Configure database pooling

The API Gateway uses Apache Commons Database Connection Pools (DBCP) for pooling database connections. For details, see http://commons.apache.org/dbcp/.

Note

Oracle recommends that if your policy interacts heavily with the database, the pool size for the database connection should be at least as big as the expected client population. This assumes the database can cope with this number of parallel connections. For example, if you are providing load from 100 parallel clients, the pool settings shown in the following example are recommended.

To configure database pooling in Policy Studio, select **External Connections** > **Database Connections** > **Add Database Connection**. For example:

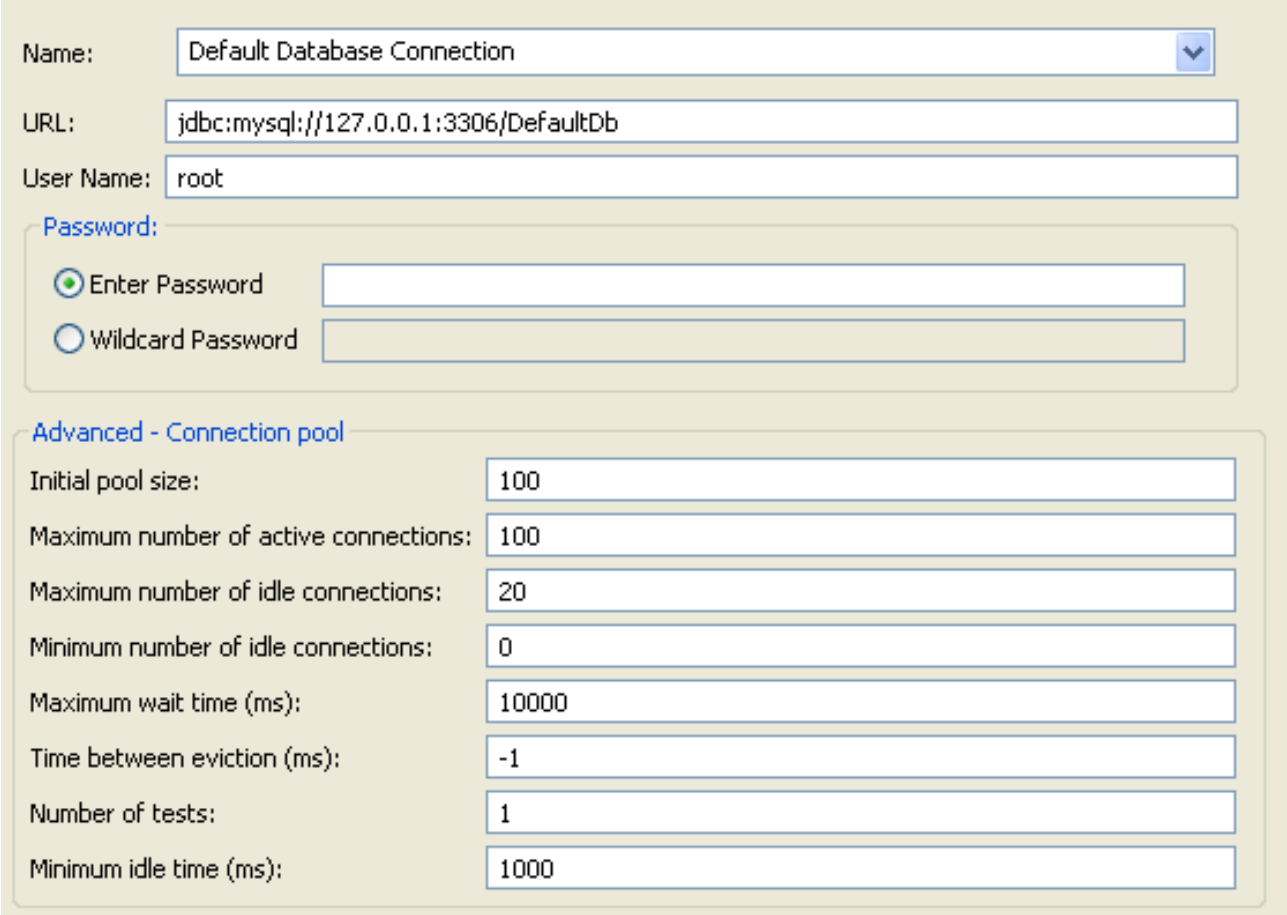

Configure HTTP keep alive

In HTTP/1.1, by default, the connection between a client and a server is maintained so that further client requests can avoid the overhead of setting up a new connection. This may or may not model the client population of a particular scenario very well. If it is acceptable to reuse TCP connections (and SSL connections on top of these), ensure your client uses HTTP/1.1, and does not opt out of the HTTP keepalive mechanism.

For the `sr` command, this means you should use the `-V1.1` and `-U1000` arguments to enable the connection be used a number of times before closing it. For details on `sr`, see the *API Gateway Policy Developer Guide*.

Note

For conformance with the HTTP/1.1 specification, the client must send a `Host` header in this configuration, so you must pass a further `-aHost:localhost` argument to `sr`. If the connection keep alive is working correctly, `sr` reports a larger number of transactions to connections in its periodic output.

Configure chunked encoding

For interoperability reasons, the API Gateway normally does not use chunked encoding when talking to a remote server. Because of the HTTP protocol, the API Gateway must send the `Content-Length` header to the server, and so must precompute the exact size of the content to be transmitted. This may be expensive.

For example, when relaying data directly from client to server, or when the message exists as an abstract XML document in the API Gateway, the size may not be immediately available. This means that the entire content from the client must be buffered, or the internal body structure must be serialized an extra time just to measure its size. By configuring a remote host for the destination server to allow HTTP 1.1, when a server is known to be advertising HTTP 1.1, chunked encoding can be used when transmitting to that server where appropriate.

For example, in the Policy Studio tree, select **Listeners** > **API Gateway** > **Sample Services**. Right-click the remote host, select **Edit**, and select **Allow HTTP 1.1**:

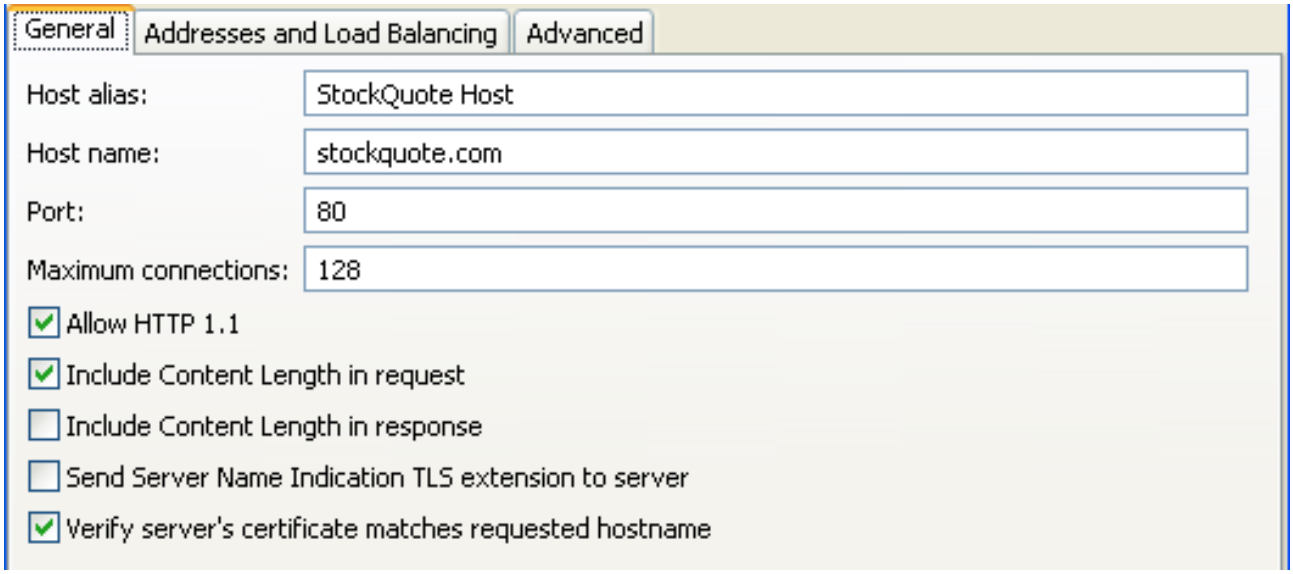

Single test client and sever

If a performance test uses one client and one target service, it is recommended that you use the following settings:

```
sysctl net.ipv4.tcp_tw_recycle=1
sysctl net.ipv4.tcp_tw_reuse=1
```

A TCP/IP connection is identified by the following properties:

- Remote IP address
- Local IP address
- Remote port
- Local port

When you have a fixed client IP address, a fixed server IP address, and a fixed server port, this leaves 65536 client ports, which means you can only have 65536 live connections at a time.

In a TCP/IP connection, one side always ends up in the `TIME_WAIT` state, which is part of the protocol. This period can last for two minutes by default. In a naive implementation, this means you can establish at most 65536 connections every 4 minutes. There are ways to mitigate this in the protocol, and most TCP stacks can reuse connections in the `TIME_WAIT` state more rapidly than this. But under certain circumstances, this can still be a problem. The `sysctl` settings shown above make the Linux kernel more aggressive in the way it recycles `TIME_WAIT` connections.

JVM memory

If processing many large XML documents, you should reduce the Java Virtual Machine (JVM) memory so that garbage collection is performed more often. Create a file called `jvm.xml` in the `conf` directory of the API Gateway installation with the following contents:

```
<ConfigurationFragment>
    <VMArg name="-Xmx100m"/>
    <VMArg name="-Xms100m"/>
    <VMArg name="-Xincgc"/>
    <Print message="JVM heap limited"/>
</ConfigurationFragment>
```

The amount of memory used by the server is limited in several ways, and one of these limits is the normal Java *heap space*. The API Gateway can use memory that is not visible to the JVM, but is *pinned* by it, so it is not released until the JVM decides to collect its garbage. If the amount of memory pinned is much larger than the size of the objects that the JVM sees (the bits pinning it down), the Java heap setting does not represent the amount of heap memory used. So the problem is worked around by reducing the Java heap space used in this scenario.

For example, by setting the Java heap size to 100 MB instead of 1024 MB, the Java garbage collector runs when the Java heap use is close to 100 MB, while the amount of memory physically pinned by the JVM may be closer to 1000 MB.

In addition, you may wish to consider reducing the number of threads that the API Gateway uses for processing incoming messages. You can do this by adding a `maxThreads="64"` attribute in the `SystemSettings` element in the `service.xml` file for the API Gateway instance (for example, in `..\apigateway\groups\group-2\instance-1`). This setting also helps in configuring the API Gateway to back off from the target service (if the API Gateway can process more load than the target service).

Number of client threads on Linux

If there are more than 200 client threads connecting, you must change the following line in the `venv` script in the `posix/bin` directory of your API Gateway installation:

```
ulimit -n unlimited ...
```

to the following:

```
ulimit -n 131072 ...
```

This prevents the too many open files error that you might see in a 200 client thread test.

Multiple connection filters

This applies if the performance test involves more than one connection filter, where the filter is routing to the same host and port with the same SSL credentials. You should create a policy containing the single connection filter, and delegate to this policy from where you normally do the routing, so you delegate to a single filter instead of a connection filter per policy.

Each connection processor caches connections independently. This is because two connection processors using different SSL certificates cannot pool their connections. They are not interchangeable from an authentication point of view. Therefore, when using multiple connection filters, there is potential to soak the target machine with too many connections.

Note

This applies to both API Gateway **Connection** and **Connect to URL** filters in Policy Studio.

Manage API Gateway users

Overview

The **Users and Groups** node in the Policy Studio tree enables you to manage API Gateway users and groups, which are stored in the API Gateway user store.

By default, the API Gateway user store contains the configuration data for managing API Gateway user information. The API Gateway user store is typically used in a development environment, and is useful for demonstration purposes. In a production environment, user information may be stored in existing user Identity Management repositories such as Microsoft Active Directory, Oracle Access Manager, CA SiteMinder, and so on. For more details, see the *API Gateway Integration Guide*.

Note

API Gateway users provide access to the messages and services protected by API Gateway. However, *Admin users* provide access to the API Gateway configuration management features available in Policy Studio, Configuration Studio, and API Gateway Manager. For more details, see *Manage Admin users*.

API Gateway users

API Gateway users specify the user identity in the API Gateway user store. This includes details such as the user name, password, and X.509 certificate. API Gateway users must be a member of at least one user group. In addition, users can specify optional attributes, and inherit attributes at the group level.

To view all existing users, select the **Users and Groups** > **Users** node in the tree. The users are listed in the table on the main panel. You can find a specific user by entering a search string in the **Filter** field.

Add API Gateway users

You can create API Gateway users on the **Users** page. Click the **Add** button on the right.

Add user details
To specify the new user details, complete the following fields on the **General** tab:

- **User Name**:
  Enter a name for the new user.
- **Password**:
  Enter a password for the new user.
- **Confirm Password**:
  Re-enter the user's password to confirm.
- **X.509 Cert**:
  Click the **X.509 Cert** button to load the user's certificate from the **Certificate Store**.

Add user attributes
You can specify optional user attributes on the **Attributes** tab, which is explained in the next section.

API Gateway user attributes

You can specify attributes at the user level and at the group level on the **Attributes** tab. Attributes specify user configuration data (for example, attributes used to generate SAML attribute assertions).

Add attributes

The **Attributes** tab enables you to configure user attributes as simple name-value pairs. The following are examples of user attributes:

- `role=admin`
- `email=niall@oracle.com`
- `dept=eng`
- `company=oracle`

You can add user attributes by clicking the **Add** button. Enter the attribute name, type, and value in the fields provided. The `Encrypted` type refers to a string value that is encrypted using a well-known encryption algorithm or cipher.

API Gateway user groups

API Gateway user groups are containers that encapsulate one or more users. You can specify attributes at the group level, which are inherited by all group members. If a user is a member of more than one group, that user inherits attributes from all groups (the superset of attributes across the groups of which the user is a member).

To view all existing groups, select the **Users and Groups** > **Groups** node in the tree. The user groups are listed in the table on the main panel. You can find a specific group by entering a search string the **Filter** field.

Add API Gateway user groups

You can create user groups on the **Groups** page. Click the **Add** button on the right to view the **Add Group** dialog.

Add group details
To specify the new group details, complete the following fields on the **General** tab:

- **Group Name**:
  Enter a name for the new group.
- **Members**:
  Click the **Add** button to display the **Add Group Member** dialog, and select the members to add to the group.

Add group attributes
You can specify optional attributes at the group level on the **Attributes** tab. For more details, see the section called “API Gateway user attributes”.

Update API Gateway users or groups

To edit details for a specific user or group, select it in the list, and click the **Edit** button on the right. Enter the updated details in the **Edit User** or **Edit Group** dialog.

To delete a specific user or group, select it in the list, and click the **Remove** button on the right. Alternatively, to delete all users or Groups, click the **Remove All** button. You are prompted to confirm all deletions.

Manage Admin users

Overview

When logging into the Policy Studio or API Gateway Manager, you must enter the user credentials stored in the local Admin user store to connect to the API Gateway server instance. Admin users are responsible for managing API Gateway instances using the API Gateway management APIs. To manage Admin users, click the **Settings > Admin Users** tab in the API Gateway Manager.

Note

Admin users provide access to the API Gateway configuration management features available in the Policy Studio and API Gateway Manager. However, *API Gateway users* provide access to the messages and services protected by the API Gateway. For more details, see *Manage API Gateway users*.

Admin user privileges

After installation, a single Admin user is defined in the API Gateway Manager with a user name of `admin`. Admin user rights in the system include the following:

- Add another Admin user.
- Delete another Admin user.
- Update an Admin user.
- Reset Admin user passwords.

Important

An Admin user *cannot* delete itself.

Remove the default Admin user
To remove the default Admin user, perform the following steps:

1. Add another Admin user.
2. Log in as the new Admin user.
3. Delete the default Admin user.

The **Admin Users** tab displays all existing Admin users. You can use this tab to add, update, and delete Admin users. These tasks are explained in the sections that follow.

Admin user roles

The API Gateway uses Role-Based Access Control (RBAC) to restrict access to authorized users based on their assigned roles in a domain. Using this model, permissions to perform specific system operations are assigned to specific roles only. This simplifies system administration because users do not need to be assigned permissions directly, but instead acquire them through their assigned roles.

For example, the default Admin user (`admin`) has the following user roles:

- `Policy Developer`
- `API Server Administrator`

- `KPS Administrator`

User roles and privileges
User roles have specific tools and privileges assigned to them. These define who can use which tools to perform what tasks. The user roles provided with the API Gateway assign the following privileges to Admin users with these roles:

Role	Tool	Privileges
`API Server Administrator`	API Gateway Manager	Read/write access to API Gateway Manager.
`API Server Operator`	API Gateway Manager	Read-only access to API Gateway Manager.
`Deployer`	Deployment scripts	Deploy a new configuration.
`KPS Administrator`	KPS Web UI	Perform create, read, update, delete (CRUD) operations on data in a Key Property Store (KPS).
`Policy Developer`	Policy Studio	Download, edit, deploy, version, and tag a configuration.

Note

A single Admin user typically has multiple roles. For example, in a development environment, a policy developer Admin user would typically have the following roles:

- `Policy Developer`
- `API Server Administrator`

Add a new Admin user

Complete the following steps to add a new Admin user to the system:

1. Click the **Settings > Admin Users** tab in the API Gateway Manager.
2. Click the **Create** button.
3. In the **Create New Admin User** dialog, enter a name for the user in the **Username** field.
4. Enter a user password in the **Password** field.
5. Re-enter the user password in the **Confirm Password** field.
6. Select roles for the user from the list of available roles (for example, `Policy Developer` and/or `API Server Administrator`).
7. Click **Create**.

Remove an Admin user

To remove an Admin user, select it in the **Username** list, and click the **Delete** button. The Admin user is removed from the list and from the local Admin user store.

Reset an Admin user password

You can reset an Admin user password as follows:

1. Select the Admin user in the **Username** list.
2. Click the **Edit** button.
3. Enter and confirm the new password in the **Password** and **Confirm Password** fields.
4. Click **OK**.

Manage Admin user roles

You can manage the roles that are assigned to specific Admin users as follows:

1. Select the Admin user in the **Username** list.
2. Click the **Edit** button.
3. Select the user roles to enable for this Admin user in the dialog (for example, `Policy Developer` and/or `API Server Administrator`).
4. Click **OK**.

Edit roles
To add or delete specific roles, you must edit the available roles in the `adminUsers.json` and `acl.json` files in the `conf` directory of your API Gateway installation.

For more details on role-based access, see *Configure Role-Based Access Control (RBAC)*. For more details on role-based access, see the *API Gateway Administrator Guide*.

Configure Role-Based Access Control (RBAC)

Overview

Role-Based Access Control (RBAC) enables you to restrict system access to authorized users based on their assigned roles. Using the RBAC model, permissions to perform specific system operations are assigned to specific roles, and system users are granted permission to perform specific operations only through their assigned roles. This simplifies system administration because users do not need to be assigned permissions directly, and instead acquire them through their assigned roles.

The API Gateway uses the RBAC permissions model to ensure that only users with the assigned role can access parts of the management services exposed by the Admin Node Manager. For example, this includes access to traffic monitoring data or making a configuration change by deploying to a group of API Gateways. The following diagram shows an overview of the RBAC model in the API Gateway:

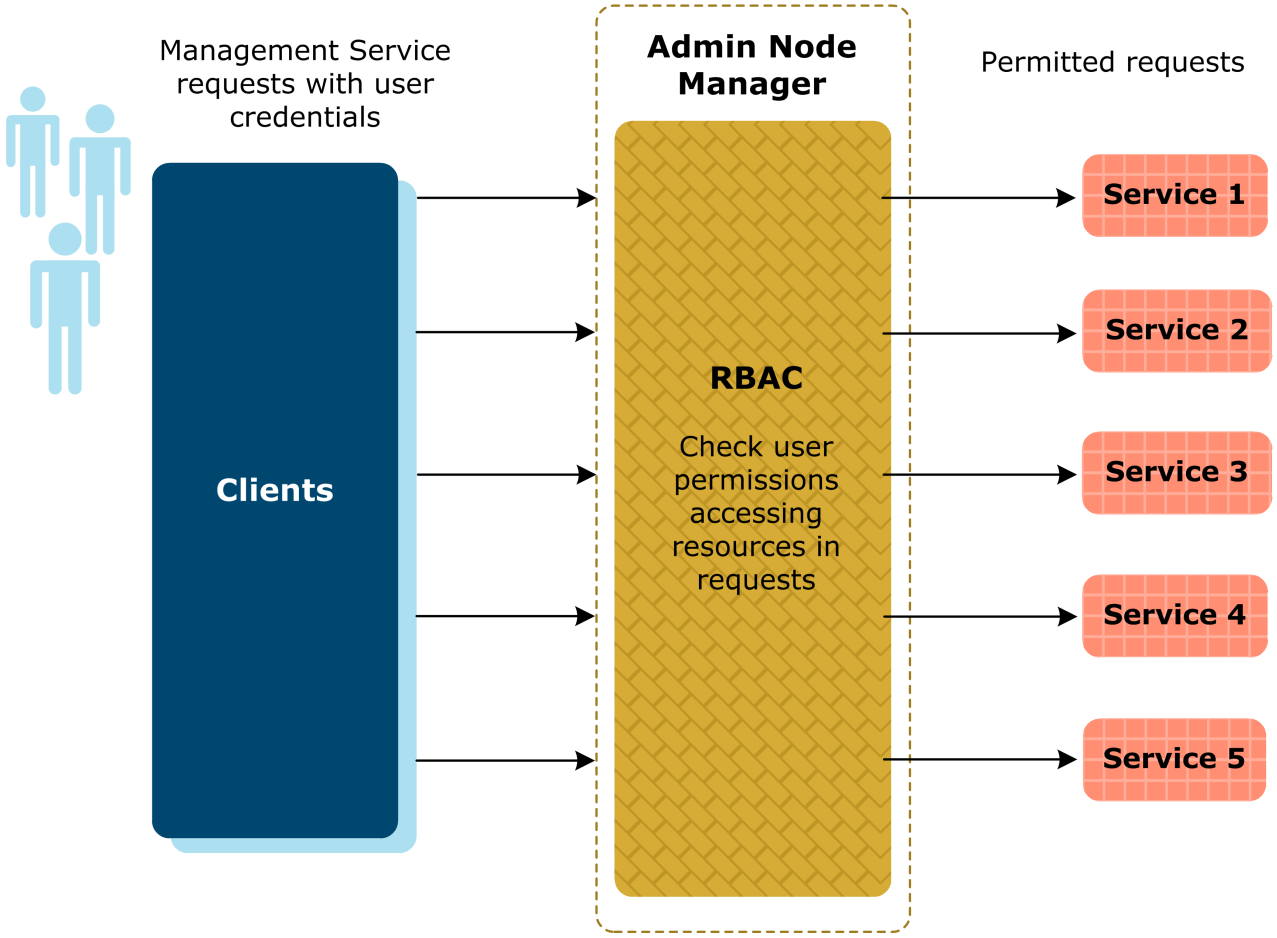

Note

This topic applies to both API Gateway and API Gateway Analytics.

API Gateway Manager
The web-based API Gateway Manager tool (`https://localhost:8090`) is a centralized dashboard for managing and monitoring the API Gateway, and is controlled by RBAC. Users connecting to this URL with different roles results in different features being displayed.

For example, users with the `API Server Administrator` role has read/write access in the API Gateway Manager tool. However, users in the `API Server Operator` have only read access in the API Gateway Manager tool.

For more details on the tools and privileges assigned to specific user roles, see the topic on *Manage Admin users*.

Protected management services
The Admin Node Manager exposes a number of REST management services, which are all protected by RBAC. For example, the exposed services and the associated tools that use them include the following:

Protected Service	Tool	Description
Traffic Monitoring Service	API Gateway Manager	Displays HTTP, HTTPS, JMS, and FTP message traffic processed by the API Gateway.
Configuration Service	API Gateway Manager	Adds and removes tags on the API Gateway.
Topology API	API Gateway Manager	Accesses and configures API Gateway domains.
Static Content Resources	API Gateway Manager	Manages UI elements in a browser.
Deployment API	Policy Studio	Deploys configurations to the API Gateway.
KPS Service	Policy Studio	Manages a Key Property Store.

RBAC user roles
User access to management services is determined by their role(s). Each role has a defined set of management services that it can access. A Management Service is defined by the URI used to access it, for example:

Role Name	Service Name	API Type	Example URI
`API Server Operator`	Topology API	REST	`/api/topology/hosts`

For full details on the default roles that have access to each Management Service, see the section called “Management service roles and permissions”.

Local admin user store

By default, all the user credentials are stored in a local admin user store in the following file:

```
INSTALL_DIR/conf/adminUsers.json
```

`INSTALL_DIR` is the directory where the API Gateway is installed as Admin Node Manager.

The following shows an example file:

```
{
  "productVersion" : "11.1.2.4.0",
  "version" : 1,
```

```
  "timestamp" : 0,
  "adminUsers" : [ {
    "id" : "user-1",
    "name" : "admin",
    "roles" : [ "role-1", "role-6", "role-7" ]
  } ],
  "credentials" : {
    "user-1" : "$ygvfNZsiLEcQJGMe8vqmqw==$2jcZcvFctadAISBuvhJqt9NfUtk7zeoOHo8cFMySjvw="
  },
  "adminUserRoles" : [ {
    "id" : "role-1",
    "name" : "API Server Administrator"
  }, {
    "id" : "role-2",
    "name" : "API Server Operator"
  }, {
    "id" : "role-5",
    "name" : "Deployer"
  }, {
    "id" : "role-6",
    "name" : "KPS Administrator"
  }, {
    "id" : "role-7",
    "name" : "Policy Developer"
  } ],
  "uniqueIdCounters" : {
    "Role" : 8,
    "User" : 2
  },
  "adminUsersVersion" : {
    "version" : 1,
    "timestamp" : 0
  }
}
```

The credentials from this file are used to authenticate and perform RBAC on all accesses to the management services. This store holds the user credentials, so their passwords can be verified, and also holds their roles. Credentials and associated roles can also be retrieved from an LDAP Directory Server (for example, Microsoft Active Directory or OpenLDAP).

Admin user password storage algorithm

Admin user passwords are in stored in the `adminUsers.json` file as a base64-encoded salted hash of the plaintext password. To authenticate a user password, the incoming password value is digested in the same way, and compared against the value in the file. If they match, the user is authenticated.

The salt is a 16 byte value generated using the SHA1PRNG pseudo-random number generator algorithm. This algorithm is provided by the Java Cryptography Extension (JCE) and is PBKDF2 with HMAC SHA1. The algorithm repeats the digest of the password along with the salt for 1024 iterations. The salt is stored with the digest allowing for password verification in the following format:

```
$<salt-value>==$<password-digest>=
```

For example:

```
$d5pSA0kJ24USYWb19b3fyw==$+NuXNIEVIZInirQmD8HI9XinrqbPcmG0G6wgNY/8CH0=
```

Because the salt is generated each time a password is stored, the same plaintext password will have a different salt and digest value. For example, the following stored passwords are all digests of the default password:

```
$rLZ8yVIDVwN/NMcuxp5B+Q==$zJBpjQH0zGHhr2uR05YTOk+hMxK1w1xyzk4J1IR95Qs=
$myCJnKEHRB5DMyyfaDA+Kw==$YKqMxFagu9kpqBP/Uug4UcwCEJgRG2nFm2RrJQnwOBc=
$hazgC9Kf/93zS0V2VvOx6Q==$5LEpsXQNSph/ZVbSpJ5UwJUo0JU9nKPSWpii+6pzxI4=
```

Configure an LDAP repository to store credentials

Alternatively, for details on how to configure an LDAP repository to store user credentials, see the following topics:

- *Active Directory for authentication and RBAC of management services*
- *OpenLDAP for authentication and RBAC of management services*

RBAC Access Control List

The Access Control List file (`acl.json`) is located in the `conf` directory of your API Gateway installation. This file lists each role and the management services that each role may access. By default, this file defines the following roles:

- `API Server Administrator`
- `API Server Operator`
- `Deployer`
- `KPS Administrator`
- `Policy Developer`

The default `admin` user is assigned the `API Server Administrator`, `KPS Administrator`, and `Policy Developer` roles by default, which together allow access to everything. For full details on the management services that each role has access to, and the permissions that must be listed in the `acl.json` file to have access to them, see the table in the section called “Management service roles and permissions”.

Important

The roles defined in the `acl.json` file should exist in the user store used to authenticate the users and load their roles and/or groups. The default roles are defined in the local Admin User store, which is used to control access to the management services using the **Protect Management Interfaces** policy. If a different user store is used (for example, an LDAP repository), the LDAP groups should be listed in the `acl.json` and `adminUsers.json` files .

Access Control List file format
Each role entry in the `acl.json` file has the following format:

```
"role-name" : [ <list_of_permission_names> ]
```

The permissions consist of operations that are defined by HTTP methods and URIs:

```
“permission-name” : { <list_of_operation_names> }
“operation-name” : {
        "methods" : [ <list of HTTP Methods> ],
        "paths" : [ <list of path-names> ]
}

“path-name” : {
        "path" : <URI>
}
```

This file entry format is described as follows:

- The permissions line is repeated for each permission the role has. To determine which permissions should be listed for each Management Service, see the table in the section called “Management service roles and permissions”.
- You can place a wildcard (*) at the end of the `path` field. For example, see the path for `dojo resources` in the example that follows. This means the role has access to all URIs that start with the URI content that precedes the *.
- In some cases, you must protect a Management Service by specifying a query string after the URI. Exact matches only are supported for query strings.

Example Access Control List file
The following example shows roles and permissions to URIs:

```
"paths" : {
        "root" : { "path" : "/" },
        "emc pages" : { "path" : "/emc/*" },
        "site images" : { "path" : "/images/*" },
        "dojo resources" : { "path" : "/dojo/*" },
        ....
        }

},

"operations" : {
        "emc_read_web" : {
        "methods" : [ "GET" ],
        "paths" : [ "emc pages", "dojo resources" ]
        },

        "common_read_web" : {
        "methods" : [ "GET" ],
        "paths" : [ "root", "site images" ]
        },
        ....
},

"permissions" : {
        "emc" : [ "common_read_web", "emc_read_web" ],
        "config" : [ "configuration" ],
        "deploy" : [ "deployment", "management" ],
        ...
},

“roles” : {
        "Policy Developer" : [ "deploy", "config"]
        ...
}
```

Configure RBAC users and roles

You can use the API Gateway Manager to configure the users and roles in the local Admin User store. Click the **Settings** > **Admin Users** to view and modify user roles (assuming you have a role that allows this). This screen is displayed as follows:

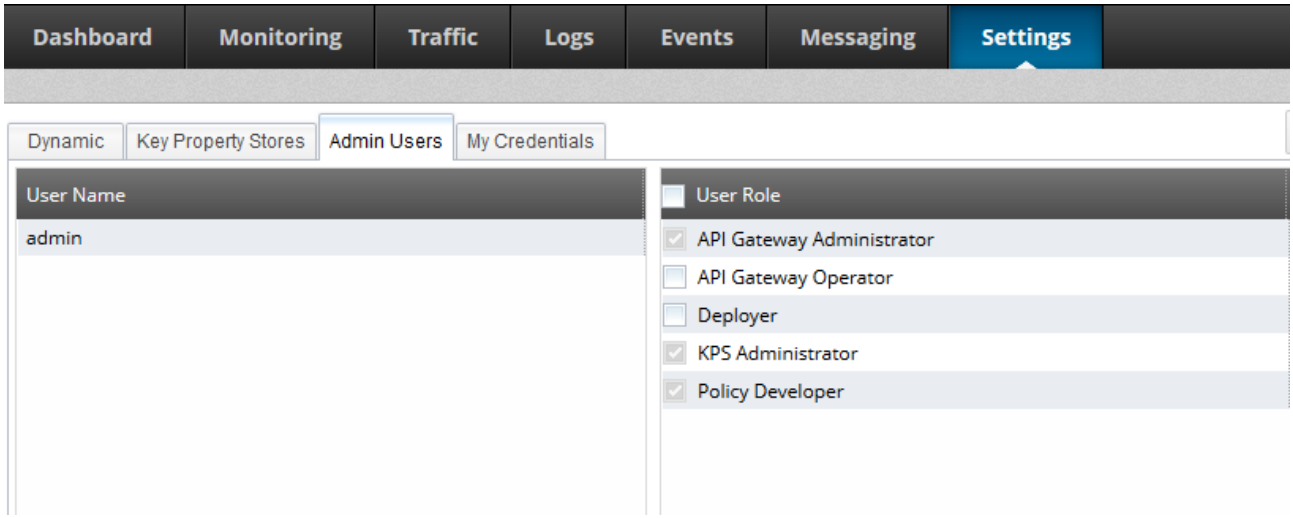

Manage RBAC user roles
When you click **Create** to create a new user, you can select the roles to assign to the that new user. New users are not assigned a default role. While users that are replicated from an LDAP repository do not require a role to be assigned to them. You can click **Edit** to changed the roles assigned to a selected user.

Add a new role to the user store
When you add a new role to the Admin User store, you must modify the available roles in the `adminUsers.json` and `acl.json` files in the `conf` directory of your Admin Node Manager installation. You must add the new role to the `roles` section of the `acl.json` file, which lists all the permissions that the new role may have.

Important

You must update the `acl.json` before you add the roles to the Admin User store. The RBAC policy object automatically reloads the `acl.json` file each time you add or remove a role in the Policy Studio.

When you update the `acl.json` file, you must restart the Admin Node Manager to reload the `acl.json` file. However, the Admin Node Manager does not need to be rebooted or refreshed if a user's roles change.

For more details on managing user roles, see the topic on *Manage Admin users*.

Management service roles and permissions

You can use the following table for reference purposes when making changes to the `acl.json` file. It defines each Management Service, and the default roles that have access to them. It also lists the permissions that must be listed in the `acl.json` file to have access to the Management Service.

Management Service	Default Roles	Permissions
API Gateway Manager (`https://localhost:8090`)	• API Server Administrator • API Server Operator	• `emc` • `mgmt`
API Gateway Manager Dashboard	• API Server Administrator	• `emc` • `mgmt`

Management Service	Default Roles	Permissions
		• `mgmt_modify` • `dashboard` • `dashboard_modify` • `deploy` • `config`
API Gateway Manager Dashboard (read-only access)	• API Server Operator	• `emc` • `mgmt` • `dashboard` • `dashboard_modify`
API Gateway Manager Monitoring	• API Server Administrator • API Server Operator	• `emc` • `mgmt` • `monitoring` • `events` • `traffic_monitor` • `settings` • `settings_modify` • `logs`
API Gateway Manager Traffic	• API Server Administrator • API Server Operator	• `emc` • `mgmt` • `traffic_monitor`
API Gateway Manager Logs	• API Server Administrator • API Server Operator	• `emc` • `mgmt` • `logs`
API Gateway Manager Events	• API Server Administrator • API Server Operator	• `emc` • `mgmt` • `monitoring` • `events`
API Gateway Manager Settings	• API Server Administrator	• `emc` • `mgmt` • `mgmt_modify` • `settings` • `settings_modify`
API Gateway Manager Settings (read-only access)	• API Server Operator	• `emc` • `mgmt`

Management Service	Default Roles	Permissions
		• `settings`
Documentation	• API Server Administrator • API Server Operator	• `emc` • `mgmt`
KPS	• KPS Administrator	• `mgmt` • `kps`
Policy Studio	• Policy Developer	• `mgmt` • `deploy` • `config`
API Gateway Configuration Deployment	• API Server Administrator • Policy Developer • Deployer	• `mgmt` • `deploy` • `config`

Active Directory for authentication and RBAC of management services

Overview

This topic explains how to use Local Directory Access Protocol (LDAP) to authenticate and perform Role-Based Access Control (RBAC) of management services. You can use the sample **Protect Management Interfaces (LDAP)** policy instead of the **Protect Management Interfaces** policy. This means that an LDAP repository is used instead of the local Admin User store for authentication and RBAC of users attempting to access management services. This topic describes how to configure the server to use an example Microsoft Active Directory LDAP repository.

Note

This topic applies to both API Gateway and API Gateway Analytics.

Step 1: create an Active Directory group

To create a new user group in Active Directory, perform the following example steps:

1. Click **Start** > **Administrative Tools** > **Active Directory Users and Computers**.
2. On the **Users** directory, right-click, and select **New** > **Group**.
3. Enter the Group name (for example, `APIGatewayAdministrator`).

You should add groups for the following default RBAC roles to give the LDAP users appropriate access to the API Gateway management services:

- `API Gateway Administrator`
- `API Gateway Operator`
- `Deployer`
- `KPS Administrator`
- `Policy Developer`

These RBAC roles are located in the `roles` section of the following file:

```
INSTALL_DIR\apigateway\conf\acl.json
```

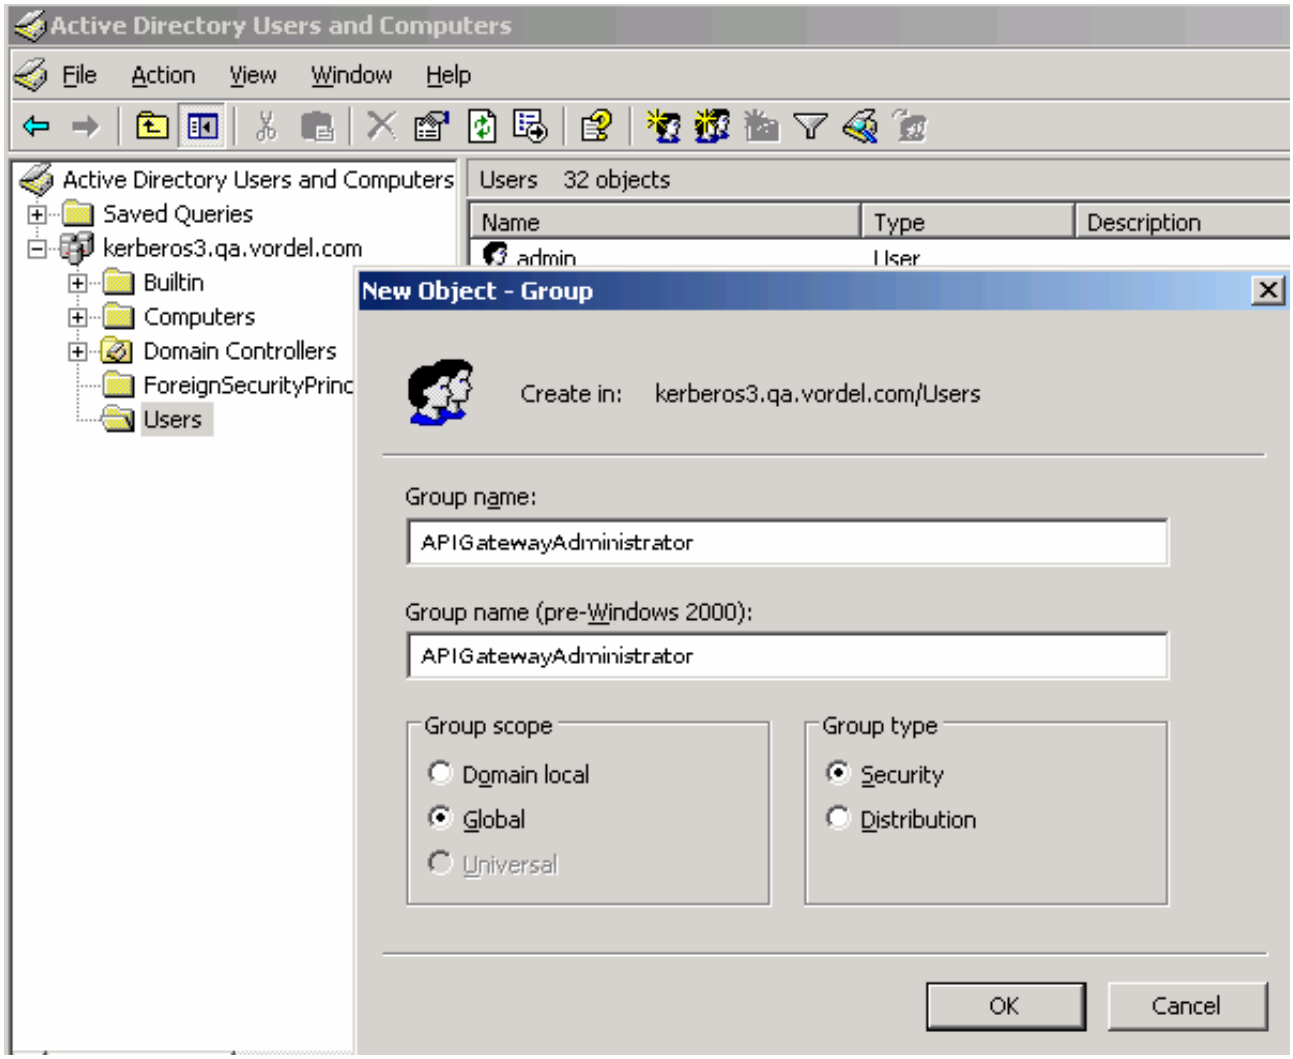

You can view the newly created groups using an LDAP Browser.

Step 2: create an Active Directory user

You will most likely be unable to create an `admin` user with a default password because the default password is not strong enough to be accepted by Active Directory. Using **Active Directory Users and Computers**, perform the following steps:

1. On the **Users** directory, right-click, and select **New** > **User**.
2. Enter a user name (for example, `admin`).

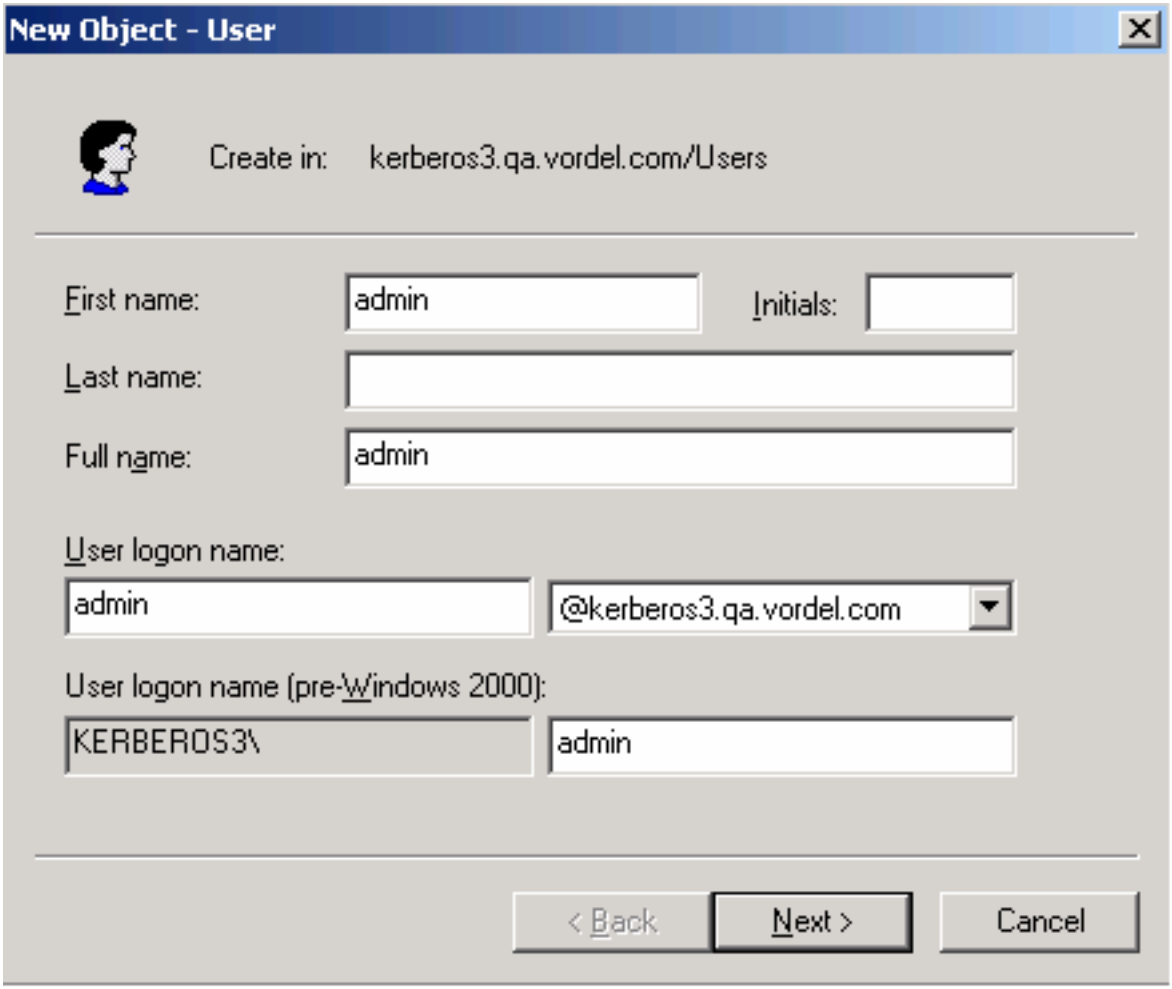

3. Click **Next**.
4. Enter a password (for example, `Oracle123`).
5. Select **User cannot change password** and **Password never expires**.
6. Ensure **User must change password at next logon** is not selected.
7. Click **Next**.
8. Click **Finish**.

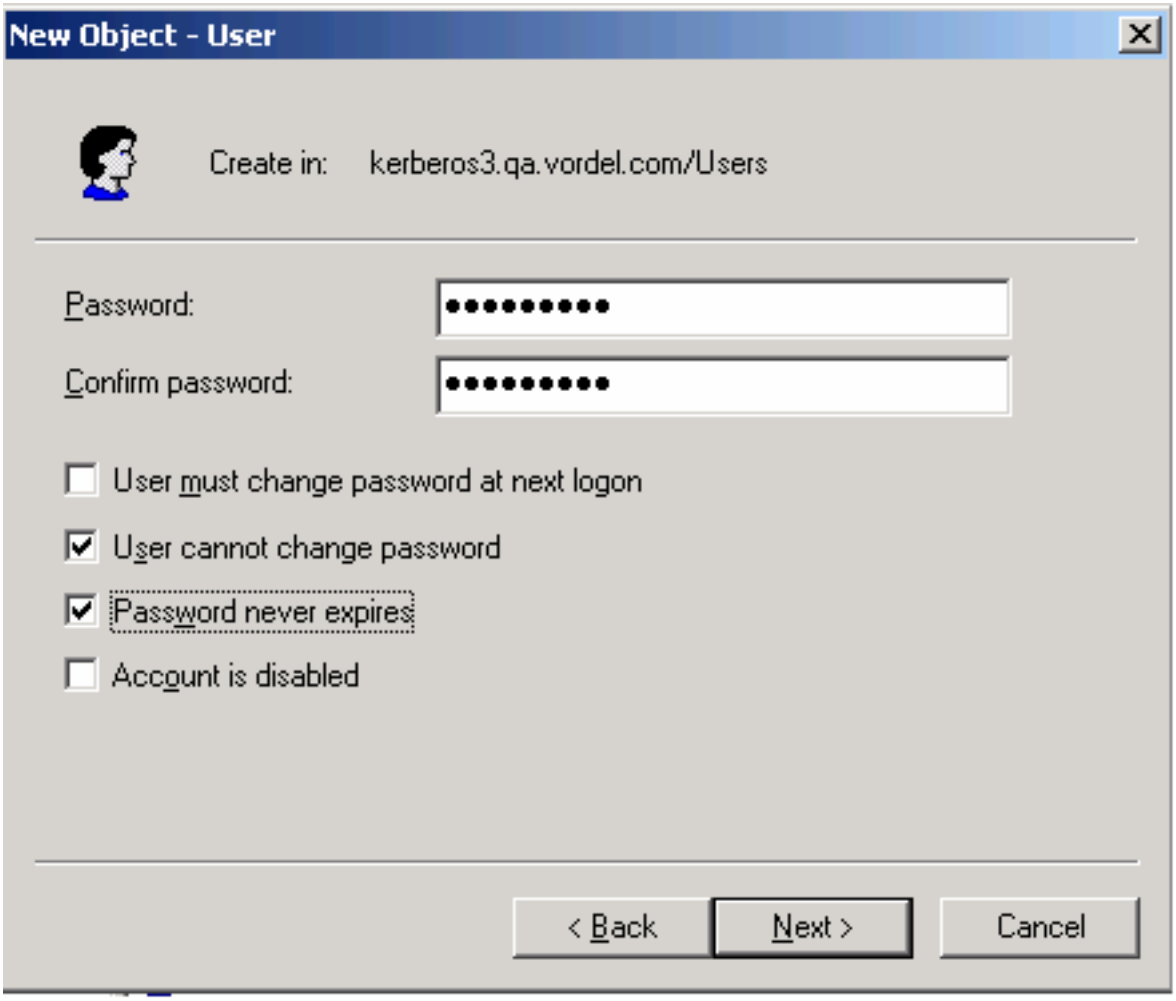

Adding the User to the Group
To make the user a member of the group using **Active Directory Users and Computers**, perform the following steps:

1. Select the **APIGatewayAdministrator** group, right-click, and select **Properties**.
2. Click the **Members** tab.
3. Click **Add**.
4. Click **Advanced**.
5. Click **Find Now**.
6. Select the `admin` user.
7. Click **OK**.

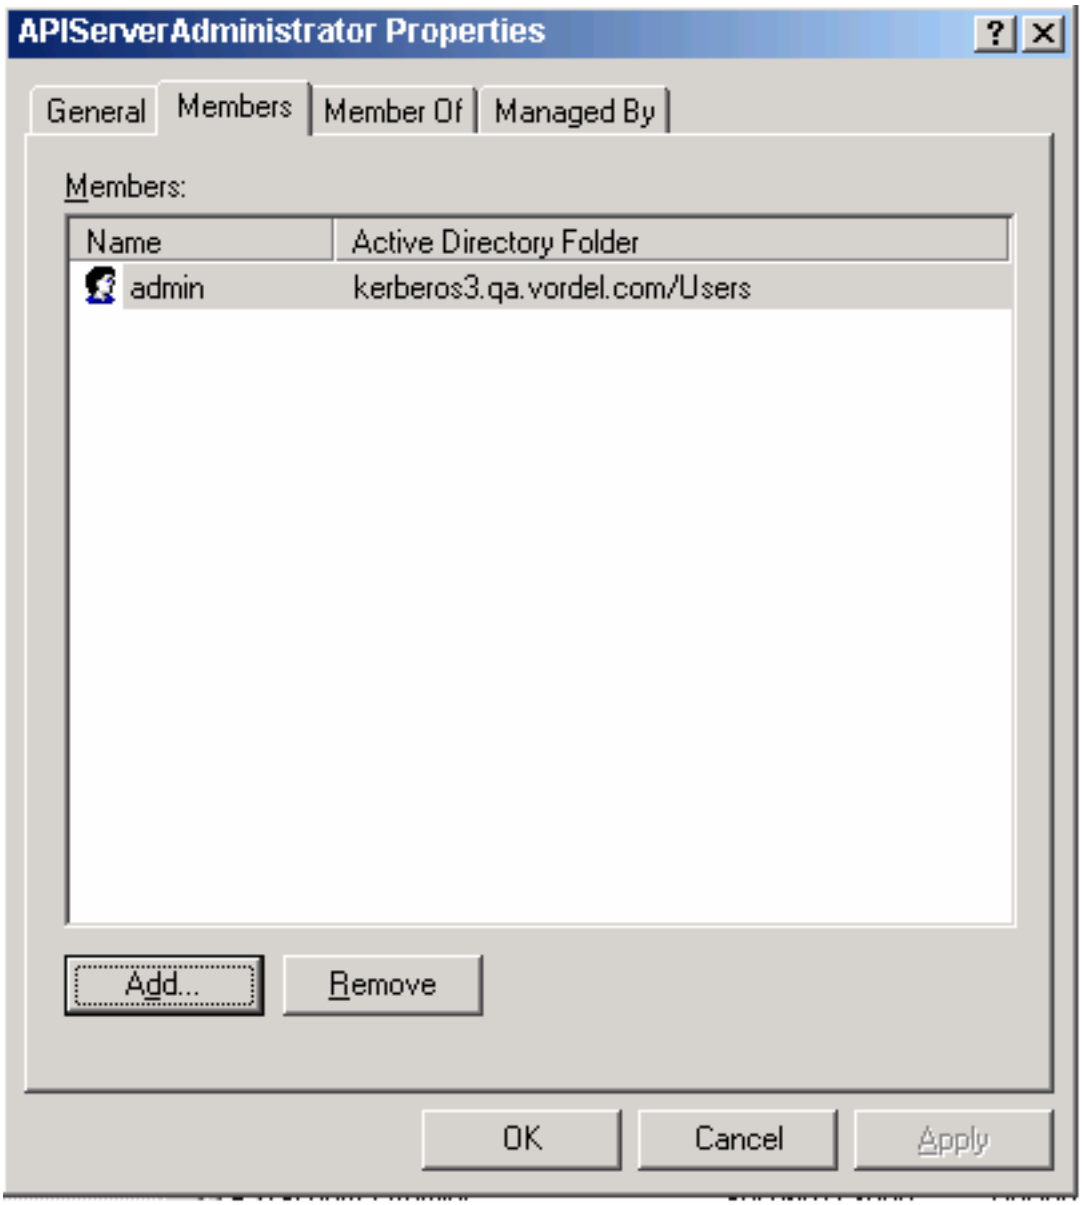

You can view the newly created user using an LDAP Browser.

Note

The `memberOf` attribute points to the Active Directory group. The user has an instance of this attribute for each group they are a member of.

Step 3: create an LDAP connection

To create an new LDAP Connection, perform the following steps:

1. In the Policy Studio, select **Open File** , and select the Admin Node Manager configuration file (for example, `INSTALL_DIR\apigateway\conf\fed\configs.xml`.
2. In the Policy Studio tree, select **External Connections** > **LDAP Connections**.
3. Right-click, and select **Create an LDAP Connection**.
4. Complete the fields in the dialog as appropriate. The specified **User Name** should be an LDAP administrator that has access to search the full directory for users. For example:

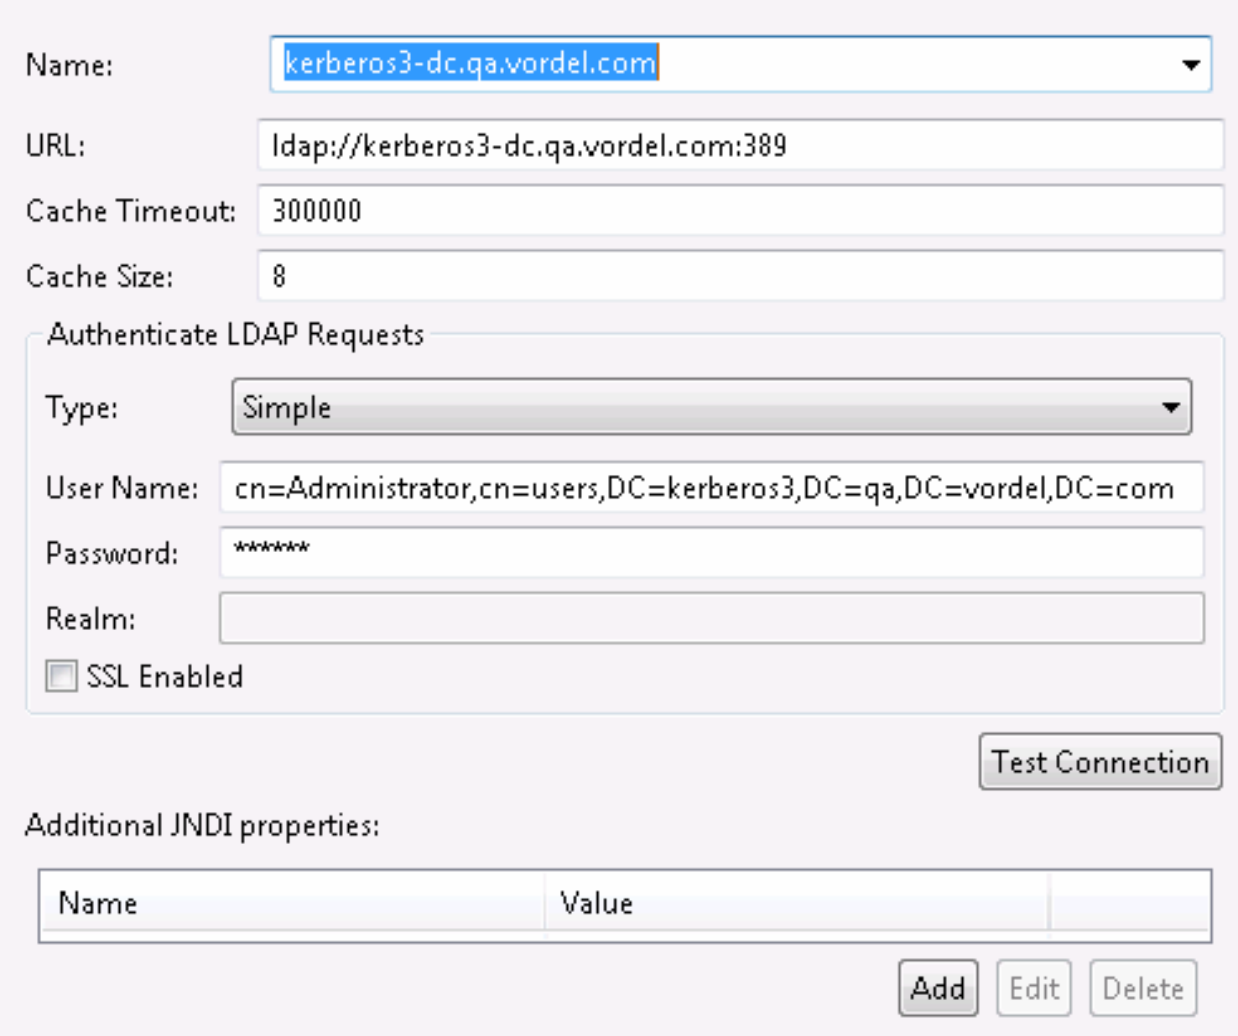

5. Click **Test Connection** to ensure the connection details are correct.

Step 4: create an LDAP repository

To create an new LDAP Repository, perform the following steps:

1. In the Policy Studio tree, select **External Connections** > **Authentication Repository Profiles** > **LDAP Repositories**.
2. Right-click, and select **Add a new Repository**.
3. Complete the following fields in the dialog:

Repository Name	Enter an appropriate name for the repository.
LDAP Directory	Use the LDAP directory created in the section called “Step 3: create an LDAP connection”.
Base Criteria	Enter the LDAP node that contains the users.
User Search Attribute	Enter `cn`. This is the username entered at login time (in this case, `admin`).
Authorization Attribute	Enter `distinguishedName`. This is the username entered at login time (`admin`). The `authentication.subject.id` message attribute is set to the value of this LDAP attribute (see example below). The `authentication.subject.id` is used as the base criteria in the filter that loads the LDAP groups (the user’s roles). This enables you to narrow the search to a particular user node in the LDAP tree. For more details, see the **Retrieve Attributes from Directory Server** filter in the section called “Step 5: create a test policy for LDAP authentication and RBAC”.

An example value of the `authentication.subject.id` message attribute is as follows:

```
CN=admin, CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com
```

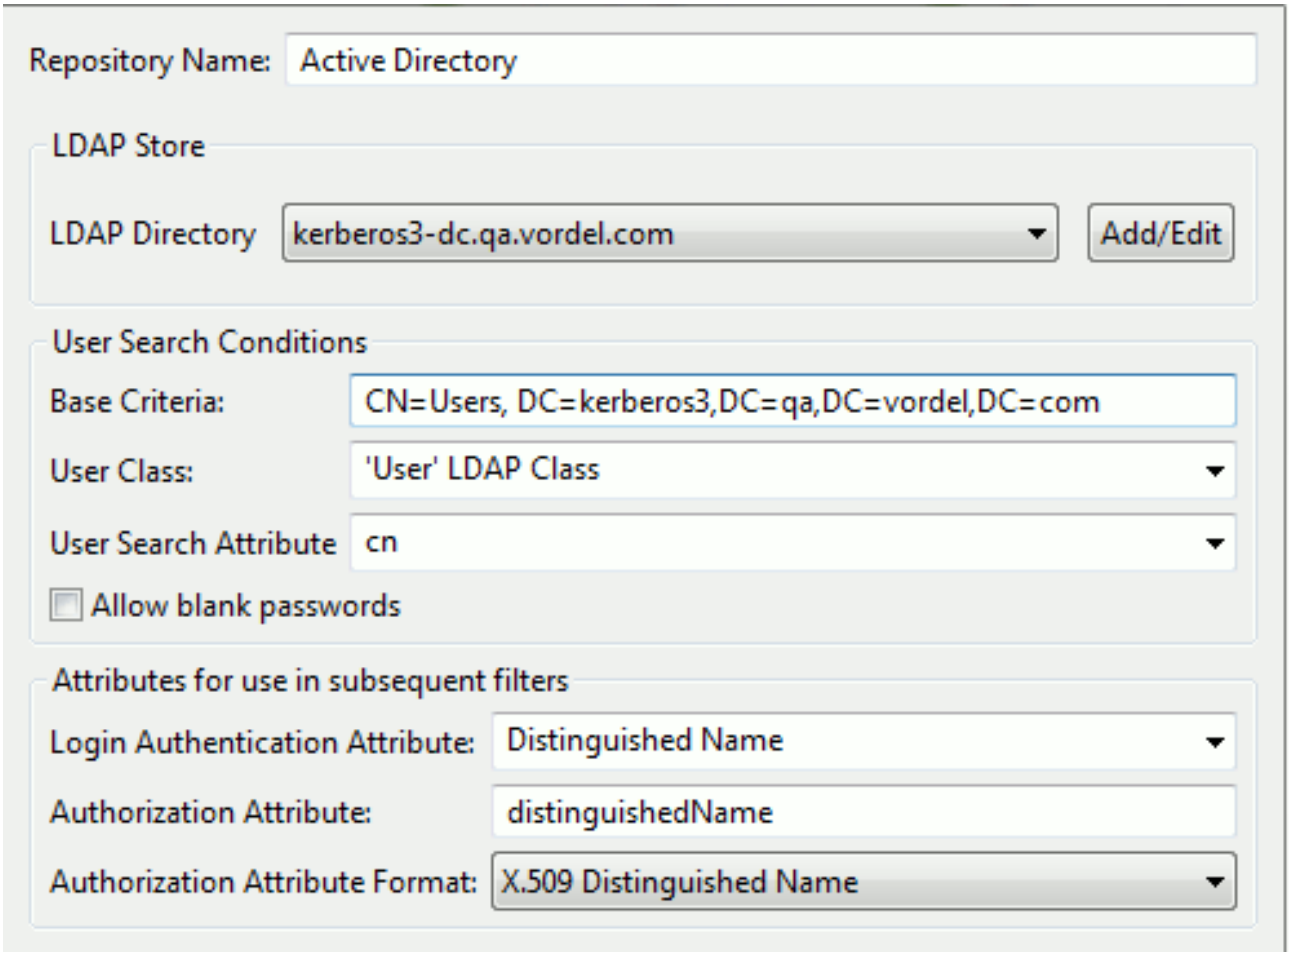

Connect to other LDAP repositories
This topic uses Microsoft Active Directory as an example LDAP repository. Other LDAP repositories such as Oracle Directory Server (formerly iPlanet and Sun Directory Server) and OpenLDAP are also supported. For an example of querying an Oracle Directory Server repository, see the **Retrieve Attributes from Directory Server** filter in the section called "Step 5: create a test policy for LDAP authentication and RBAC". For details on using OpenLDAP, see *OpenLDAP for authentication and RBAC of management services*.

Step 5: create a test policy for LDAP authentication and RBAC

To avoid locking yourself out of the Policy Studio, you can create a test policy for LDAP authentication and RBAC, which is invoked when a test URI is called on the server (and not a management services URI). For an example policy, select **Policies** > **Management Services** > **Sample LDAP Policies** > **Protect Management Interfaces (LDAP)** when the Admin Node Manager configuration is loaded.

Create the test policy
Perform the following steps:

1. Select **Open File** and load the Admin Node Manager configuration file in the Policy Studio. For example:

   ```
   INSTALL_DIR/apigateway/conf/fed/configs.xml
   ```

2. Right-click the **Policies** node in the tree view of the Policy Studio, and select **Add Policy**.

3. Enter a suitable name (for example `Test Policy`) for the new policy in the **Name** field, and click **OK**.
4. Click the new policy in the tree to start configuring the policy filters. You can configure the policy by dragging the required filters from the filter palette on the right, and dropping them on to the policy canvas.

For more details on creating policies, see the *API Gateway Policy Developer Guide*.

Configure the test policy
Configure the test policy with the following filters:

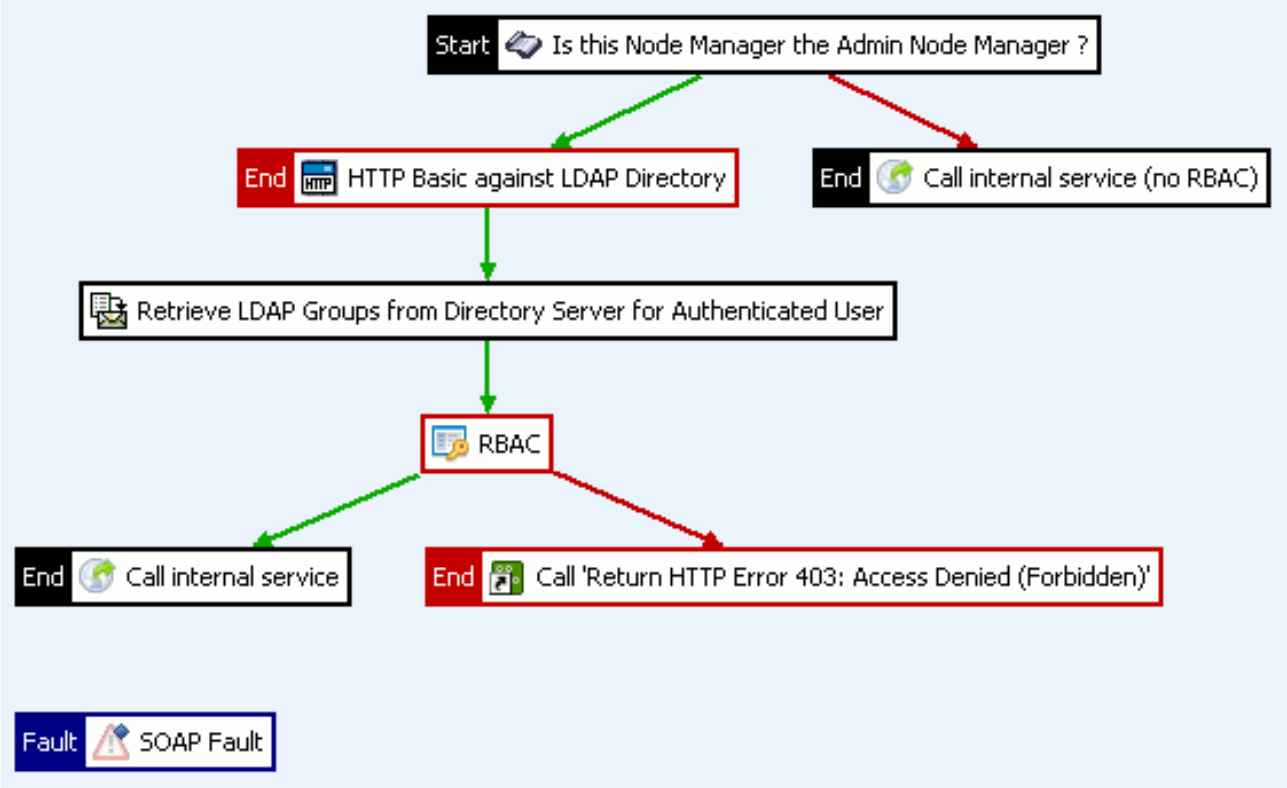

Scripting Language filter
This includes the following settings:

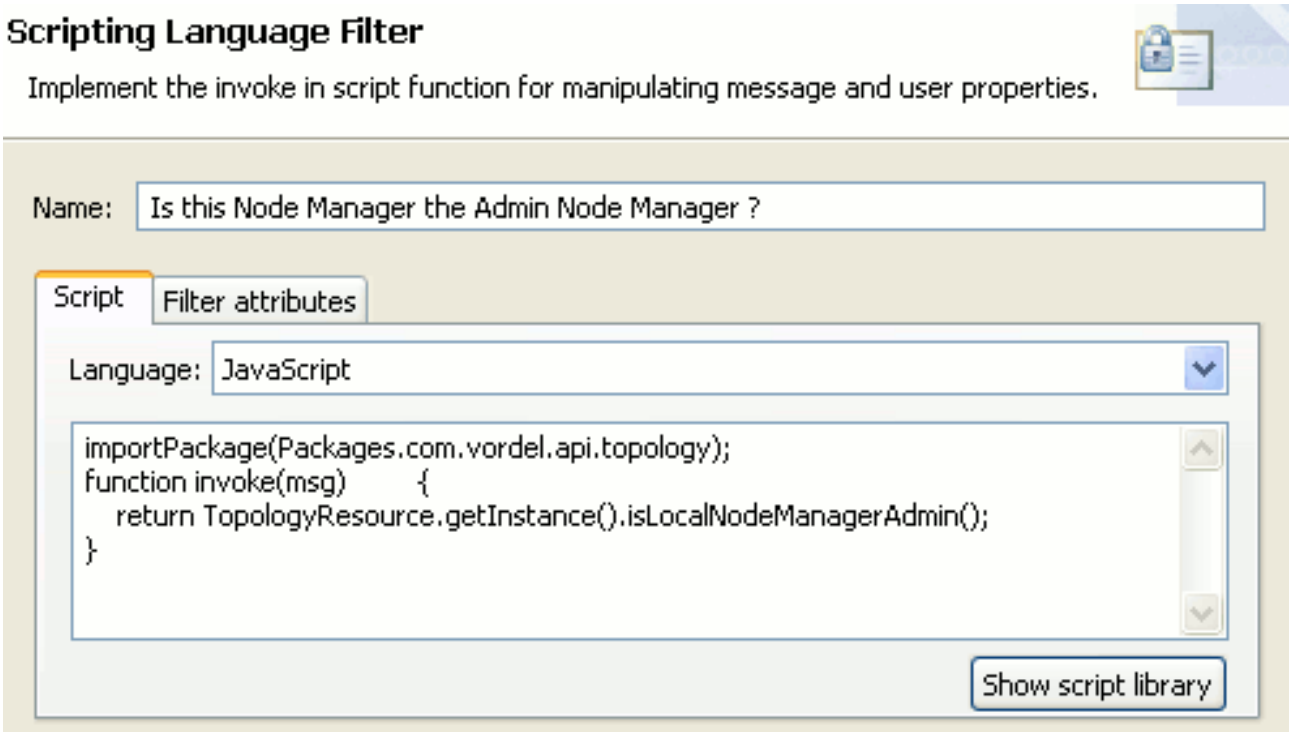

The **Scripting Language** filter performs the following tasks:

1. Returns true if the Node Manager is the Admin Node Manager and passes control to the **HTTP Basic Authentication** filter.
2. Otherwise, calls the **Call Internal Service (no RBAC)** filter without adding the `authentication.subject.role` and `authentication.subject.id` HTTP headers.

Call Internal Service (no RBAC) filter
This filter is called without adding any HTTP headers as follows:

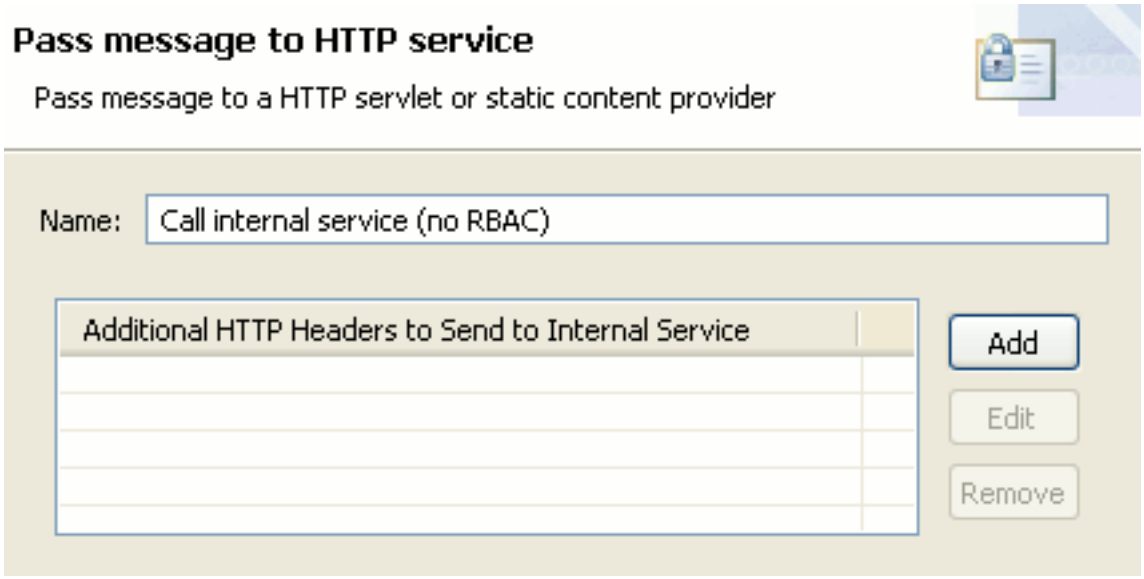

HTTP Basic Authentication filter
This filter uses the LDAP repository configured in the section called “Step 4: create an LDAP repository”, and includes the following settings:

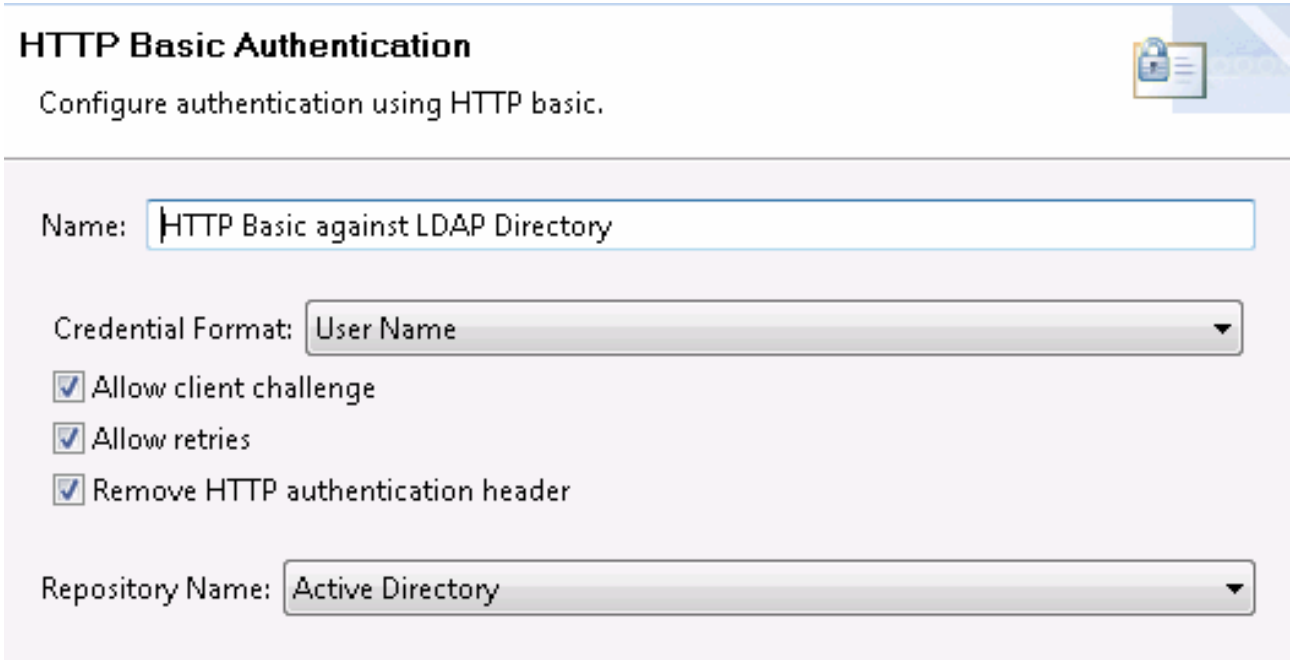

The **HTTP Basic Authentication** filter performs the following tasks:

1. Connects to the LDAP directory using the connection details specified in the LDAP directory.
2. Finds the user using the specified base criteria and search filter.
3. If the user is found, verifies the user's name and password against the LDAP repository by performing a bind.
4. If authentication fails, always throws a 401. This allows retry for browser users.

5. The `distinguishedName` (DName) is held in the `authentication.subject.id` message attribute. This is specified by the **Authorization Attribute** field in the LDAP repository configuration.
6. The user's roles (LDAP groups) are not available yet.

Retrieve Attributes from Directory Server filter
On the **Database** tab, this filter uses the LDAP directory configured in the section called "Step 3: create an LDAP connection", and includes the following settings:

- **Base Criteria**:
  `${authentication.subject.id}`
- **Search Filter**:
  `(objectclass=User)`
- **Attribute Name**:
  `memberOf`

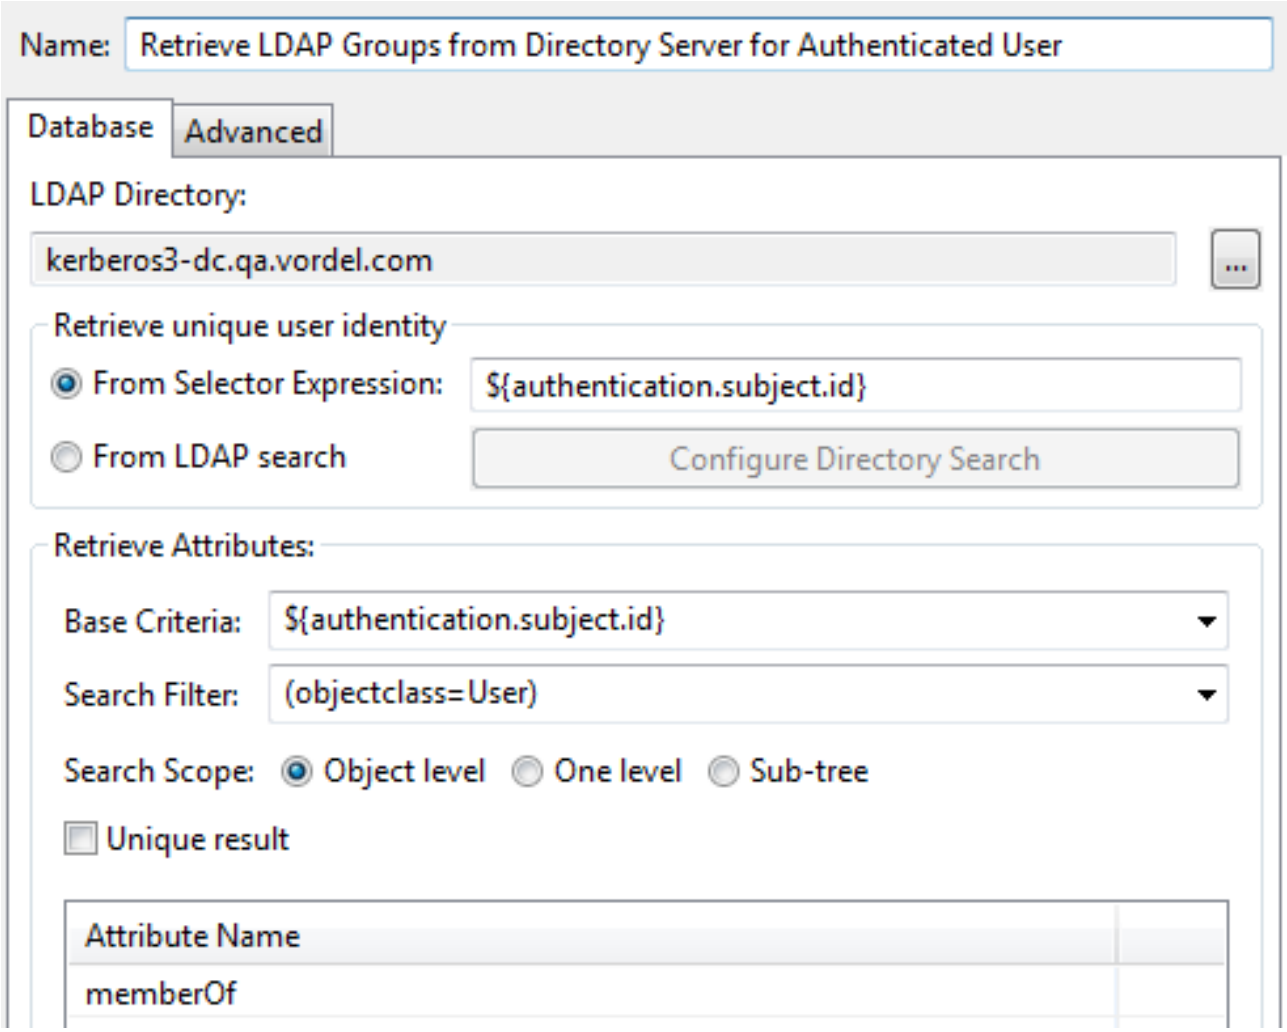

Note

On the **Advanced** tab, you must ensure that the **Enable legacy attribute naming for retrieved attributes** setting is selected.

The **Retrieve Attributes from Directory Server** filter performs the following tasks:

1. Using the user's DName as the search start point, find the user's `memberOf` attribute, and load the LDAP groups for the user.
2. If the user is in one group, the group name is contained in the `user.memberOf` message attribute.
3. If the user is in multiple (n) LDAP groups, the group names are held in `user.memberOf.1` ... `user.memberOf.n` message attributes.

Alternatively, the following screen shows an example of querying an Oracle Directory Server repository. The following **Search Filter** setting returns the authenticated user's groups instead of the user object:

```
(&(objectclass=groupOfUniqueNames)(uniqueMember=${authentication.subject.id}))
```

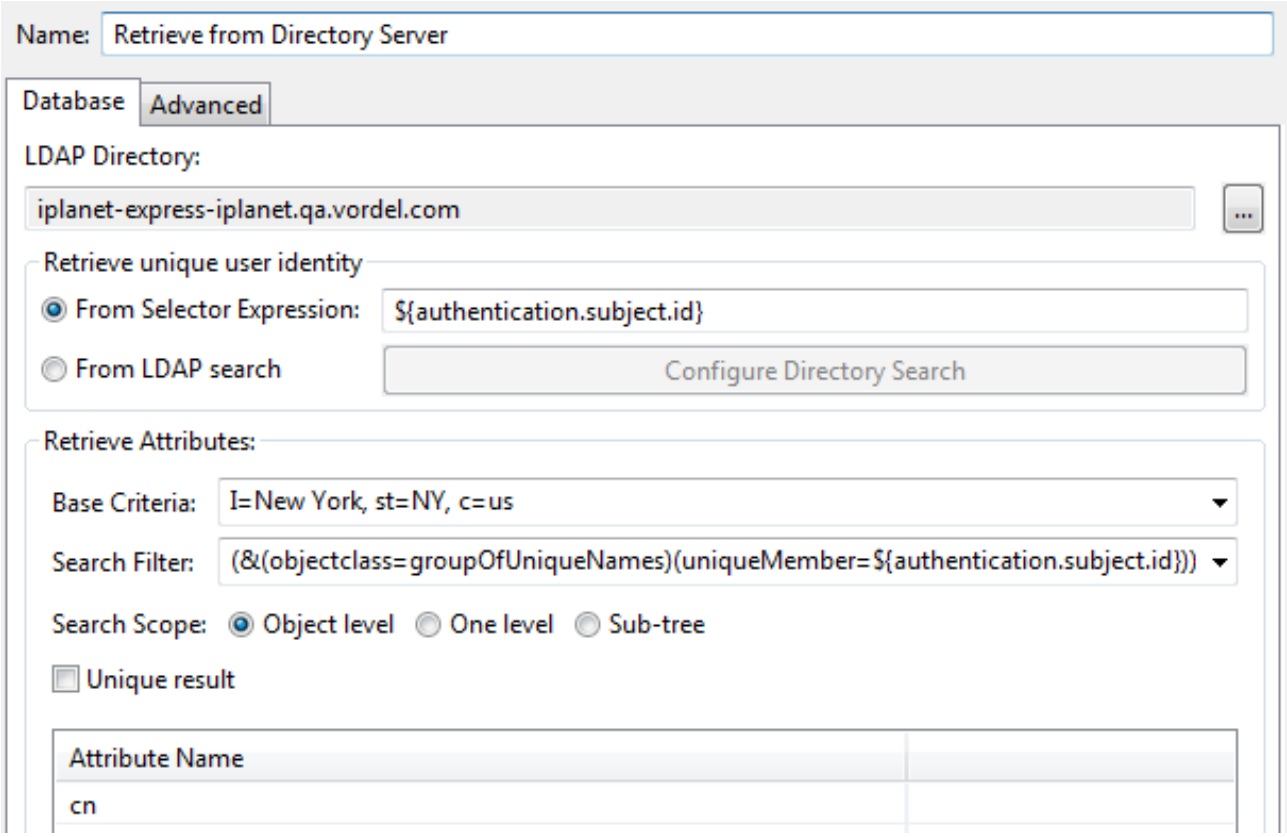

You should be able to query any LDAP directory in this way. Assuming that the user's groups or roles can be retrieved as attributes of an object, the query does not need to be for the user object.

Management Services RBAC filter

This filter includes a the following setting:

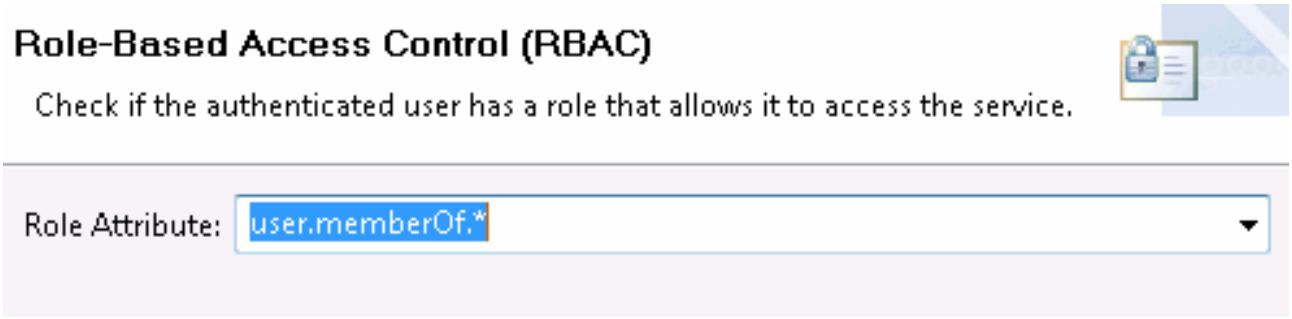

The **Management Services RBAC** filter performs the following tasks:

1. Reads the roles from the `user.memberOf.*` message attribute. It understands the meaning of the wildcard, and loads the roles as required. It creates a string version of the roles, and places it in the `authentication.subject.role` message attribute for consumption by the **Call Internal Service** filter, which receives the roles as an HTTP header value.
2. Determines which Management Service URI is currently being invoked.
3. Returns true if one of the roles has access to the Management Service currently being invoked, as defined in the `acl.json file`.
4. Otherwise, returns false and the **Return HTTP Error 403: Access Denied (Forbidden)** policy is called. This message content of this filter is shown when a valid user has logged into the browser, but their role(s) does not give them access to the URI they have invoked. For example, this occurs if a new user is created and they have not yet been assigned any roles.

Test the policy configuration
To test this policy configuration, perform the following steps:

1. Update the `acl.json` file with the new LDAP group as follows:

```
"CN=APIGatewayAdministrator,CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com" : [
  "emc", "mgmt", "mgmt_modify", "dashboard", "dashboard_modify", "deploy" "config",
  "monitoring", "events", "traffic_monitor", "settings", settings_modify", "logs" ]
```

2. Update the `adminUsers.json` file with the new role as follows:

```
{
  "name" : "CN=APIGatewayAdministrator,CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com",
  "id" : "role-8"
}
```

   And increase the number of roles, for example:

```
"uniqueIdCounters" : {
    "Role" : 9,
    "User" : 2
},
```

3. In the Policy Studio tree, select **Listeners** > **Node Manager** > **Add HTTP Services**, and enter a service name (for example, `LDAP Test`).
4. Right-click the HTTP service, and select **Add Interface** > **HTTP**.
5. Enter an available port to test the created policy (for example, 8888), and click **OK**.
6. Right-click the HTTP service, and select **Add Relative Path**.
7. Enter a relative path (for example, `/test`).

8. Set the **Path Specify Policy** to the **Protect Management Interfaces (LDAP)** policy, and click **OK**.
9. Close the connection to the Admin Node Manager file, and restart the Admin Node Manager so it loads the updated configuration.
10. Use API Gateway Explorer to call `http://localhost:8080/test`.
11. Enter the HTTP Basic credentials (for example, username `admin` and password `Oracle123`). If authentication is passed, the Admin Node Manager should return an HTTP 404 code (not found).

Important

Do not use the **Admin Users** tab in the API Gateway Manager to manage user roles because these are managed in LDAP.

Step 6: use the LDAP policy to protect management services

If the authentication and RBAC filters pass, you can now use this policy to protect the management interfaces. To ensure that you do not lock yourself out of the server, perform the following steps:

1. Make a copy of the `conf/fed` directory contents from the server installation, and put it into a directory accessible from the Policy Studio.
2. Make another backup copy of the `conf/fed` directory, which will remain unmodified.
3. In the Policy Studio, select **File** > **Open**, and browse to `configs.xml` in the first copy of the `fed` directory.
4. Under the **Listeners** > **Management Services** node, select the `/` and the `/configuration/deployments` relative paths, and set the **Path Specify Policy** to the **Protect Management Interfaces (LDAP)** policy.
5. Remove the previously created `LDAP Test` HTTP Services.
6. Close the connection to the file.
7. Copy the `fed` directory back to the Admin Node Manager's `conf` directory.
8. Reboot the Admin Node Manager.
9. Start the Policy Studio, and connect to the Admin Node Manager using `admin` and password `Oracle123` (the LDAP user credentials). You should now be able to edit API Gateway configurations as usual.

Add an LDAP user with limited access to management services

You can add an LDAP user with limited access to management services. For example, assume there is already a user named `Fred` defined in Active Directory. `Fred` has the following DName:

```
CN=Fred,CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com
```

`Fred` belongs to an existing LDAP group called `TraceAnalyzers`. He can also belong to other LDAP groups that have no meaning for RBAC in the API Gateway. The `TraceAnalyzers` LDAP group has the following DName:

```
CN=TraceAnalyzers,CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com
```

The user `Fred` should be able to read server trace files in a browser. No other access to management services should be given to `Fred`.

Adding Limited Access Rights
You must perform the following steps to allow `Fred` to view the trace files:

1. Add the following entry in the `roles` section in the `acl.json` file:

```
"CN=TraceAnalyzers,CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com" :
        [ "emc", "mgmt", "logs" ]
```

2. Update the `adminUsers.json` file with the new role as follows:

```
{
  "name" : "CN=TraceAnalyzers,CN=Users,DC=kerberos3,DC=qa,DC=vordel,DC=com",
  "id" : "role-8"
}]
```

   And increase the number of roles, for example:

```
"uniqueIdCounters" : {
    "Role" : 9,
    "User" : 2
},
```

3. Restart the Admin Node Manager so that the `acl.json` and `adminUsers.json`file updates are picked up.
4. Enter the following URL in your browser:
   `http://localhost:8090/`
5. Enter user credentials for `Fred` when prompted in the browser.
6. The API Gateway Manager displays a **Logs** tab enabling access to the trace files that `Fred` can view.

Note

`Fred` is not allowed to access the server APIs used by the Policy Studio. If an attempt is made to connect to the server using the Policy Studio with his credentials, an `Access denied` error is displayed. No other configuration is required to give user `Fred` the above access to the management services. Other users in the same LDAP group can also view trace files without further configuration changes because the LDAP group is already defined in the `acl.json` file.

OpenLDAP for authentication and RBAC of management services

Overview

This topic explains how to use Local Directory Access Protocol (LDAP) to authenticate and perform Role-Based Access Control (RBAC) of management services. You can use the sample **Protect Management and Interfaces (LDAP)** policy instead of the **Protect Management Interfaces** policy. This means that an LDAP repository is used instead of the local Admin User store for authentication and RBAC of users attempting to access the management services. This topic describes how to reconfigure the server to use OpenLDAP as the LDAP repository, and to use the Apache Directory Studio as an LDAP browser.

Note

This topic applies to both API Gateway and API Gateway Analytics.

Prerequisites
This example assumes that you already have configured connection to the OpenLDAP server and setup your organization groups and users that you wish to use to perform RBAC. For example:

- **LDAP URL**: `ldap://openldap.qa.oracle.com:389`
- **User**: `cn=admin,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE`
- **Password**: `oracle`

Step 1: create an OpenLDAP group for RBAC roles

To create a new user group in OpenLDAP, perform the following steps:

1. Select the `cn=admin,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE` directory.
2. Right-click, and select **New** > **New Entry**.
3. Select **Create entry from scratch**.
4. Click **Next**.
5. Add an `organizationalUnit` object class.
6. Click **Next**.
7. Set the **Parent** to `o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE`.
8. Set the **RDN** to `ou = RBAC`.
9. Click **Next**.
10. Click **Finish**.

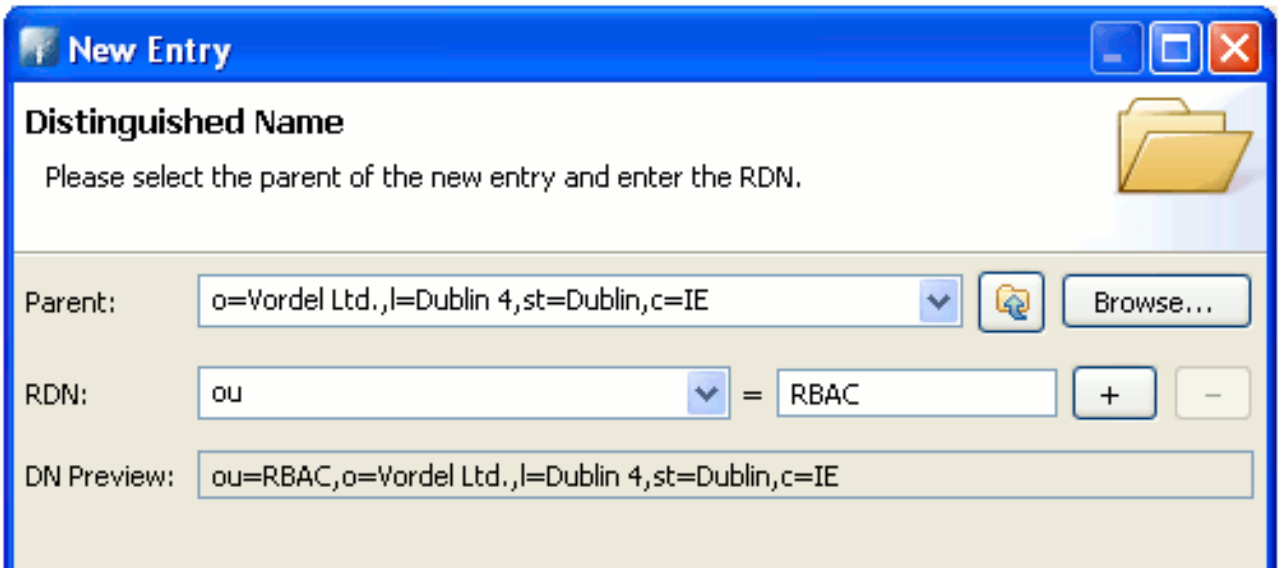

You can view the new group using an LDAP Browser. For example:

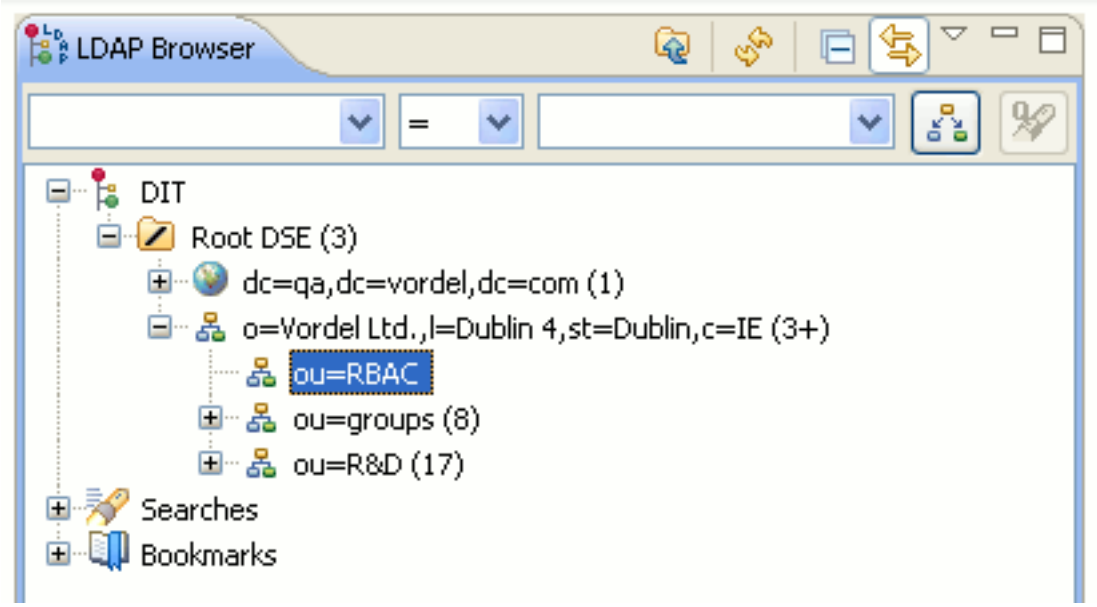

Step 2: add RBAC roles to the OpenLDAP RBAC group

You must add the following default RBAC roles to the `ou=RBAC,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE` group to give the LDAP users appropriate access to the API Gateway management services:

- `API Gateway Administrator`
- `API Gateway Operator`
- `KPS Administrator`
- `Policy Developer`
- `Deployer`

These RBAC roles are located in the `roles` section of the `acl.json` file.

Adding Roles to the RBAC Directory
To add these RBAC roles to the OpenLDAP RBAC group, perform the following steps:

1. Select the `cn=admin,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE` directory.
2. Right-click, and select **New** > **New Entry**.
3. Select **Create entry from scratch**.
4. Click **Next**.
5. Add a `groupOfNames` object class.
6. Click **Next**.
7. Set the **Parent** to `ou=RBAC,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE`.
8. Set the **RDN** to `cn = Policy Developer`.
9. Click **Next**.
10. In the **DN Editor** dialog, set `admin` as first member of the following group: `cn=admin,ou=R&D,o=Vordel Ltd.,l=Dublin 4,st=Dublin,c=IE`. You can change the member Distinguished Name at any time.
11. Click **OK**.
12. Click **Finish**.

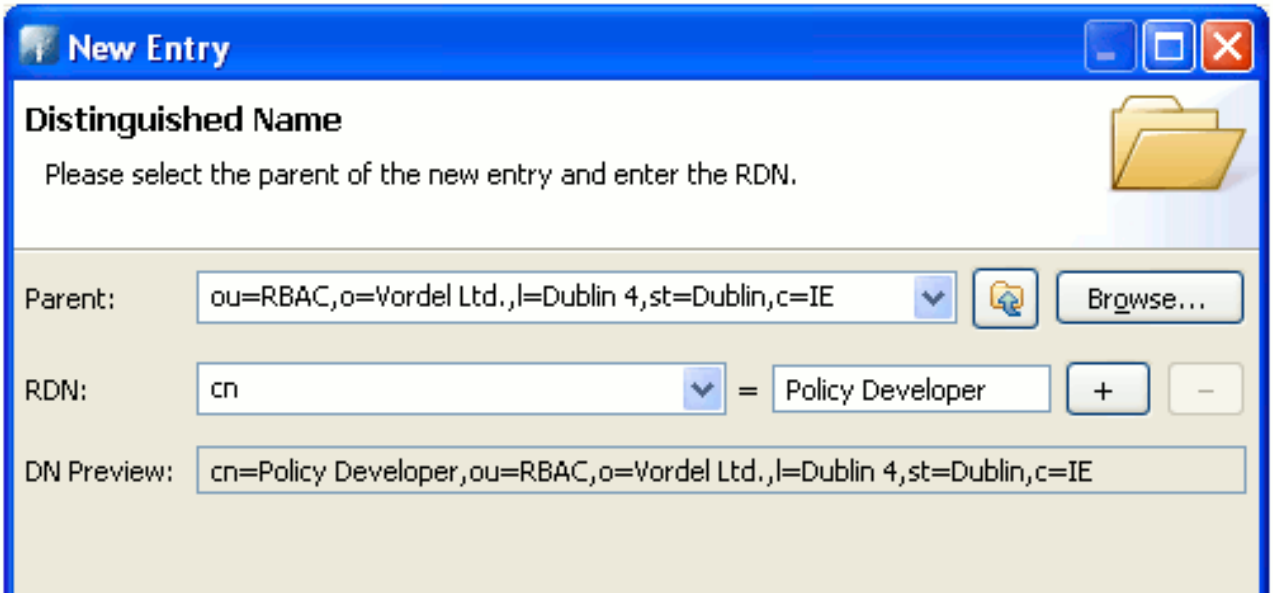

You can view the role in the OpenLDAP group in an LDAP Browser. For example:

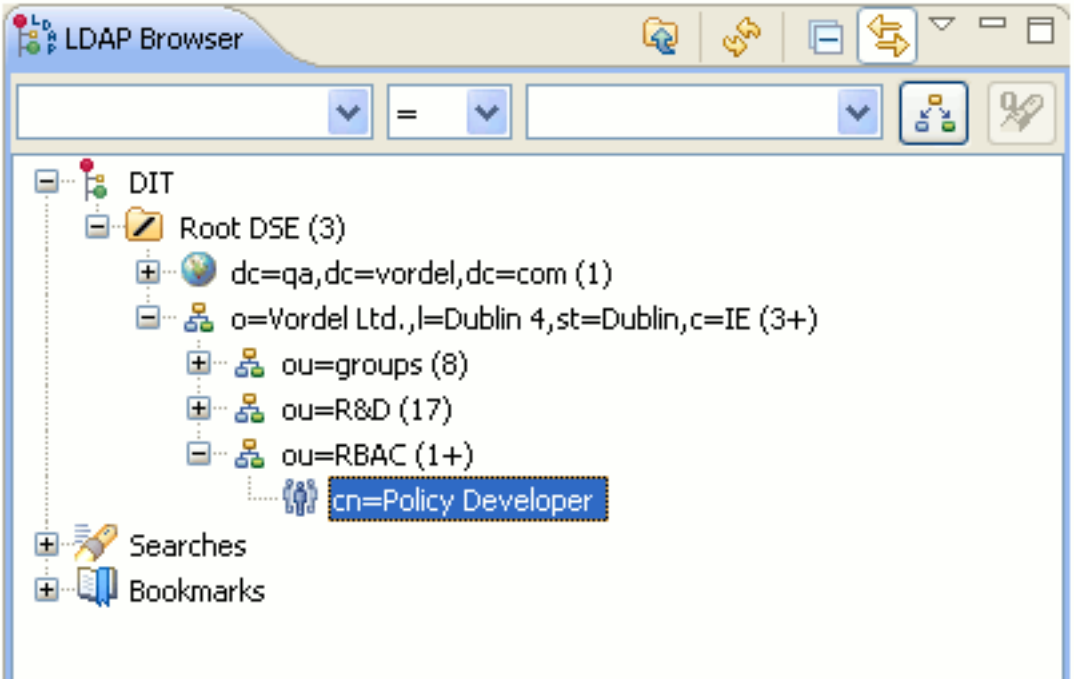

Add other roles to the RBAC directory

You can repeat these steps to add other roles to the RBAC directory. Alternatively, you can copy the `Policy Developer` entry, and paste it into the RBAC directory, renaming the entry with required RBAC role name. For example:

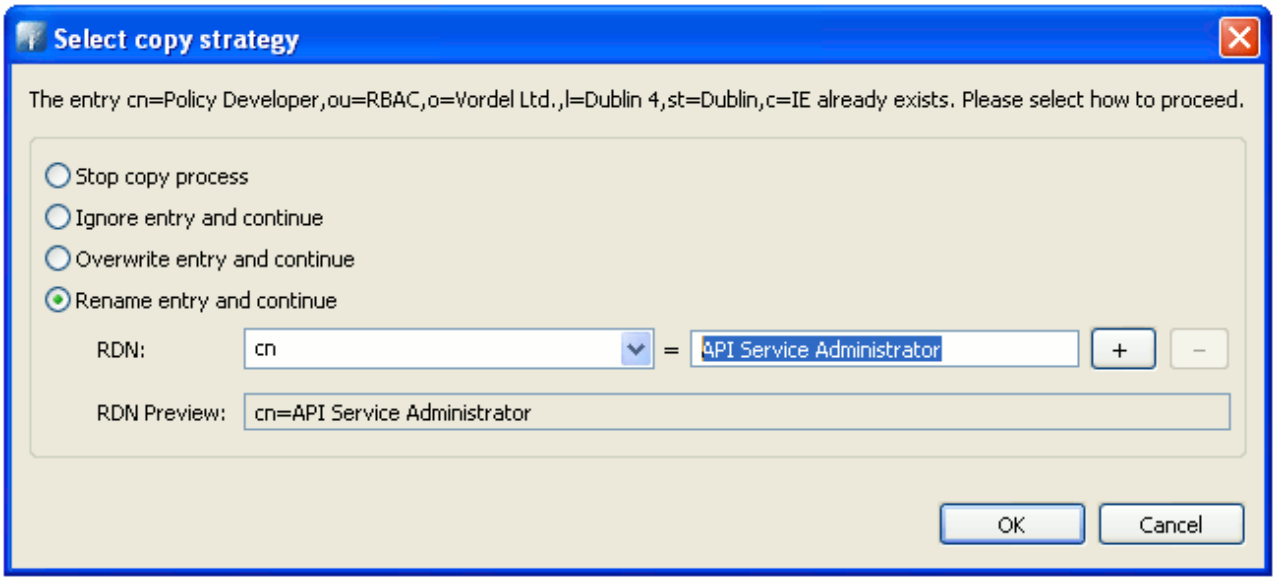

Note

You should have the RBAC directory ready to add members to the role entries. By default, the `admin` user (“`cn=admin,ou=R&D,o=Vordel Ltd.,l=Dublin 4,st=Dublin,c=IE`”) is already a member of the role entries.

The following example shows the added roles:

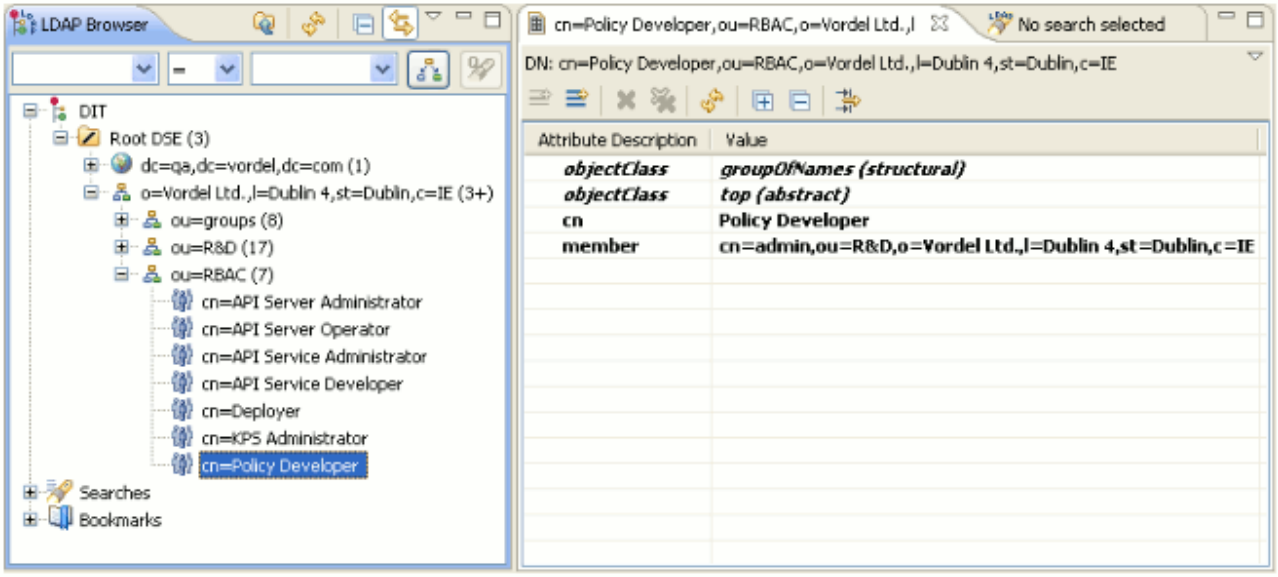

Now you can add new users to the RBAC role entries. The member attribute value should contain the user Distinguished Name. This is explained in the next section.

Step 3: add users to the OpenLDAP RBAC group

To add a user to the OpenLDAP RBAC group, perform the following steps:

1. Select the required RBAC group (for example, `cn=API Gateway Administrator`) to view the group details.
2. Right-click the list of group attributes, and select **New Attribute**.
3. Enter `member` in the attribute type.

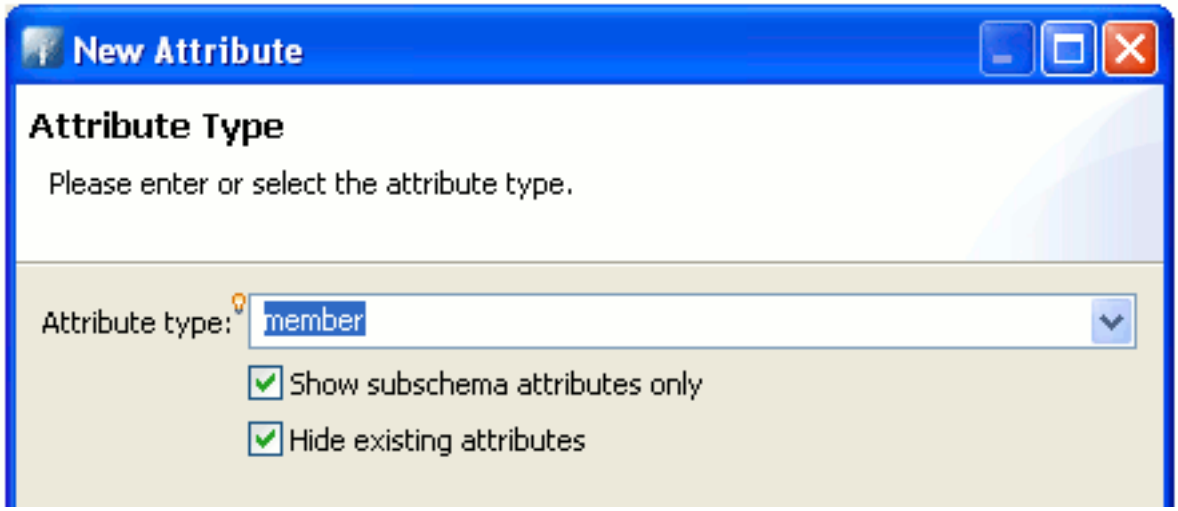

4. Click **Finish**.
5. In the **DN Editor** dialog, enter the user Distinguished Name (for example, `cn=joe.bloggs,o=Vordel Ltd.,l=Dublin 4,st=Dublin,c=IE`).
6. Click **OK**.

The `cn=joe.bloggs,o=Vordel Ltd.,l=Dublin 4,st=Dublin,c=IE` new user has been added to the RBAC `API Gateway Administrator` role.

Step 4: create an LDAP connection

To create an new LDAP Connection, perform the following steps:

1. In the Policy Studio, select **Open File** , and select the Admin Node Manager configuration file (for example, `INSTALL_DIR\apigateway\conf\fed\configs.xml`).
2. In the Policy Studio tree, select **External Connections** > **LDAP Connections**.
3. Right-click, and select **Create an LDAP Connection**.
4. Complete the fields in the dialog as appropriate. For example:

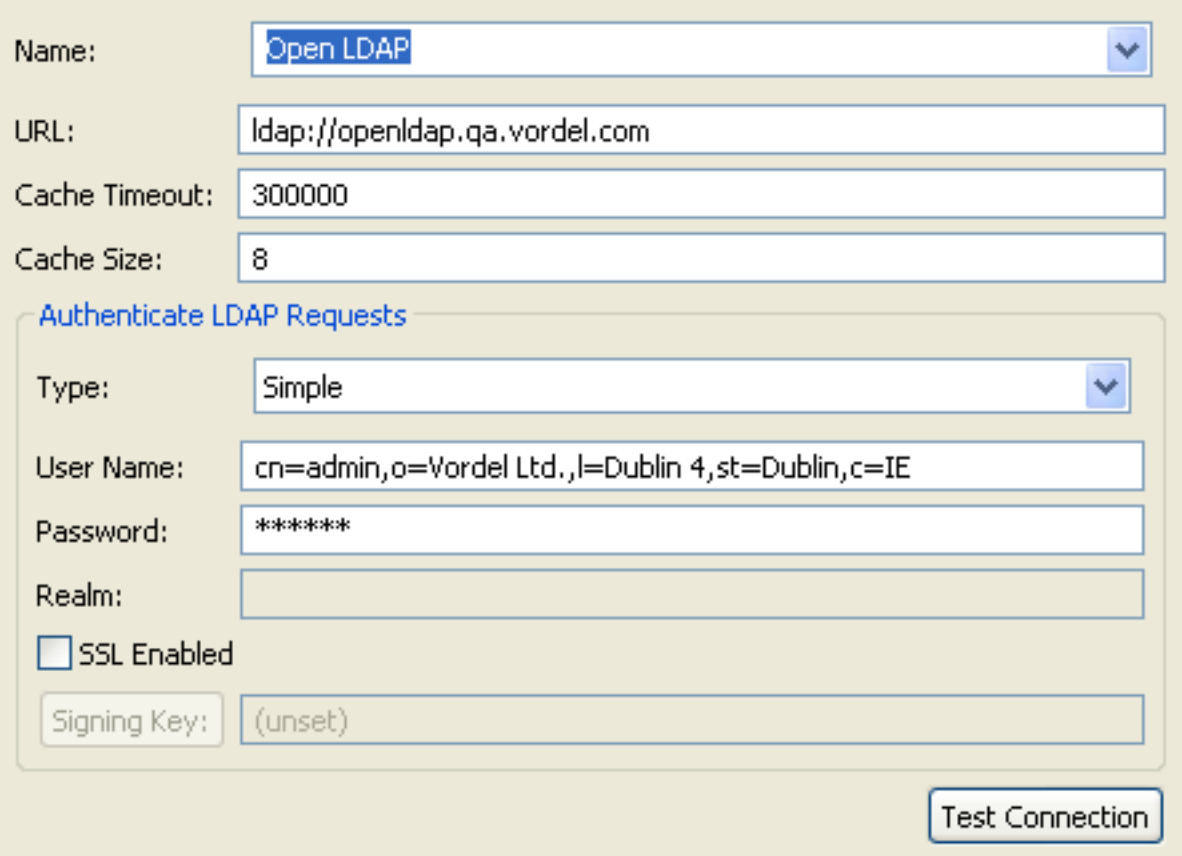

Note

The specified **User Name** should be an LDAP administrator that has access to search the full directory for users.

5. Click **Test Connection** to ensure the connection details are correct.

Step 5: create an OpenLDAP repository

To create an new OpenLDAP Repository, perform the following steps:

1. In the Policy Studio tree, select **External Connections** > **Authentication Repository Profiles** > **LDAP Repositories**.
2. Right-click, and select **Add a new Repository**.
3. Complete the following fields in the dialog:

Repository Name	Enter an appropriate name for the repository.
LDAP Directory	Use the LDAP directory created in the section called "Step 4: create an LDAP connection".
Base Criteria	Enter the LDAP node that contains the users (for example, see the LDAP Browser screen in the section called "Step 3: add users to the OpenLDAP RBAC group").
User Search Attribute	Enter `cn`. This is the username entered at login time (in this case, `admin`).
Authorization Attribute	Enter `cn`. The `authentication.subject.id` message attribute is set to the value of this LDAP attribute (for example, `cn=admin,ou=R&D,o=Vordel Ltd.,l=Dublin 4,st=Dublin,c=IE`. The `authentication.subject.id` is used as the base criteria in the filter used to load the LDAP groups (the user's roles). This allows you to narrow the search to a particular user node in the LDAP tree. For more details, see the **Retrieve Attributes from Directory Server** filter in the section called "Step 6: create a test policy for LDAP authentication and RBAC".

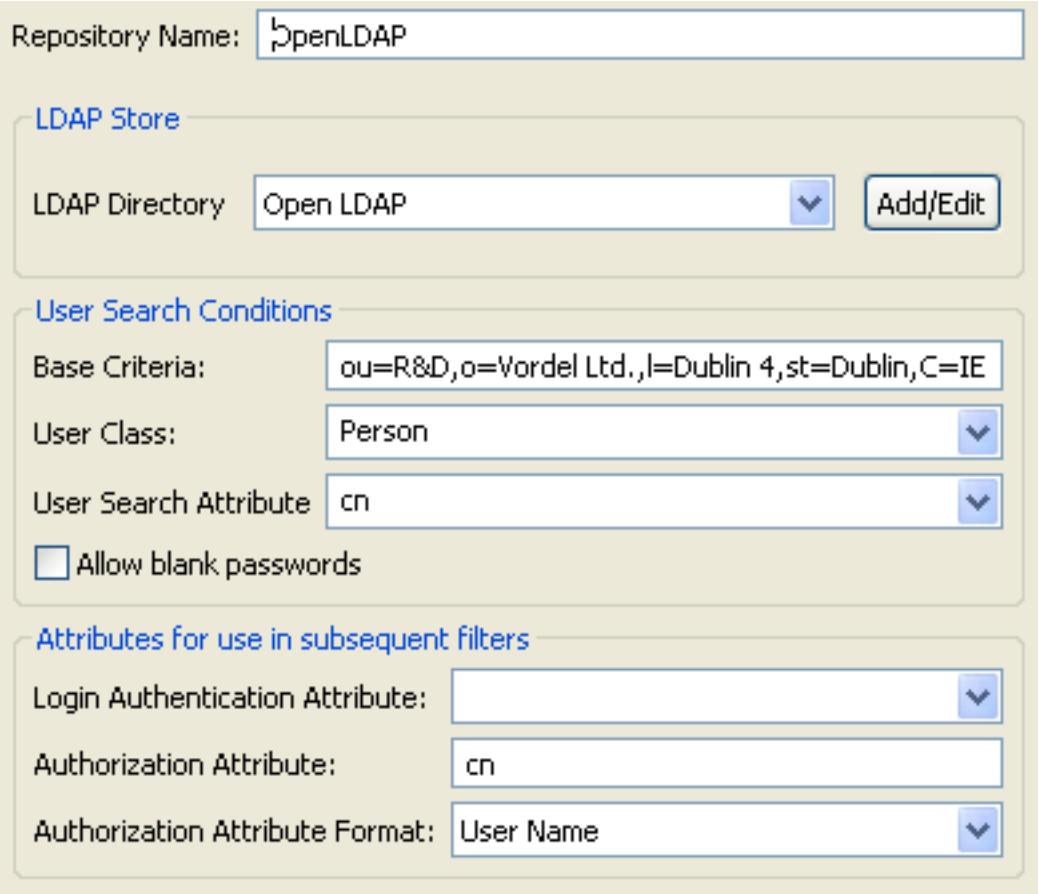

Connect to other LDAP repositories
This topic uses OpenLDAP as an example LDAP repository. Other LDAP repositories such as Oracle Directory Server (formerly iPlanet and Sun Directory Server) and Microsoft Active Directory are also supported. For details on using a Microsoft Active Directory repository, see *Active Directory for authentication and RBAC of management services*. For an example of querying an Oracle Directory Server repository, see the **Retrieve Attributes from Directory Server** filter in the section called "Step 5: create a test policy for LDAP authentication and RBAC".

Step 6: create a test policy for LDAP authentication and RBAC

To avoid locking yourself out of the Policy Studio, you can create a test policy for LDAP authentication and RBAC, which is invoked when a test URI is called on the server (and not a management services URI). For an example policy, select **Policies** > **Management Services** > **Sample LDAP Policies** > **Protect Management Interfaces (LDAP)** when the Admin Node Manager configuration is loaded.

Create the test policy
Perform the following steps:

1. Select **Open File** and load the Admin Node Manager configuration file in the Policy Studio. For example:

   ```
   INSTALL_DIR/apigateway/conf/fed/configs.xml
   ```

2. Right-click the **Policies** node in the tree view of the Policy Studio, and select **Add Policy** .

3. Enter a suitable name (for example `Test Policy`) for the new policy in the **Name** field, and click the **OK** button. The new policy is now visible in the tree view.
4. Click the new policy in the tree view to start configuring the filters for the policy. You can easily configure the policy by dragging the required filters from the filter palette on the right of the Policy Studio, and dropping them on to the policy canvas.

For more details on creating policies, see the *API Gateway Policy Developer Guide*.

Configure the test policy
Configure the test policy with the following filters:

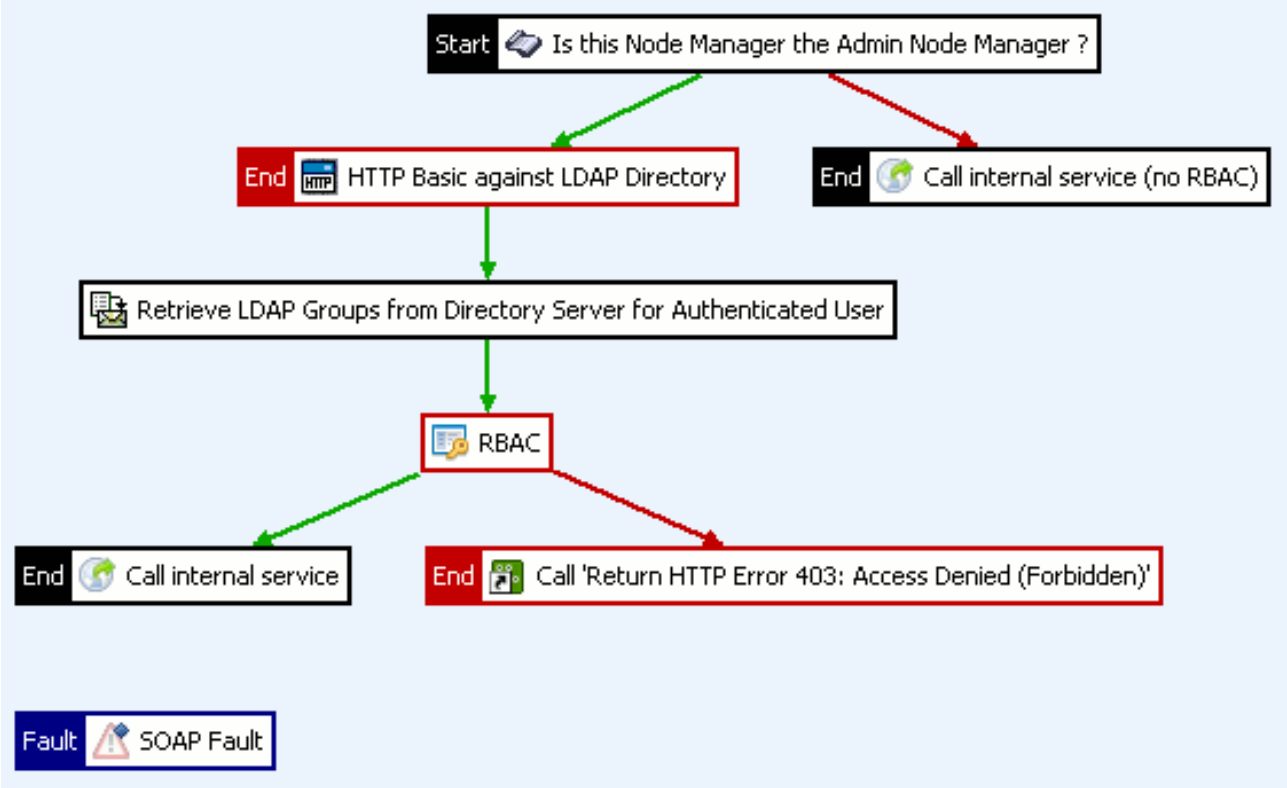

Scripting Language filter
This includes the following settings:

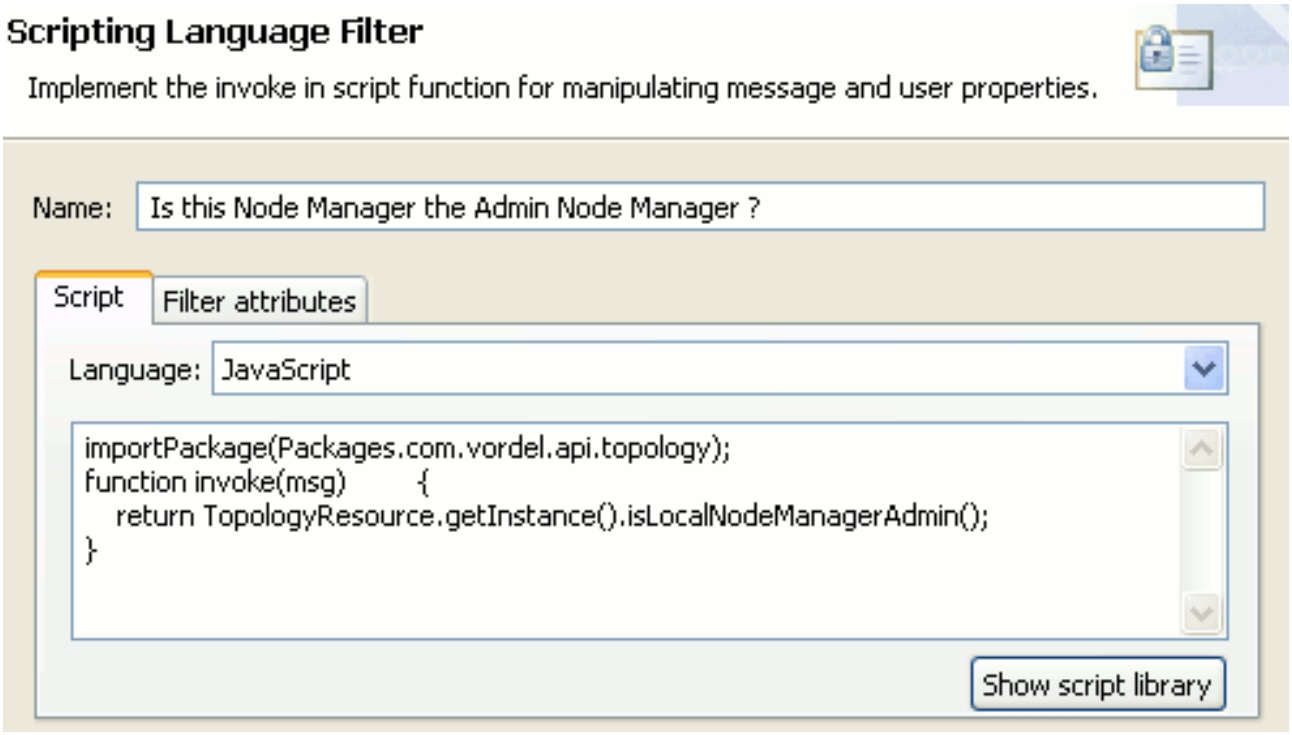

The **Scripting Language** filter performs the following tasks:

1. Returns true if the Node Manager is the Admin Node Manager, and passes control to the **HTTP Basic Authentication** filter.
2. Otherwise, calls the **Call Internal Service (no RBAC)** filter without adding the `authentication.subject.role` and `authentication.subject.id` HTTP headers.

Call Internal Service (no RBAC) filter
This filter is called without adding any HTTP headers as follows:

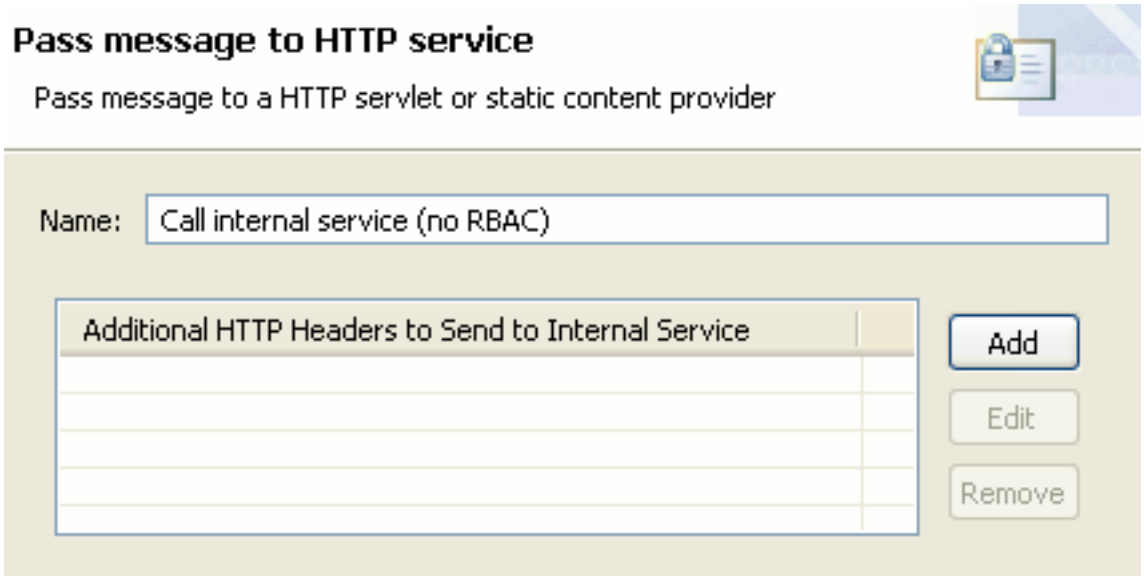

HTTP Basic Authentication filter
This filter uses the LDAP repository configured in the section called “Step 5: create an OpenLDAP repository”, and includes the following settings:

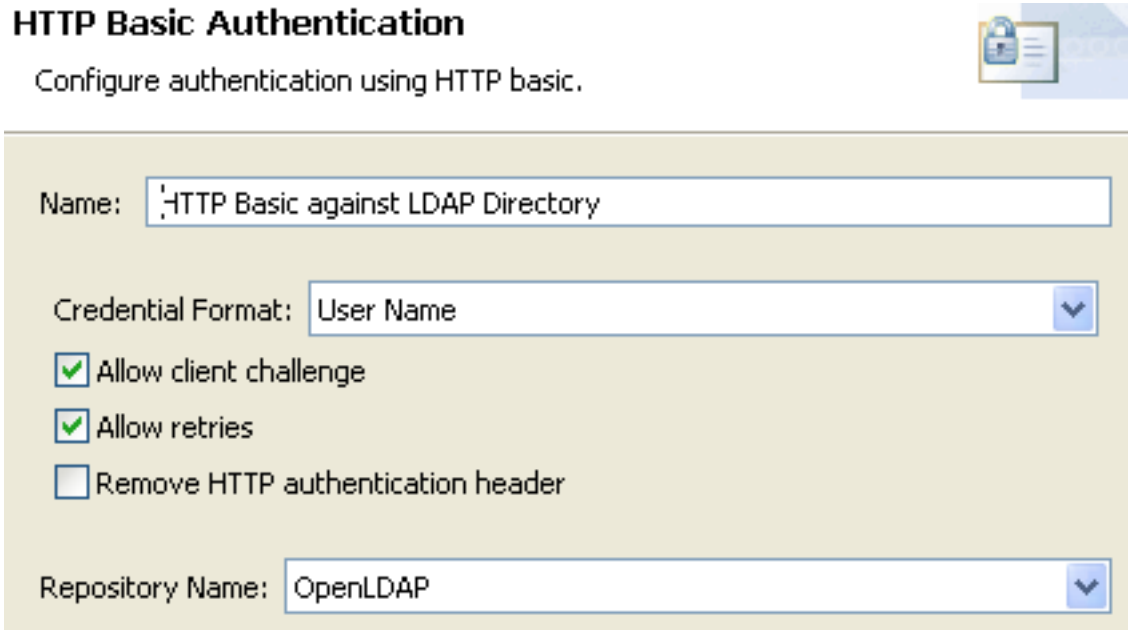

The **HTTP Basic Authentication** filter performs the following tasks:

1. Connects to the LDAP directory using the connection details specified in the LDAP directory.
2. Finds the user using the specified base criteria and search filter.
3. If the user is found, verifies the user's name and password against the LDAP repository by performing a bind.
4. If authentication fails, always throws a 401. This allows retry for browser users.
5. The `distinguishedName` (DName) is held in the `authentication.subject.id` message attribute. This is

specified by the **Authorization Attribute** field in the LDAP repository configuration.

6. The user's roles (LDAP groups) are not available yet.

Retrieve Attributes from Directory Server filter
On the **Database** tab, this filter uses the LDAP directory configured in the section called "Step 4: create an LDAP connection", and includes the following settings:

- **Retrieve Unique User Identity From Selector Expression**:
  `${authentication.subject.id}`
- **Base Criteria**:
  `ou=RBAC,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE`
- **Search Filter**:
  `(&(objectclass=groupOfNames)(member=CN=${authentication.subject.id},ou=R&D,o=Vordel Ltd.,l=Dublin 4,st=Dublin,C=IE))`
- **Attribute Name**:
  `cn`

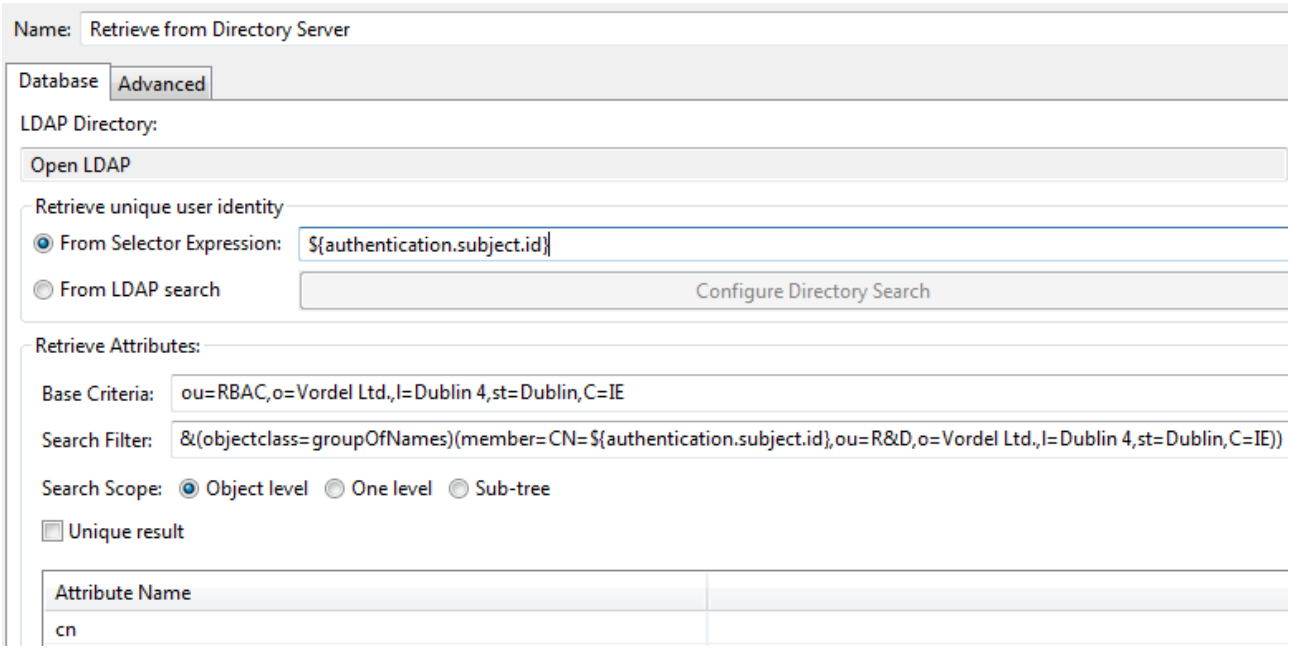

Note

On the **Advanced** tab, you must ensure that the **Enable legacy attribute naming for retrieved attributes** setting is selected.

The **Retrieve Attributes from Directory Server** filter performs the following tasks:

1. Using the user's DName as the search start point, finds the user's `cn` attribute, and loads the LDAP groups for the user.
2. If the user is in one group, the group name is contained in the `user.cn` message attribute.
3. If the user is in multiple (n) LDAP groups, the group names are contained in `user.cn.1` ... `user.cn.n` message attributes.

Management Services RBAC filter
This filter includes a the following setting:

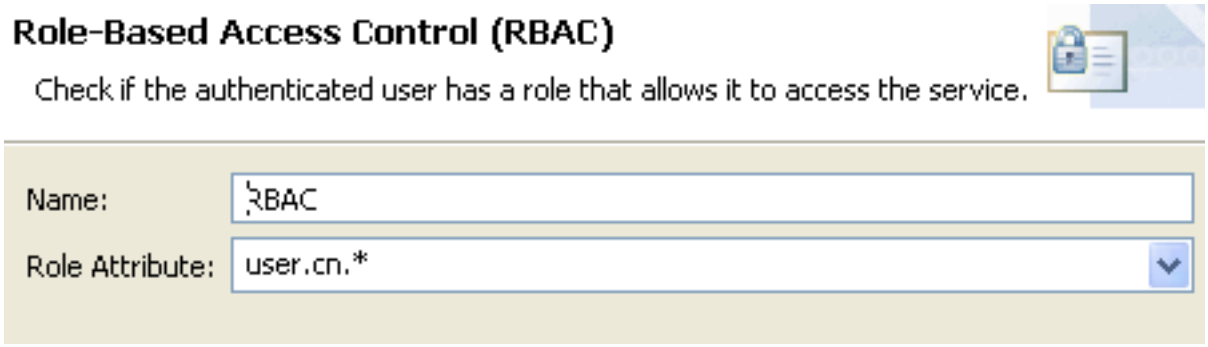

The **Management Services RBAC** filter performs the following tasks:

1. Reads the roles from the `user.cn.*` message attribute. It understands the meaning of the wildcard, and loads the roles as required. It creates a string version of the roles, and places it in the `authentication.subject.role` message attribute for consumption by the **Call Internal Service** filter, which receives the roles as an HTTP header value.
2. Determines which Management Service URI is currently being invoked.
3. Returns true if one of the roles has access to the Management Service currently being invoked, as defined in the `acl.json` file.
4. Otherwise, returns false and calls the **Return HTTP Error 403: Access Denied (Forbidden)** policy. The message content of this filter is shown when a valid user has logged into the browser, but their role(s) does not give them access to the URI they have invoked. For example, this occurs if a new user is created and they have not yet been assigned any roles.

Call Internal Service filter
This filter includes a the following settings:

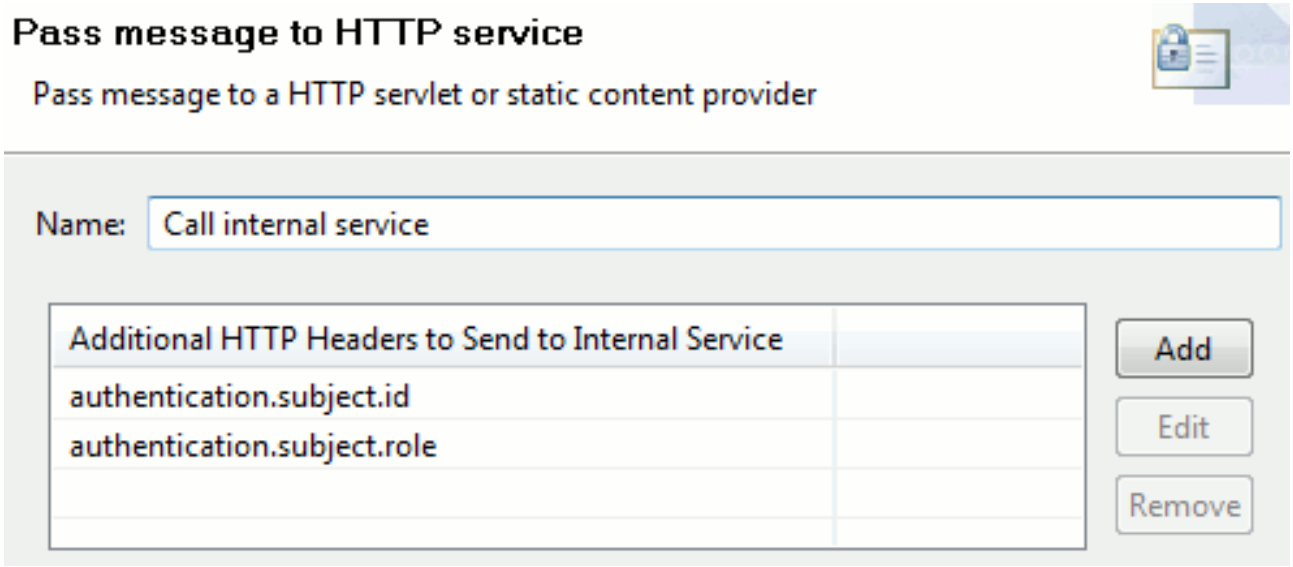

The **Call Internal Service** filter sends the message to the internal service with the following additional HTTP headers:

- `authentication.subject.role`
- `authentication.subject.role`

Note

The `authentication.subject.id` message attribute is specified using `${authentication.subject.orig.id}` because `authentication.subject.id` holds the full DName, and the `cn` only needs to be passed to the services.

Test the policy configuration
To test this policy configuration, perform the following steps:

1. In the Policy Studio tree, select **Listeners** > **Node Manager** > **Add HTTP Services**, and enter a service name (for example, `LDAP Test`).
2. Right-click the HTTP service, and select **Add Interface** > **HTTP**.
3. Enter an available port to test the created policy (for example, `8888`), and click **OK**.
4. Right-click the HTTP service, and select **Add Relative Path**.
5. Enter a relative path (for example, `/test`).
6. Set the **Path Specify Policy** to the **Protect Management Interfaces (LDAP)** policy, and click **OK**.
7. Close the connection to the Admin Node Manager file and reboot the Admin Node Manager so it loads the updated configuration.
8. Use API Gateway Explorer to call `http://localhost:8080/test`.
9. Enter the HTTP Basic credentials (for example, username `admin` and password `Oracle123`). If authentication is passed, the Admin Node Manager should return an HTTP 404 code (not found).

Step 7: use the OpenLDAP policy to protect management services

If the authentication and RBAC filters pass, you can now use this policy to protect the management interfaces. To ensure that you do not lock yourself out of the server, perform the following steps:

1. Make a copy of the `conf/fed` directory contents from the server installation, and put it into a directory accessible from the Policy Studio.
2. Make another backup copy of the `conf/fed` directory, which will remain unmodified.
3. In the Policy Studio, select **File** > **Open**, and browse to `configs.xml` in the first copy of the `fed` directory.
4. Under the **Listeners** > **Management Services** node, select the `/` and the `/configuration/deployments` relative paths, and set the **Path Specify Policy** to the **Protect Management Interfaces (LDAP)** policy.
5. Remove the previously created `LDAP Test` HTTP Services.
6. Close the connection to the file.
7. Copy the `fed` directory back to the Admin Node Manager's `conf` directory.
8. Reboot the Admin Node Manager.
9. Start the Policy Studio, and connect to the Admin Node Manager using `admin` with its LDAP password (for example, `Oracle123`). You should now be able to edit API Gateway configurations as usual.

Configure a DNS service with wildcards for virtual hosting

Overview

The Domain Name System (DNS) is a hierarchical distributed naming system for resources on the Internet or a private network. It translates domain names to numerical IP addresses used to locate services and devices worldwide, and associates details with domain names assigned to each resource. You can use wildcard DNS records to specify multiple domain names by using an asterisk in the domain name (such as `*.example.com`).

A virtual host is a server, or pool of servers, that can host multiple domain names (for example, `company1.api.example.com` and `company2.api.example.com`). To use virtual hosting for the API Gateway, you need to be able to use different host names to refer to the same destination server. A convenient way to do this is by using a DNS service to specify wildcard entries for physical host names (for example, by setting `*.example.com` to `192.0.2.11`).

This setting enables you to run more than one website, or set of REST APIs, on a single host machine (in this case, `192.0.2.11`). Each domain name can have its own host name, paths, APIs, and so on. For example:

```
https://company1.api.example.com:8080/api/v1/test
https://company2.api.example.com:8080/api/v2/test
```

Note

This topic explains how to set up DNS wildcards for virtual hosts. This is a prerequisite for configuring the API Gateway for virtual hosts. For more details, see the *API Gateway Policy Developer Guide*.

DNS workflow

When a client makes an HTTP request to a specific URL, such as `http://www.example.com`, the client must decide which IP address to connect to. In this case, the client makes a request on the local name service for the address of the `www.example.com` machine. This usually involves the following workflow:

1. Check the `/etc/hosts` file for the specified name, and if that fails, make a DNS request.
2. Send the DNS request to the default DNS server for the client system, which refers the client to the server for `example.com` (referral), or makes the request on the client's behalf (recursion).
3. The DNS server that manages the `www.example.com` domain responds with a record that specifies the IP address of `www.example.com`.
4. The client contacts that IP address, and includes a `Host` header when making its request (`Host: www.example.com`).
5. The server can use this `Host` header to distinguish requests to different sites (virtual hosts) that use the same physical hardware and HTTP server process.

You can divide a parent domain such as `example.com` into subdomains, one for each hosted site. This provides each site with its own distinct namespace (for example, `company1.example.com`, `company2.example.com`, `company3.example.com`, and so on).

BIND DNS software

Internet Systems Consortium (ISC) BIND is the de facto standard DNS server used on the Internet. It works on a wide variety of LINUX and UNIX systems, and on Microsoft Windows. Windows Server operating systems can also use the Windows DNS service. The examples in this topic describe the configuration for BIND only. For more details, see http://www.isc.org/downloads/bind/.

UNIX systems generally enable the configuration of BIND to be installed using a package manager. For example, for Ubuntu systems, you can use the following command:

```
$ sudo apt-get install bind9
```

The service and packages are generally called `bind`, `bind9`, or `named`. For example, on Ubuntu, you can restart BIND using the following command:

```
$ sudo /etc/init.d/bind9 start
$ sudo /etc/init.d/bind9 stop
$ sudo /etc/init.d/bind9 restart
```

Configure a wildcard domain

You can configure BIND using the `named.conf` configuration file, which is typically installed in one of the following locations:

```
/etc/named.conf
/etc/bind/named.conf
```

This configuration file is typically set up with `include` entries to make configuration and upgrade easier, and which should be easy to follow. The simple example described in this topic uses a flat file only. There are two main parts to the configuration:

- `options` to control the behavior of the service
- `zone` indicates how each autonomous part of the domain name tree should behave

The example shown in this topic uses the simplest options.

Configure DNS options

The following are some example DNS options that you can configure for your installation in the `named.conf` file:

```
options {
    directory "."; // e.g., /var/named
    listen-on port 8866 { any; }; // remove this for production
    pid-file "named.pid";
    allow-recursion { any; };
    allow-transfer { any; };
    allow-update { any; };
    forwarders { 10.253.253.253; 10.252.252.252; };
};
```

The example options are described as follows:

Option	Description
`directory`	This is normally set to a directory on a writeable filesystem, such as `/var/bind`. This example is set to the current directory (`"."`).
`listen-on port`	This example is set to `8866` for testing, but you should leave this blank for real DNS use in a production environment. Most DNS clients such as Web browsers always request on the standard DNS port 53. You can also configure debugging tools such as `dig` and `nslookup` to try other ports.
`pid-file`	This example sets up `named` to run in the current directory (or `/var/named` if configured), and gives it a name to store its lock file.
`allow-`	The example `allow-` options are permissive, and allow any external host to

Option	Description
	use this name server, update it dynamically, and transfer a dump of its domains. This makes it unsuitable for external exposure. To restrict these `allow-` options to the local system, change the `any` settings to `127.0.0.1`.
`forwarders`	Requests that cannot be serviced locally are forwarded to the specified servers, which are the default DNS servers for the site. Specifying `forwarders` enables you to use this name server as your local DNS. It can forward requests for sites outside your domain (for example, `example.com`) to the forwarding name servers. This enables your test environment to coexist with your normal name servers.

Configure default zones

The next step is to configure default zones for the `127.0.0.1` address. For example:

```
zone "localhost" IN {
     type master;
     file "localhost.zone";
     allow-transfer { any; };
};

zone "0.0.127.in-addr.arpa" IN {
     type master;
     file "127.0.0.zone";
     allow-transfer { any; };
};
```

These settings configure addresses for mapping `localhost` to `127.0.0.1`, and mapping `127.0.0.1` back to `localhost` when required. These example options are described as follows:

Option	Description
`type`	Specifies that this is the `master` server for `localhost`.
`file`	Specifies the domain zone file that contains the records for the domain (for more details, see the section called “Configure domain zone files”).
`allow-transfer`	Specifies addresses that are allowed to transfer zone information from the zone server. The default `any` setting means that the contents of the domain can be transferred to anyone.

Configure logging

The following settings provide some default logging configuration:

```
logging {
   channel default_syslog {
     file "named.log";
     severity debug 10;
   };
};
```

For example, you can change the `file` setting to `/var/log/named.log`, or change the severity level.

Configure a wildcard domain

You must now configure your wildcard domain. Here you specify the name of the domain for which to serve wildcard addresses. For example:

```
zone "example.com." IN {
   type master;
   file "example.zone";
   allow-transfer { any;
   };
};
```

This is almost identical to the `localhost` zone already configured.

Configure domain zone files

Finally, your `zone` configuration now includes references to two separate domain zone files (in this case, `localhost.zone` and your wildcard zone (`example.zone`). These domain zone files use a standard format, which is defined in RFC 1035: http://www.ietf.org/rfc/rfc1035.txt. For more details, see also Wikipedia: http://en.wikipedia.org/wiki/Zone_file.

The domain zone files dictate how the server responds to requests for data in the specified zone. For example, your basic `localhost.zone` domain is as follows:

```
@  IN SOA  root.acmecorp.com. (
            20131021  ; serial number of zone file (yyyymmdd##)
            3H        ; refresh time
            15M       ; retry time in case of problem
            1W        ; expiry time
            1D )      ; maximum caching time in case of failed lookups
   IN NS    @
   IN A     127.0.0.1
```

In this file, `@` is shorthand for the domain, and describes the first (and only) record in the file. `@` specifies the name of the zone, and has the following associated records:

- `SOA`—specifies details about the zone, including various serial and timer settings
- `NS`—specifies that the name server for this domain is `localhost`
- `A`—specifies that the address for localhost is `127.0.0.1`

Your wildcard domain is similar. For example, the contents of `example.zone` are as follows:

```
@  IN SOA  . root.acmecorp.com. (
             20130903  ; serial number of zone file (yyyymmdd##)
             604800    ; refresh time
             86400     ; retry time in case of problem
             2419200   ; expiry time
             604800)   ; maximum caching time in case of failed lookups
   IN NS     .
   IN A      192.168.0.10
*  IN A      192.168.0.10
```

This domain has the `SOA`, `NS`, and `A` records like the `localhost` zone, but also adds a `*` record. This matches any subdomain of `example.com` to resolve to the specified IP address.

Manage embedded ActiveMQ messaging

Overview

The **Messaging** view in API Gateway Manager enables you to manage the Apache ActiveMQ messaging broker that is embedded in the API Gateway instance. For example, this includes managing JMS message queues, topics, subscribers, and consumers, monitoring server connections, and so on. For more details on Apache ActiveMQ, see http://activemq.apache.org/. For details on the embedded ActiveMQ architecture, see the *API Gateway Concepts Guide*.

Important

The **Messaging** view is disabled by default. Before you can use the features in this view, you must first enable the ActiveMQ messaging broker and configure a shared directory in Policy Studio. For more details, see *Embedded ActiveMQ settings*.

Manage messaging queues

The **Queues** tab displays the messaging queues that exist in the selected API Gateway instance, along with queue statistics. For example, these include the number of messages in the queue, the number of consumers, and so on:

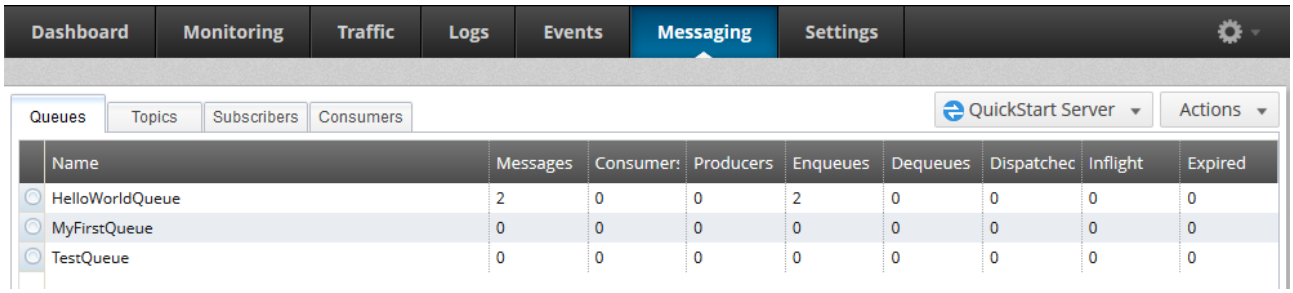

You can use the **Actions** drop-down menu on the right of the screen to perform the following tasks:

- Create a new queue
- Delete an existing queue
- Purge a queue (remove all its messages)
- Refresh the list of queues

For example, to create a new queue, select **Actions** > **Create Queue**, and enter a name for the queue in the dialog. The newly created queue is displayed in the list of queues.

Manage messages in a queue

When you select a queue name on the **Queues** tab, this displays the messages contained in the queue. For example:

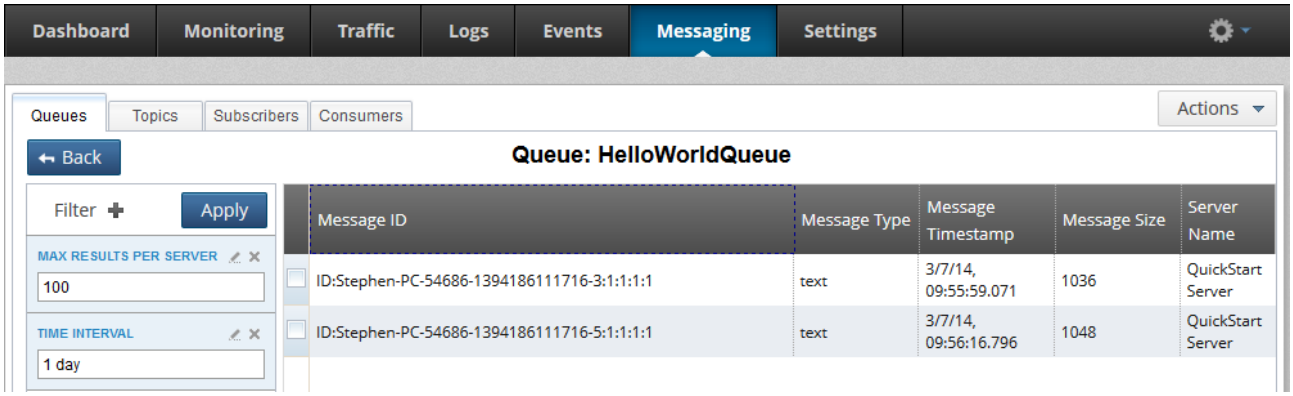

You can filter the messages displayed in this screen using the filter panel on the left. By default, you can configure the number of messages displayed, or the time interval for messages. Click the **Filter** button to add more viewing options (for example, **JMS Message ID** or **JMS Type**).

You can use the **Actions** drop-down menu on the right of the screen to perform the following tasks:

- Create a message in the queue
- Copy selected messages to a different queue
- Move selected messages to a different queue
- Delete selected messages from the queue

Create a new message

To create a new message in the selected queue, perform the following tasks:

1. Select **Actions** > **Create Message**, and complete the following **Message Details**:

Setting	Description
Destination API Gateway	Specifies the required destination API Gateway for the message (for example `APIGateway1`).
Message Type	Specifies the required type of the message (`Text` or `Binary`). Defaults to `Text`.
JMS Reply To	Specifies the optional destination to send a reply message to (for example, JMS queue or topic). It is the responsibility of the application that consumes the message from the queue (JMS consumer) to send the message back to the specified destination.
Time to live	By default, the message never expires. However, if a message becomes obsolete after a certain period, you can set an optional expiration time (in milliseconds). If the specified value is `0`, the message never expires.
JMS Delivery Mode	Select one of the following delivery modes: • **Persistent**: Instructs the JMS provider to ensure that a message is not lost in transit if the JMS provider fails. A message sent with this delivery mode is logged to persistent storage when it is sent. This is the default mode. • **Non-persistent**: Does not require the JMS provider to store the message. With this mode the message may be lost if the JMS provider fails.
JMS Correlation ID	Specifies an optional identifier used to correlate response messages with the

Setting	Description
	corresponding request messages. For example, you could enter the `${id}` message attribute selector to correlate request messages with their correct response messages.
JMS Priority	Specifies an optional message priority level to deliver urgent messages first. Priority levels are from `0` (lowest) to `9` (highest). If you do not specify a priority level, the default level is 4.
JMS Type	Specifies an optional user-defined message header that is defined as an arbitrary string. This is used to identify the message structure or payload (for example, the XML schema that the payload conforms to).

2. Click next to specify **User-defined properties**. Custom properties may be required for compatibility with other messaging systems. You can use message attribute selectors as property values. For example, you can create a property called `AuthNUser`, and set its value to `${authenticated.subject.id}`. Other applications can then filter on this property—such as only consume messages where `AuthNUser` equals `admin`. To add a new property, enter the property name and property value in the text box. You can click the green plus sign to add more properties.
3. Click next to specify the message payload content. If you selected a **Message Type** of `Text`, you can enter a **Text Payload** (for example, `Hello World!`). Alternatively, if you selected a **Message Type** of `Binary`, click **Select File**, and browse to the file that you wish to send.
4. Click **Send**.

View message contents

When you have created a message in a queue, the message is displayed in the list of messages on the **Queues** tab. You can click a message in the list to display a detailed view of the message properties and message body contents. To save the message body to a file, click the **Download** link on the bottom right of the screen.

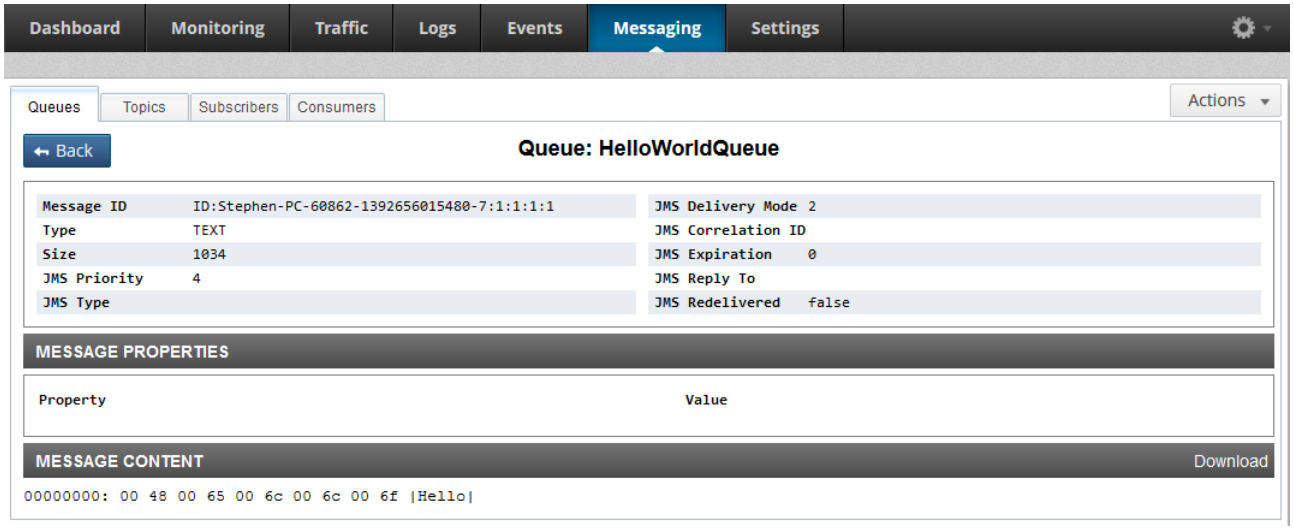

Manage messaging topics

The **Topics** tab displays the messaging topics that exist in the selected API Gateway instance, or across the selected API Gateway group. The **Actions** drop-down menu on the right enables you to perform the following tasks:

- Create a new topic
- Delete a topic
- Refresh the list of topics

To create a new topic, select **Actions** > **Create Topic**, and enter the required topic name in the dialog.

Unlike queues, where messages posted to the queue are kept until someone reads them, messages posted to the topic are broadcasted to all topic subscribers, and then immediately removed from the topic. This means that when you click a newly created topic in the list, the **Actions** menu on the right includes only the **Create Message** action to enable you to post a new message to the topic. Delete or move actions are not required because the topic is always empty.

Manage messaging subscribers

The **Subscribers** tab displays the list of durable subscribers that have registered to receive messages from a specified publisher (for example, messaging topic). In this model, the subscriber and the publisher are not aware of each other. The **Actions** drop-down menu on the right enables you to perform the following tasks:

- Create a new subscriber
- Delete a subscriber
- Refresh the list of subscribers

Create a new subscriber

To create a new subscriber, select **Actions** > **Create Subscriber**, and complete the following details:

Setting	Description
Subscriber Name	Specifies a name for this subscriber (for example `HelloWorldTopicSubscriber`).
Client ID	Specifies a unique client ID used for the subscriber connection. If you do not specify a client ID, one is generated by default. For example: `ID:<hostname>-60862-1392656015480-28:1` The client ID is required for durable subscriptions, which enable a client to disconnect or fall over while consuming a topic, and retrieve any missed messages when it reconnects. To achieve this, the broker needs the client ID to identify which messages are pending consumption.
Destination	Specifies the JMS destination being subscribed to (for example, topic name).
Selector	Specifies a JMS selector used to attach a filter to a subscription and perform content-based routing. JMS selectors are defined using SQL 92 syntax, and typically apply to message headers. For example: `JMSType = 'car' AND color = 'blue' AND weight > 250` For more details, see http://activemq.apache.org/selectors.html.

Note

JMS selectors and filters in ActiveMQ are in no way related to API Gateway selectors and filters. For more details on API Gateway selectors and filters, see the *API Gateway Policy Developer Guide*.

The following screen shows a newly created subscriber:

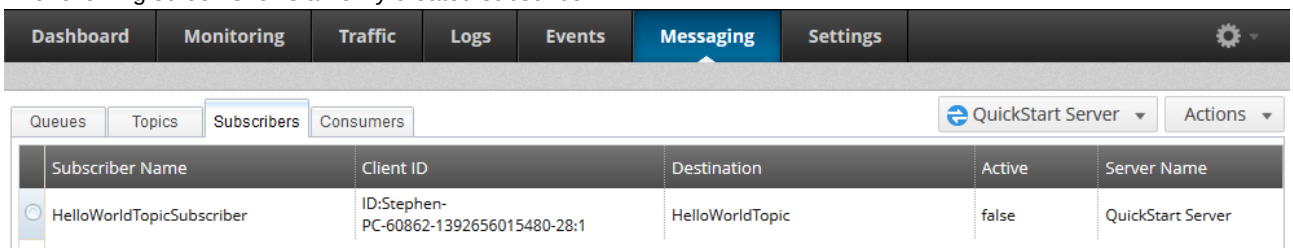

Manage messaging consumers

The **Consumers** tab displays the list of JMS consumers that are currently connected to this embedded ActiveMQ server. This includes details such as client and connection IDs, destination type and name, and server name.

Then when you click a consumer in the list, this displays more detailed information on that consumer. For example:

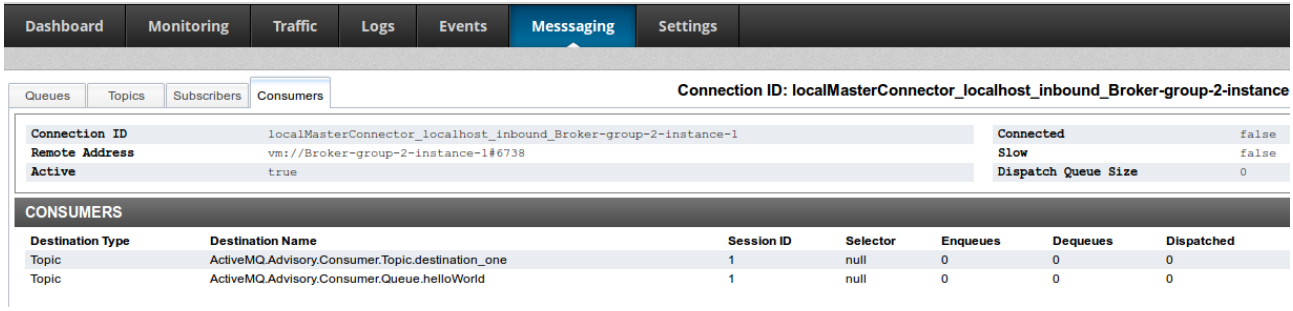

General settings

Overview

The top-level **General** settings screen in Policy Studio enables you to set global configuration settings to optimize the behavior of the API Gateway for your environment.

To configure the **General** settings, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** > **General**. Alternatively, in the Policy Studio tree, select the **Server Settings** node, and click **General**. To confirm updates to these settings, click **Apply changes** at the bottom right of the screen.

After changing any settings, you must deploy to the API Gateway for the changes to be enforced. You can do this in the Policy Studio main menu by selecting **Server** > **Deploy**. Alternatively, click the **Deploy** button in the toolbar, or press F6.

Settings

You can configure the following settings in the **General** screen:

Setting	Description
Tracing level	Enables you to set the trace level for the API Gateway at runtime. Select the appropriate option from the **Trace Level** drop-down list. Defaults to `INFO`.
Active timeout	When the API Gateway receives a large HTTP request, it reads the request off the network when it becomes available. If the time between reading successive blocks of data exceeds the **Active Timeout** specified in milliseconds, the API Gateway closes the connection. This guards against a host closing the connection in the middle of sending data. For example, if the host's network connection is pulled out of the machine while in the middle of sending data to the API Gateway. When the API Gateway has read all the available data off the network, it waits the **Active Timeout** period before closing the connection. Defaults to `30000` milliseconds. You can configure this setting on a per-host basis using the **Remote Hosts** interface. For more details, see the *API Gateway Policy Developer Guide*.
Date format	Configures the format of the date for the purposes of tracing, logging, and reporting. Defaults to `MM.dd.yyyy HH:mm:ss,SSS`. For more information, see: `http://docs.oracle.com/javase/6/docs/api/java/text/SimpleDateFormat.html`
Cache refresh interval	Configures the number of seconds that the server caches data loaded from an external source before refreshing data from that source. Defaults to `5` seconds. To disable the cache, set this to `0`. This cache applies to attributes retrieved from external databases, LDAP directories, internal user stores, and IBM Tivoli. It also applies to query results for authentication against LDAP or databases, and to certificate revocation lists for certificate validation (CRL and XKMS only).
Transaction timeout	A configurable transaction timeout that detects slow HTTP attacks (slow header write, slow body write, slow read) and rejects any transaction that keeps the worker threads occupied for an excessive amount of time. The default value is `240000` milliseconds.
Maximum sent bytes per transaction	The maximum number of bytes sent in a transaction. This is a configurable maximum length for the transmitted data on transactions that API Gateway can handle. This helps to prevent denial-of-service (DoS) attacks. This setting limits the entire amount of data sent over the link, regardless of whether it consists of body, headers, or request line. The default value is `10` MB (`10485760` bytes).
Maximum received	The maximum number of bytes received in a transaction. This is a configurable maximum

Setting	Description
bytes per transaction	length for the received data on transactions that API Gateway can handle. This helps to prevent denial-of-service (DoS) attacks. This setting limits the entire amount of data received over the link, regardless of whether it consists of body, headers, or request line. The default value is `10` MB (`10485760` bytes).
Idle timeout	The API Gateway supports HTTP 1.1 persistent connections. The **Idle Timeout** specified in milliseconds is the time that the API Gateway waits after sending a message over a persistent connection before it closes the connection. Typically, the host tells the API Gateway that it wants to use a persistent connection. The API Gateway acknowledges this instruction and decides to keep the connection open for a certain amount of time after sending the message to the host. If the connection is not reused within the **Idle Timeout** period, the API Gateway closes the connection. Defaults to `15000` milliseconds. You can configure this setting on a per-host basis using the **Remote Hosts** interface. For more details, see the *API Gateway Policy Developer Guide*.
LDAP service provider	Specifies the service provider used for looking up an LDAP server (for example, `com.sun.jndi.ldap.LdapCtxFactory`). The provider is typically used to connect to LDAP directories for certificate and attribute retrieval.
Maximum memory per request	The maximum amount of memory in bytes that the API Gateway can allocate to each request. This setting helps protect against denial of service caused by undue pressure on memory. **Tip** As a general rule for XML messages, if you need to process a message of size N, you should allocate 6–7 times N amount of memory. You also can configure this setting at the HTTP/S interface level. For more details, see the *API Gateway Policy Developer Guide*.
Realm	Specifies the realm for authentication purposes.
Schema pool size	Sets the size of the Schema Parser pool.
Server brand	Specifies the branding to be used in the API Gateway.
SSL session cache size	Specifies the number of idle SSL sessions that can be kept in memory. You can use this setting to improve performance because the slowest part of establishing the SSL connection is cached. A new connection does not need to go through full authentication if it finds its target in the cache. Defaults to `32`. If there are more than 32 simultaneous SSL sessions, this does not prevent another SSL connection from being established, but means that no more SSL sessions are cached. A cache size of `0` means no cache, and no outbound SSL connections are cached.
Token drift time	Specifies the number of seconds drift allowed for WS-Security tokens. This is important in cases where the API Gateway is checking the date on incoming WS-Security tokens. It is likely that the machine on which the token was created is out-of-sync with the machine on which the API Gateway is running. The drift time allows for differences in the respective machine clock times.
Allowed number of operations to limit XPath transforms	Specifies the total number of node operations permitted in XPath transformations. Complex XPath expressions (or those constructed together with content to produce expensive processing) might lead to a denial-of-service risk. Defaults to `4096`.
Input encodings	Click the browse button to specify the HTTP content encodings that the API Gateway instance can accept from peers. The available content encodings include `gzip` and `deflate`.

Setting	Description
	Defaults to no context encodings. For more details, see the *API Gateway Policy Developer Guide*.
Output encodings	Click the browse button to specify the HTTP content encodings that the API Gateway instance can apply to outgoing messages. The available content encodings include `gzip` and `deflate`. Defaults to no context encodings. For more details, see the *API Gateway Policy Developer Guide*.
Server's SSL cert's name must match name of requested server	Ensures that the certificate presented by the server matches the name of the host address being connected to. This prevents host spoofing and man-in-the-middle attacks. This setting is enabled by default.
Send desired server name to server during TLS negotiation	Specifies whether to add a field to outbound TLS/SSL calls that shows the name that the client used to connect. For example, this can be useful if the server handles several different domains, and needs to present different certificates depending on the name that the client used to connect. This setting is not selected by default.
Add correlation ID to outbound headers	Specifies whether to insert the correlation ID in outbound messages. For the HTTP transport, this means that an `X-CorrelationID` header is added to the outbound message. This is a transaction ID that is tagged to each message transaction that passes through the API Gateway, and which is used for traffic monitoring in the API Gateway Manager web console. You can use the correlation ID to search for messages in the console. You can also access the its value using the `id` message attribute in an API Gateway policy. An example correlation ID value is `Id-54bbc74f515d52d71a4c0000`. This setting is selected by default.

MIME/DIME settings

Overview

The MIME/DIME settings list a number of default common content types that are used when transmitting MIME messages. You can configure the API Gateway's **Content Type** filter to accept or block messages containing specific MIME types. Therefore, the contents of the MIME types library act as the set of all MIME types that the API Gateway can filter messages with.

All of the MIME types listed in the table are available for selection in the **Content Type** filter. For example, you can configure this filter to accept only XML-based types, such as `application/xml`, `application/*+xml`, `text/xml`, and so on. Similarly, you can block certain MIME types (for example, `application/zip`, `application/octet-stream`, and `video/mpeg`). For more details on configuring this filter, see the *API Gateway Policy Developer Guide*.

Configuration

To configure the MIME/DIME settings, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** > **General** > **MIME/DIME**. Alternatively, in the Policy Studio tree, select the **Server Settings** node, and click **General** > **MIME/DIME**. To confirm updates to these settings, click **Apply changes** at the bottom right of the screen.

The MIME/DIME settings screen lists the actual MIME types on the left column of the table, together with their corresponding file extensions (where applicable) in the right column.

To add a new MIME type, click the **Add** button. In the **Configure MIME/DIME Type** dialog, enter the new content type in the **MIME or DIME Type** field. If the new type has a corresponding file extension, enter this extension in the **Extension** field. Click the **OK** button when finished.

Similarly, you can edit or delete existing types using the **Edit** and **Delete** buttons.

Namespace settings

Overview

The API Gateway exposes global settings that enable you to configure which versions of the SOAP and WSSE specifications it supports. You can also specify which attribute is used to identify the XML Signature referenced in a SOAP message.

To configure the namespace settings, in the Policy Studio tree, select the **Server Settings** node, and click **General** > **Namespaces**. Alternatively, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** > **General** > **Namespaces**. To confirm updates to these settings, click **Apply changes** at the bottom right of the screen.

SOAP Namespace

The **SOAP Namespace** tab can be used to configure the SOAP namespaces that are supported by the API Gateway. In a similar manner to the way in which the API Gateway handles WSSE namespaces, the API Gateway will attempt to identify SOAP messages belonging to the listed namespaces in the order given in the table.

The default behavior is to attempt to identify SOAP 1.1 messages first, and for this reason, the SOAP 1.1 namespace is listed first in the table. The API Gateway will only attempt to identify the message as a SOAP 1.2 message if it can't be categorized as a SOAP 1.1 message first.

Signature ID Attribute

The **Signature ID Attribute** tab allows you to list the supported attributes that can be used by the API Gateway to identify a Signature reference within an XML message.

An XML-signature `<signedInfo>` section may reference signed data via the `URI` attribute. The `URI` value may contain an id that identifies data in the message. The referenced data will hold the "URI" field value in one of its attributes.

By default, the server uses the `Id` attribute for each of the WSSE namespaces listed above to locate referenced signed data. The following sample XML Signature illustrates the use of the `Id` attribute:

```
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
 <soap:Header>
  <dsig:Signature id="Sample" xmlns:dsig="http://www.w3.org/2000/09/xmldsig#">
   <dsig:SignedInfo>
    ...
    <dsig:Reference URI="#Oracle:sLmDCph3tGZ10">
      ...
    </dsig:Reference>
   </dsig:SignedInfo>
    ....
  </dsig:Signature>
 </soap:Header>
 <soap:Body>
  <getProduct wsu:Id="Oracle:sLmDCph3tGZ10"
      xmlns:wsu="http://schemas.xmlsoap.org/ws/2003/06/utility">
   <Name>SOA Test Client</Name>
   <Company>Company</Company>
  </getProduct>
 </soap:Body>
</soap:Envelope>
```

It is clear from this example that the Signature reference identified by the `URI` attribute of the `<Reference>` element refers to the nodeset identified with the `Id` attribute (the `<getProduct>` block).

Because different toolkits and implementations of the XML-Signature specification can use attributes other than the `Id`

attribute, the API Gateway allows the user to specify other attributes that should be supported in this manner. By default, the API Gateway supports the `Id`, `ID`, and `AssertionID` attributes for the purposes of identifying the signed content within an XML Signature.

However you can add more attributes by clicking the **Add** button and adding the attribute in the interface provided. The priorities of attributes can be altered by clicking the **Up** and **Down** buttons. For example, if most of the XML Signatures processed by the API Gateway use the `ID` attribute, this attribute should be given the highest priority.

WSSE Namespace

The **WSSE Namespace** tab is used to specify the WSSE (and corresponding WSSU) namespaces that are supported by the API Gateway.

The API Gateway attempts to identify WS Security blocks belonging to the WSSE namespaces listed in this table. It first attempts to locate Security blocks belonging to the first listed namespace, followed by the second, then the third, and so on until all namespaces have been utilized. If no Security blocks can be found for any of the listed namespaces, the message will be rejected on the grounds that the API Gateway does not support the namespace specified in the message. To add a new namespace, click the add button.

Note

Every WSSE namespace has a corresponding WSSU namespace. For example, the following WSSE and WSSU namespaces are inextricably bound:

WSSE Namespace	`http://schemas.xmlsoap.org/ws/2003/06/secext`
WSSU Namespace	`http://schemas.xmlsoap.org/ws/2003/06/utility`

First, enter the WSSE namespace in the **Name** field. Then enter the corresponding WSSU namespace in the **WSSU Namespace** field.

HTTP Session settings

Overview

The **HTTP Session** settings enable you to configure session management settings for the selected cache. For example, you can configure the period of time before expired sessions are cleared from the `HTTP Sessions` cache, which is selected by default.

To configure HTTP session settings, select the **Server Settings** node in the Policy Studio tree, and click **General** > **HTTP Session**. Alternatively, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** > **General** > **HTTP Session**. To confirm updates to these settings, click **Apply changes** at the bottom right of the screen.

Configuration

Configure the following session settings:

Cache:
Specifies the cache that you wish to configure. Defaults to `HTTP Sessions`. To configure a different cache, click the button on the right, and select the cache to use. The list of currently configured caches is displayed in the tree.

To add a cache, right-click the **Caches** tree node, and select **Add Local Cache** or **Add Distributed Cache**. Alternatively, you can configure caches under the **Libraries** node in the Policy Studio tree. For more details, see the *API Gateway Policy Developer Guide*.

Clear Expired Sessions Period:
Enter the number of seconds before expired sessions are cleared from the selected cache. Defaults to `60`.

Transaction audit log settings

Overview

One of the most important features of a server-based product is its ability to maintain highly detailed and configurable logging. It is crucial that a record of each and every transaction is kept, and that these records can easily be queried by an administrator to carry out detailed transaction analysis. In recognition of this requirement, the API Gateway provides detailed logging to a number of possible locations.

You can configure the API Gateway so that it logs information about all requests. Such information includes the request itself, the time of the request, where the request was routed to, and the response that was returned to the client. The logging information can be written to the console, log file, local/remote syslog, and/or a database, depending on what is configured in the logging settings.

The API Gateway can also digitally sign the logging information it sends to the log files and the database. This means that the logging information can not be altered after it has been signed, thus enabling an irreversible audit trail to be created.

Warning

The transaction audit log includes the complete contents of HTTP requests, including HTTP headers, body, and attachments. This may include sensitive information. You must ensure that appropriate safeguards are in place to protect this information in the different audit log locations.

Configure log output

To edit the default API Gateway logging settings in the Policy Studio tree, select the **Server Settings** node, and click **Logging** > **Transaction Audit Log**. Alternatively, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** .

You can configure the API Gateway to log to the following locations described in this topic.

Log to Text File

To configure the API Gateway to log in text format to a file, click the **Text File** tab, and select the **Enable logging to file** checkbox. You can configure the following fields:

- **File Name**:
  Enter the name of the text-based file that the API Gateway logs to. The default is `transactionLog`.
- **File Extension**:
  Enter the file extension of the log file. Defaults to `.log`.
- **Directory**:
  Enter the directory of the log file in this field. By default, all log files are stored in the `/logs/transaction` directory of your API Gateway installation.
- **File Size (MB)**:
  Enter the maximum size that the log file grows to. When the file reaches the specified limit, a new log file is created. By default, the maximum file size is 1000 MB.
- **Roll Log Daily**:
  Specify whether to roll over the log file at the start of each day. This is enabled by default.
- **Number of Files**:
  Specify the number of log files that are stored. The default number is 20.
- **Format**:
  You can specify the format of the logging output using the values entered here. You can use selectors to output logging information that is specific to the request. The default logging format is as follows:

```
${level} ${timestamp} ${id} ${text} ${filterType} ${filterName}
```

The available logging properties are described as follows:

- **`level`**:
  The log level (`fatal`, `fail`, `success`).
- **`timestamp`**:
  The time that the message was processed in user-readable form.
- **`id`**:
  The unique transaction ID assigned to the message.
- **`text`**:
  The text of the log message that was configured in the filter itself. In the case of the **Log Message Payload** filter, the `${payload}` selector contains the message that was sent by the client.
- **`filterName`**:
  The name of the filter that generated the log message.
- **`filterType`**:
  The type of the filter that logged the message.
- **`ip:`**
  The IP address of the client that sent the request.

- **Signing Key**:
  To sign the log file, select a **Signing Key** from the Certificates Store that is used in the signing process. By signing the log files, you can verify their integrity at a later stage.

To confirm updates to these settings, click **Apply changes** at the bottom right of the screen.

Log to XML File

To configure the API Gateway to log to an XML file, click the **XML File** tab, and select the **Enable logging to XML file** checkbox. The log entries are written as the values of XML elements in this file. You can view historical XML log files (not the current file) as HTML for convenience by opening the XML file in your default browser. The `/logs/xsl/MessageLog.xsl` stylesheet is used to render the XML log entries in a more user-friendly HTML format.

You can configure the following fields on the **XML File** tab:

- **File Name**:
  Enter the name of the text-based file that the API Gateway logs to. By default, the log file is called `oracle`.
- **File Extension**:
  Enter the file extension of the log file in this field. By default, the log file is given the `.log` extension.
- **Directory**:
  Enter the directory of the log file in this field. By default, all log files are stored in the `/logs/transaction` directory of your API Gateway installation.
- **File Size**:
  Enter the maximum size that the log file grows to. When the file reaches the specified limit, a new log file is created. By default, the maximum file size is 1000 kilobytes.
- **Roll Log Daily**:
  Specify whether to roll over the log file at the start of each day. This is enabled by default.
- **Number of Files**:
  Specify the number of log files that are persisted. The default number is 20.
- **Signing Key**:
  To sign the log file, select a **Signing Key** from the Certificates Store that will be used in the signing process. By signing the log files, you can verify their integrity at a later stage.

Log to Database

Using this option, you can configure the API Gateway to log messages to an Oracle, SQL Server, or MySQL relational database.

Note

Before configuring the API Gateway to log to a database, you must first create the database tables that the API Gateway writes to. For details on setting up tables for supported databases, see the *API Gateway Installation and Configuration Guide*.

When you have set up the logging database tables, you can configure the API Gateway to log to the database. Click the **Database** tab, and select **Enable logging to database**. You can configure the following fields on the **Database** tab:

- **Connection**:
  Select an existing database from the **Connection** drop-down list. To add a database connection, click the **External Connections** button on the left, right-click the **Database Connections** tree node, and select **Add a Database Connection**. For more details, see the *API Gateway Policy Developer Guide*.
- **Signing Key**:
  You can sign log messages stored in the database to ensure that they are not tampered with. Click **Signing Key** to open the list of certificates in the Certificate Store, and select the key to use to sign log messages.

Log to Local Syslog

To configure the API Gateway to send logging information to the local UNIX syslog, click the **Local Syslog** tab, and select the **Enable logging to local UNIX Syslog** checkbox. You can configure the following fields:

- **Select Syslog server**:
  Select the local syslog facility that the API Gateway should log to. The default is `LOCAL0`.
- **Format**:
  You can specify the format of the log message using the values (including selectors) entered in this field. For details on the properties that are available, see the section called “Log to Text File”.

Log to Remote Syslog

To configure the API Gateway to send logging information to a remote syslog, click the **Remote Syslog** tab, and select the **Enable logging to Remote Syslog** checkbox. You can configure the following fields:

- **Syslog Server**
  Select a previously configured **Syslog Server** from the list. For details on how to configure Syslog Server, see the chapter on External Connections in the *API Gateway Policy Developer Guide*.
- **Format**:
  You can specify the format of the log message using the values (including properties) entered in this field. For details on the properties that are available, see the section called “Log to Text File”.

Log to System Console

To configure the API Gateway to send logging information to the system console, click the **System Console** tab, and select the **Enable logging to system console** checkbox. For details on how to use the **Format** field to configure the format of the log message, see the section called “Log to Text File”.

Transaction access log settings

Overview

The access log records a summary of the request and response messages that pass through the API Gateway. By default, the API Gateway records this in the `access.log` file in the `log` directory. This file rolls over with a version number added for each new version of the file (for example, `access.log.0`, `access.log.1`, and so on).

The transaction access log file format is based on that used by Apache HTTP Server. This means that the log file can be consumed by third-party Web analytics tools such as Webtrends to generate charts and statistics.

Access log format

The syntax used to specify the access log file is based on the syntax of available patterns used by the access log files in Apache HTTP Server. For example:

```
%h %l %u %t "%r" %s %b
```

The log format strings in this example are explained as follows:

%h	Remote hostname or IP address.
%l	Remote logical username.
%u	Remote user that was authenticated (for example, Distinguished Name of a certificate).
%t	Date and time of the request in Common Log Format.
%r	First line of the request that originated at the client.
%s	HTTP status code returned to the client in the response.
%b	Bytes sent, excluding HTTP headers.

The following extract from the `access.log` file illustrates the resulting log format:

```
s1.oracle.com - lisa [09/05/2012:18:24:48 00] "POST / HTTP/1.0" 200 429
s2.oracle.com - dave [09/05/2012:18:25:26 00] "POST / HTTP/1.0" 200 727
s3.oracle.com - fred [09/05/2012:18:27:12 00] "POST / HTTP/1.0" 200 596
.................
.................
.................
```

For more details on Apache HTTP Server access log formats, see the following:

- http://httpd.apache.org/docs/current/logs.html
- http://httpd.apache.org/docs/current/mod/mod_log_config.html#formats

Configure the access log

To configure the access log in the Policy Studio tree, select the **Server Settings** node, and click **Logging** > **Transaction Access Log**. Alternatively, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** > **Logging** > **Transaction Access Log**. To confirm updates to these settings, click **Apply changes** at the bottom right of the window.

You can configure the following fields to enable the server to write an access log to file:

Writing to Transaction Access Log:
Select whether to configure the API Gateway instance to start writing event data to the transaction access log. This setting is disabled by default.

File name:
Enter the name of the access log file. When the file rolls over (because the maximum file size has been reached, or because the date has changed), a suitable increment is appended to the file name. Defaults to `access`.

File extension:
Enter the file extension for the log file. Defaults to `.log`.

Directory:
Enter the directory for the access log file. Defaults to the `logs/access` directory of your product installation.

File size (MB):
Specify the maximum size that the log file is allowed reach before it rolls over to a new file. Defaults to `1000` MB.

Roll log daily:
Select whether to roll over the log file at the start of each day. This is enabled by default.

Number of log files:
Specify the number of log files that are stored. Defaults to `20`.

Format:
Enter the access log file format. This is based on the syntax used in Apache HTTP Server access log files, for example:

```
%h %l %u %t "%r" %s %b
```

For more details, see the section called “Access log format”.

Important

These settings configure the access log at the API Gateway level. You must also configure the access log at the service level on a specific relative path.

For example, in the Policy Studio tree, select the relative path, right-click it in the **Resolvers** pane, and select **Edit**. Then click the **Logging Settings** tab, and select **Include in server access log records**. For more details, see the *API Gateway Policy Developer Guide*.

Redact sensitive details from the access log

The default syntax for the access log is as follows:

```
%h %l %u %t "%r" %s %b
```

The `%r` format string results in the entire HTTP request line being added to the access log file, including the query string. For example:

```
127.0.0.1 - - [02/07/2014:12:39:29 00]  "POST /healthcheck?name=value HTTP/1.0" 200 19
```

The query string may contain sensitive information (for example, credit card number, or social security number). If you do not wish the query string to be included in the access log, it is recommended that you use the following format instead:

```
%h %l %u %t "%m %U% %H" %s %b
```

For example, this results in the following output instead:

```
127.0.0.1 - - [02/07/2014:12:39:29 00]  "POST /healthcheck HTTP/1.0" 200 19
```

The "`%m %U %H`" options log the method, path, and HTTP version. This results in the same output as `%r`, but without the query string.

To confirm updates to these settings, click **Apply changes** at the bottom right of the screen. Click **Deploy** in the toolbar to deploy the updated configuration to the API Gateway.

Transaction event log settings

Overview

The **Transaction Event Log** provides a summary of each API Gateway message transaction, which is written to a log file, and used to generate metrics for API Gateway monitoring (for example, in the **Monitoring** view in API Gateway Analytics). In a distributed system with multiple API Gateway instances running, the events data is written to separate transaction event log files for each API Gateway instance.

The event log file data is processed by the local Node Manager every 5 minutes, aggregated into the appropriate metrics data, and then written to a database. API Gateway Analytics can use the data from the database to report on analytics in the system. Event log file data is written in JSON format, which also enables it to be integrated with third-party logging tools such as Splunk.

Note

Node Manager processing of event log data is not enabled by default. You must enable the Node Manager to write metrics to the database. For more details, see the section called “Configure the Node Manager to process event logs and update the metrics database”.

For more details how metrics are displayed, see the following:

- *Monitoring and reporting with API Gateway Analytics*

Transaction event log formats

Event log files are located in the `events` directory of your API Gateway installation by default. For example:

```
INSTALL-DIR/apigateway/events/group-2_instance-1.log
```

When each event log file has been processed (every 5 minutes), it can be moved to a `processed` directory. For example:

```
INSTALL-DIR/apigateway/events/processed
```

By default, files are deleted after being processed.

Entries in the transaction event log file are generated for different event types (for example, `header`, `system`, `transaction`, and `custom`).

Event log header entries

Event log `header` entries contain details about the creation of the log file. For example, this includes when the log file is created, and on which host, domain, group, instance, and so on.

The following example shows the JSON format used for `header` events:

```
{
    "type": "header",
    "logCreationTime": "2015-01-23 12:25:00.120",
    "hostname": "Stephen-PC",
    "domainId": "cfbe55d1-be45-4968-8b4b-f06a4db858b8",
    "groupId": "group-2",
    "groupName": "QuickStart Group",
    "serviceId": "instance-1",
    "serviceName": "QuickStart Server",
```

```
    "version": "v7.4.0-Internal"
}
...
```

Event log system entries

Event log `system` entries contain details about the API Gateway system. For example, this includes details such as the amounts of disk space, memory, and CPU.

The following example shows the JSON format used for `system` events:

```
{
    "type": "system",
    "time": 1422015900120,
    "diskUsed": 30,
    "instUsed": 1,
    "sysCpu": 5,
    "instMem": 533436,
    "sysMem": 4641996,
    "sysMemTotal": 16759240
}
...
```

Event log transaction entries

Event log `transaction` entries contain details about a specific message transaction. For example, this includes details such as the protocol, method, bytes sent and received, IP addresses, ports, service name, and so on.

The following example shows the JSON format used for an HTTP `transaction` event with a service context and inbound and outbound transaction legs:

```
{
    "type": "transaction",
    "time": 1425291330502,
    "path": "/stockquote.asmx",
    "protocol": "http",
    "protocolSrc": "8080",
    "duration": 1842,
    "status": "success",
    "serviceContexts": [
        {
            "service": "StockQuote",
            "monitor": true,
            "client": null,
            "org": null,
            "app": null,
            "method": "GetQuote",
            "status": "success",
            "duration": 1824
        }
    ],
    "customMsgAtts": {},
    "correlationId": "4038f4540400788ebe4f84ca",
    "legs": [
        {
            "uri": "/stockquote.asmx",
            "status": 200,
            "statustext": "OK",
            "method": "POST",
            "vhost": null,
            "wafStatus": 0,
```

```
            "bytesSent": 1278,
            "bytesReceived": 612,
            "remoteName": "127.0.0.1",
            "remoteAddr": "127.0.0.1",
            "localAddr": "127.0.0.1",
            "remotePort": "49104",
            "localPort": "8080",
            "sslsubject": null,
            "leg": 0,
            "timestamp": 1425291328660,
            "duration": 1843,
            "serviceName": "StockQuote",
            "subject": null,
            "operation": "GetQuote",
            "type": "http",
            "finalStatus": "Pass"
        },
        {
            "uri": "/stockquote.asmx",
            "status": 200,
            "statustext": "OK",
            "method": "POST",
            "vhost": null,
            "wafStatus": 0,
            "bytesSent": 736,
            "bytesReceived": 1202,
            "remoteName": "www.webservicex.net",
            "remoteAddr": "173.201.44.188",
            "localAddr": "10.142.10.142",
            "remotePort": "80",
            "localPort": "49438",
            "sslsubject": null,
            "leg": 1,
            "timestamp": 1425291329916,
            "duration": 566,
            "serviceName": "StockQuote",
            "subject": null,
            "operation": "GetQuote",
            "type": "http",
            "finalStatus": null
        }
    ]
}
...
```


Configure the transaction event log

To configure the transaction event log in the Policy Studio tree, select the **Server Settings** node, and click **Logging** > **Transaction Event Log**. Alternatively, in the Policy Studio main menu, you can select **Tasks** > **Manage Gateway Settings**.

Configure the following fields to enable the API Gateway instance to write a transaction event log to a file:

Writing to Transaction Event Log:
Enables writing to an event log for all message transactions received by the API Gateway. This setting is enabled by default, and is required for API Gateway Analytics. For example, you could deselect this setting to optimize performance.

Write transaction event logs to directory:
Specifies the directory where transaction event logs are written. Defaults to `${environment.VDISTDIR}/events`.

Note

If transaction event logs are being used to populate the metrics database, you must also update the `sourceEventLogDir` property in the Node Manager configuration if you change this directory. For more details, see *Monitoring and reporting with API Gateway Analytics*.

System event frequency (secs):
Specifies how often in seconds that a system entry is written to each event log file. Defaults to `60` seconds. For more details, see the section called “Event log system entries”.

Maximum disk space for event logs (MB):
Specifies the maximum amount of disk space used for event logs. When the directory reaches the specified limit, the oldest log files are deleted. Defaults to `1024` MB.

Check disk space interval (secs):
Specifies how often the amount of available disk space used for event logs is checked. Defaults to `600` seconds.

Select the message attributes to be stored in transaction events:
Enables you to specify custom message attributes to write to the transaction event logs (for example, the HTTP request URI). To specify an attribute, click **Add**, and enter the attribute name in the dialog.

Note

If the local Node Manager is configured to write transaction event logs to the metrics database, these attributes are written to the `transaction_data` table. Data from this table is not displayed in API Gateway Analytics. However, you can access this data using SQL queries.

You can disable writing of this data to the database by configuring the `processCustomMessageAttributes` option in the `esexplorer` tool. For details, see the section called “Configure additional options for event log processing in the Node Manager”.

To confirm updates to these settings, click **Apply changes** at the bottom right of the window. Click **Deploy** in the toolbar to deploy the updated configuration to the API Gateway.

Embedded ActiveMQ settings

Overview

The **Embedded ActiveMQ** settings enable you to configure settings for the Apache ActiveMQ messaging broker that is embedded in each API Gateway instance. You can also configure multiple embedded ActiveMQ brokers to work together as a network of brokers in a group of API Gateway instances. For more details on Apache ActiveMQ, see http://activemq.apache.org/.

To configure Embedded ActiveMQ settings, select the **Server Settings** node in the Policy Studio tree, and click **Messaging** > **Embedded ActiveMQ**. Alternatively, in the Policy Studio main menu, select **Tasks** > **Manage Gateway Settings** > **Messaging** > **Embedded ActiveMQ**. To apply updates to these settings, click **Apply changes** at the bottom right of the screen.

General messaging settings

Configure the following ActiveMQ messaging settings:

Enable Embedded ActiveMQ Broker:
Specifies whether to enable starting up the ActiveMQ broker that is embedded in the API Gateway instance. This is not selected by default.

Address:
Specifies the IP address used to open a listening socket for incoming ActiveMQ connections. Defaults to `0.0.0.0`, which specifies that all interface addresses should be used.

Port:
Specifies the TCP port for incoming ActiveMQ connections. Defaults to `${env.BROKER.PORT}`, which enables the port number to be environmentalized. This means that the port number is specified in the `envSettings.props` file on a per-server basis. For more details, see the *Deployment and Promotion Guide*. Alternatively, you can enter the port number directly in this field (for example, `61616`).

Shared Directory:
Specifies the location of the shared directory in your environment that is used by multiple embedded ActiveMQ brokers. This setting is required, and must be configured for high availability and failover. Defaults to `<install-dir>/messaging-shared`.

SSL settings

Configure the following settings to secure the communication with JMS clients, and between multiple embedded ActiveMQ brokers:

Enable SSL:
Specifies whether to use Secure Sockets Layer (SSL) to secure the communication with JMS clients, and between ActiveMQ brokers.

Server Cert:
When **Enable SSL** is selected, click to select the server certificate with a private key that is used for SSL communication between ActiveMQ brokers. For details on importing certificates into the certificate store, see *Manage certificates and keys*.

Accepted cipher suites:
When **Enable SSL** is selected, select which cipher suites should be accepted by the JMS server when the SSL communication is being established.

Note
If no cipher suites are selected, the default cipher suites from the Java Security Socket Extension (JSSE) are used.

Require Client Certificates:
When **Enable SSL** is set, specifies whether to require client certificates for client SSL authentication. For example, for mutual (two-way) SSL communication, you must trust the issuer of the client certificate by importing the client certificate issuer into the certificate store. For details on importing certificates, see *Manage certificates and keys*.

When **Require client certificates** is selected, you can then select the root certificate authorities that are trusted for mutual (two-way) SSL communication between ActiveMQ brokers. For details on importing certificates into the API Gateway certificate store, see *Manage certificates and keys*.

Authentication settings

Configure the following to specify authentication settings between multiple embedded ActiveMQ brokers:

Note
The authentication settings are also used by features on the **Messaging** tab in the API Gateway Manager web console (for example, sending messages and managing durable topic subscriptions). For more details, see *Manage embedded ActiveMQ messaging*.

Authenticate broker and client connections with the following policy:
When username/password credentials are provided for inter-broker communication, they are delegated to the selected policy for authentication. By default, no policy is selected. To select a policy, click the button on the right, and select a pre-configured policy in the dialog.

Username:
Specifies the username credential when connecting to other ActiveMQ brokers.

Password:
Specifies the password credential when connecting to other ActiveMQ brokers.

Communicate with brokers in the same group:
Every API Gateway instance belongs to a group. This setting specifies whether to communicate only with ActiveMQ brokers in the same API Gateway group. This is the default setting.

Or brokers outside the group registered with the following alias:
Specifies an alias name used to communicate with other ActiveMQ brokers registered with the same alias. This setting enables communication with ActiveMQ brokers that belong to different API Gateway groups.

Traffic monitoring settings

Overview

The **Traffic Monitor** settings enable you to configure the traffic monitoring available in the web-based API Gateway Manager tool. For example, you can configure where the data is stored and what message transaction details are recorded in the message traffic log.

To access the **Traffic Monitor** settings, click the **Sever Settings** node in the Policy Studio tree, and click **Monitoring** > **Traffic Monitor**. Alternatively, in the main menu select **Tasks** > **Manage Gateway Settings**, and click **Monitoring** > **Traffic Monitor**. To confirm updates to these settings, click **Apply changes** at the bottom right of the screen.

To access traffic monitoring in API Gateway Manager, go to `http://localhost:8090`, and click the **Traffic** button in the toolbar.

Configuration

You can configure the following **Traffic Monitor** settings:

Enable Traffic Monitor:
Select whether to enable the web-based Traffic Monitor tool. This is enabled by default.

Transaction Store Location Settings:
Enter the **Transaction Directory**: that stores the traffic monitoring data files and database. You must enter the location relative to the API Gateway instance directory, which is the location where the server instance runs (for example, `<install-dir>/apigateway/groups/<group-id>/<instance-id>`). The **Transaction Directory** defaults to `conf/opsdb.d`. If this directory or the database does not exist when the API Gateway starts up, they are recreated automatically.

Persistence Settings:
You can configure the following data persistence settings:

Record inbound transactions	Select whether to record inbound message transactions received by the API Gateway. This is enabled by default.
Record outbound transactions	Select whether to record outbound message transactions sent from the API Gateway to remote hosts. This is enabled by default.
Record policy path	Select whether to record the policy path for the message transaction, which shows the filters that the message passes through. This is enabled by default.
Record trace	Select whether to record the trace output for the message transaction. This is enabled by default.
Record sent data for transactions	Select whether to record the sent payload data for the message transaction. This is enabled by default.
Record received data for transactions	Select whether to record the received payload data for the message transaction. This is enabled by default.

Note

These settings are global for all traffic passing through the API Gateway. You can override these persistence settings at the port level when configuring an HTTP or HTTPS interface. For more details, see the *API Gateway Policy Developer Guide*.

Details of inbound and outbound transactions are also written to the transaction event log. If recording of inbound or outbound transactions is disabled in **Traffic Monitor** settings, transaction data will not be written to the event log. For more details, see *Transaction event log settings*.

Transaction Expiry Settings:
The following settings configure when a new transaction file is created. This occurs when the maximum transaction file size is exceeded or when the maximum number of transactions is reached (whichever occurs first).

Target transaction file size	Enter the target size of the transaction file, and select the units from the list. Defaults to 8 MB. When this limit is exceeded, no new transactions are added to the file, and a new file is created. **Note** Transaction files may exceed this limit due to continuing traffic on existing transactions. The current file is closed only when all current transactions using the file are complete. This means that the specified file size may be exceeded, but no more transactions are stored in the file when this limit is reached. For example, with the default of 8 MB, assume the current file size is 7 MB, and a 100 MB transaction is being logged, making a total of 107 MB. No more transactions are logged in this file, and a new file is created.
Maximum number of transactions per file	Enter the maximum number of transactions stored in the file. Defaults to `100000` transactions. When this limit is reached, no new transactions are written. The current file is closed when all current transactions using the file are completed.
Maximum number of transaction files	Enter the maximum number of transaction files to store on disk at any one time. Defaults to `128`. When this limit is reached, old files that have no open transactions are purged. **Note** The number of transaction files may be exceeded if the oldest file still has open transactions.

Real-time monitoring metrics

Overview

You can configure real-time monitoring metrics for an API Gateway instance. For example, this enables you to specify monitoring of messages at the level of API services, methods, clients, and remote hosts. This is important when managing APIs because of requirements to bill clients for their API usage.

When real-time monitoring is enabled, monitoring data is stored in API Gateway memory and displayed in the API Gateway Manager web console. API Gateway Manager uses the configured real-time monitoring metrics to display graphical reports in its **Dashboard** and **Monitoring** views. For more details on viewing real-time metrics, see *Monitor services in API Gateway Manager*.

To configure real-time monitoring settings in the Policy Studio tree, select the **Server Settings** node, and click **Monitoring** > **Real Time Monitoring**. Alternatively, in the main menu, you can select **Tasks** > **Manage Gateway Settings**.

Enable monitoring

Configure the following general settings:

Enable real-time monitoring:
This enables real-time monitoring globally for the API Gateway instance in the API Gateway Manager web console. This setting must be enabled to display monitoring data in the **Dashboard** and **Monitoring** views in API Gateway Manager, and is selected by default. To disable real-time monitoring, deselect this setting.

System metrics update frequency (secs):
Specifies how often in seconds that system metrics are measured (for example, CPU, disk space, and memory usage). Defaults to 3 seconds.

Configure real-time metrics

Configure the following settings in the **Real Time Monitoring Limits** section:

Note

Real-time monitoring may have a negative impact on API Gateway performance. To optimize performance, disable monitoring for one or more metrics.

You should set the maximum services, methods, and remote hosts to values that will never be reached in normal operation. These settings protect the API Gateway by setting an upper limit on the amount of memory consumed by real-time monitoring.

Service:
Enables real-time monitoring of metrics data on the **API Services** tab. This is enabled by default.

Maximum Services:
Specifies the maximum number of API services that are monitored by the API Gateway. When the maximum is reached, the API Gateway stops collecting metrics for new services. Defaults to `10000`.

Method:
Enables real-time monitoring of metrics data on the **API Methods** tab. This is enabled by default.

Note

To enable method monitoring, you must ensure that service monitoring is also enabled. Disabling service monitoring also disables method monitoring.

Maximum Methods:
Specifies the maximum number of API methods that are monitored by the API Gateway. When the maximum is reached, the API Gateway stops collecting metrics for new methods. Defaults to `100000`.

Remote Host:
Enables real-time monitoring of metrics data on the **Remote Hosts** tab. This is enabled by default. For more details on remote hosts, see the *API Gateway Policy Developer Guide*.

Maximum Remote Hosts:
Specifies the maximum number of remote hosts that are monitored by the API Gateway. When the maximum is reached, the API Gateway stops collecting metrics for new remote hosts. Defaults to `10000`.

Client:
Enables real-time monitoring of metrics data on the **Clients** tab. This is enabled by default.

Maximum Clients:
Specifies the maximum number of clients that are monitored by the API Gateway. When the maximum is reached, the API Gateway stops collecting metrics for new clients. Defaults to `10000`.

Note

The number of unique clients that communicate with an API Gateway is potentially unbounded. The maximum number of clients is therefore a soft limit. When this is reached, monitoring stops for the oldest client and begins for the newest client instead.

For the other maximum values (services, methods, and remote hosts), exceptions are thrown and logged when the limits are reached.

To confirm updates to these settings, click **Apply changes** at the bottom right of the window. Click **Deploy** in the toolbar to deploy the updated configuration to the API Gateway.

Important

You must restart the API Gateway instance after changing any of the maximum values (for example, **Maximum Services**, **Maximum Methods**, **Maximum Clients**, or **Maximum Remote Hosts**).