

Universal Messaging Concepts

Version 10.1

October 2017

This document applies to Universal Messaging Version 10.1 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2013-2018 Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Table of Contents

Overview.....	7
Architecture.....	9
Architecture Overview.....	10
Support for Open Standards.....	11
Realms.....	11
Zones.....	12
Overview of Zones.....	12
Administration of Zones.....	14
Examples of Message Propagation.....	15
Messaging Paradigms supported.....	17
umTransport API.....	20
Communication Protocols and RNAMEs.....	27
Management.....	29
Administration and Management.....	30
JMX Administration and Management.....	31
Deployment.....	37
Deployment.....	38
Server.....	39
Performance and Tuning.....	39
Server Failover / High Availability.....	40
Data Routing.....	40
Federation Of Servers.....	41
Proxy Servers and Firewalls.....	42
Server Memory for Deployment.....	42
Server Parameters.....	42
Server Security for Deployment.....	48
Deployment.....	49
Connecting to multiple realms using SSL.....	49
The Java Virtual Machine.....	52
The Network.....	56
The Operating System.....	58
The Realm Server.....	59
Periodic Logging of Server Status.....	61
Client.....	63
Using HTTP/HTTPS.....	63
Browser / Applet Deployment.....	64
Browser Plugins.....	65
Client Jars.....	65
Client Security.....	67

Client Parameters.....	67
Multiplexing Sessions.....	69
Language Deployment Tips.....	69
Silverlight Application Deployment.....	69
JavaScript Application Deployment.....	70
Security.....	71
Security Overview.....	72
Authentication.....	72
Authentication Overview.....	72
Using SASL.....	88
Overview.....	88
Client-side Authentication.....	89
Server-side Authentication.....	90
Client Negotiation.....	90
Directory Backend.....	90
Internal User Repository.....	91
LDAP.....	91
Server JAAS Authentication with Software AG Security infrastructure component.....	92
Converting a .jks Key Store to a .pem Key Store.....	95
Access Control Lists.....	95
Security Policies.....	95
Access Control Lists (ACLs).....	95
Using SSL.....	98
SSL Concepts.....	98
SSL Encryption.....	99
Client SSL Configuration.....	100
JMS Client SSL Configuration.....	102
Setting up and using FIPS.....	104
Performance, Scalability and Resilience.....	109
Overview of Performance, Scalability and Resilience.....	110
Clustered Server Concepts.....	111
Clusters: An Overview.....	111
Client Connection.....	117
Masters and Slaves.....	118
Quorum.....	119
Election of a new Master.....	120
Message Passing.....	125
Outages and Recovery.....	126
Creating Clustered Resources.....	126
Inter-Cluster Connections.....	127
Clusters with Sites.....	127

Disaster Recovery using Clusters with Sites.....	134
Shared Storage Configurations.....	135
Setting Up a High Availability (HA) Failover Cluster.....	136
Separating Client and Cluster Communication.....	138
Multicast: An Overview.....	140
Shared Memory (SHM).....	142
Scalability.....	142
Realm Benchmarks.....	142
Failover.....	147
Connections Scalability.....	148
MQTT: An Overview.....	149
AMQP.....	153
AMQP Overview.....	154
AMQP Plugin Configuration.....	156
AMQP Security Configuration.....	158
AMQP Message Transformation.....	159
Commonly Used Features.....	161
Overview.....	162
Sessions.....	162
Channel Attributes.....	162
Storage Properties of Channels and Queues.....	167
Channel Publish Keys.....	169
Queue Attributes.....	171
Events.....	173
Named Objects, Shared Named Objects and Serial Named Objects.....	173
Event Filtering.....	178
Event Filtering using Selectors.....	179
Event Filtering using Subjects and Principals.....	182
Event Dictionaries.....	185
Durable Consumers.....	185
Durable Event Browsing and Purging.....	185
Native Communication Protocols.....	186
Comet Communication Protocols.....	189
Google Protocol Buffers.....	191
Using the Shared Memory Protocol.....	192
Rollback.....	193
Out-of-Memory Protection.....	193
Pause Publishing.....	195

Overview

This guide describes the underlying concepts of the Universal Messaging product, with the focus on the following areas of design and functionality:

- Architecture
- Management
- Performance, Scalability and Resilience
- Deployment
- Security
- MQTT (MQ Telemetry Transport)

1 Architecture

■ Architecture Overview	10
■ Support for Open Standards	11
■ Realms	11
■ Zones	12
■ Messaging Paradigms supported	17
■ umTransport API	20
■ Communication Protocols and RNAMEs	27

Architecture Overview

Universal Messaging is a Message Orientated Middleware product that guarantees message delivery across public, private and wireless infrastructures. Universal Messaging has been built from the ground up to overcome the challenges of delivering data across different networks. It provides its guaranteed messaging functionality without the use of a web server or modifications to firewall policy.

Universal Messaging design supports both *broker-based* and *umTransport* communication, and thus comprises client and server components.

Broker-Based Communication

The standard UM "broker-based" client component can be subdivided into messaging clients, comet clients and management clients. The server component has specific design features to support each of these classifications of client as well as Scheduling and Triggers, Plugins, Federation, Clustering and Low Latency IO.

Server Components

The Universal Messaging realm server is a heavily optimized Java process capable of delivering high throughput of data to large numbers of clients while ensuring latencies are kept to a minimum. In addition to supporting the client types described below the Universal Messaging realm server has a number of built in features to ensure its flexibility and performance remains at the highest levels.

Client Components

Universal Messaging supports 3 client types:

- Messaging clients
- Comet clients
- Management clients

Each client type has been developed using open protocols with specific attention paid to performance and external deployment. Each client type has been specifically designed to transparently pass through firewalls and other security infrastructure while providing its own inherent security features.

Messaging Clients

Universal Messaging messaging clients support synchronous and asynchronous middleware models. Publish Subscribe and Queues functionality are all supported and can be used independently or in combination with each other. Universal Messaging messaging clients can be developed in a wide range of languages on a wide range of platforms. Java, C# and C++ running on Windows, Solaris and Linux are all supported. Mobile devices and Web technologies such as Silverlight all exist as native messaging clients.

WebSocket, Comet and LongPolling for JavaScript Clients

In addition to our native binary wire protocol, Universal Messaging also supports text based delivery for languages that do not support binary data. Used in conjunction with Universal Messaging server plug-in technology, Comet and Long Polling clients use HTTP and persistent connections to deliver asynchronous Publish Subscribe functionality to clients. JavaScript clients can also use WebSocket as a delivery approach although this is not yet sufficiently supported in users' browsers to warrant a reliance on it over Comet/long Polling.

Management Clients

Universal Messaging provides a very extensive and sophisticated management API written in Java. Management clients can construct resources (channels, ACLs, queues etc.) and query management data (throughput, cluster state, numbers of connections etc.) directly from Universal Messaging realm servers.

umTransport Communication

Universal Messaging offers, in addition to its standard full-featured client-server API, an extremely lightweight client-client communication API known as the umTransport API (currently available for Java and C++). See the section ["umTransport API" on page 20](#) for a description of the umTransport model.

Support for Open Standards

Universal Messaging supports many open standards at different levels from network protocol support through to data definition standards.

At the network level Universal Messaging will run on an TCP/IP enabled network supporting normal TCP/IP based sockets, SSL enabled TCP/IP sockets, HTTP and HTTPS, providing multiple communications protocols (see ["Communication Protocols and RNAMEs" on page 27](#)).

Universal Messaging provides support for the JMS standard.

Realms

A Universal Messaging Realm is the name given to a single Universal Messaging server. Universal Messaging realms can support multiple network interfaces, each one supporting different Universal Messaging protocols.

A Universal Messaging Realm can contain many Channels or Message Queues.

Universal Messaging provides the ability to create clusters of realms that share resources (see ["Messaging Paradigms supported" on page 17](#)) within the namespace. Cluster objects can be created, deleted and accessed programmatically or through the Universal Messaging Enterprise Manager.

Objects created within a cluster can be accessed from any of the realms within the cluster and Universal Messaging ensures that the state of each object is maintained by all realms within a cluster. The clustering technology used within Universal Messaging ensures an unsurpassed level of reliability, resilience, scalability.

Realms can also be added to one another within the namespace. This allows the creation of a federated namespace (see ["Federation Of Servers" on page 41](#)) where realms may be in different physical location, but accessible through one physical namespace.

Zones

Overview of Zones

Zones provide a logical grouping of one or more Realms which maintain active connections to each other. Realms can be a member of zero or one zone, but a realm cannot be a member of more than one zone. Realms within the same zone will forward published channel messages to other members of the same zone, if there is necessary interest on corresponding nodes.

Note: The forwarding of messages between realms in a zone applies only to messages on channels, not to messages on queues. Messages on queues are not forwarded between realms in a zone.

Clusters can also be members of zones. In this case, the cluster as a whole is treated as a single member of a zone. Members of a zone should be well-connected, i.e., always connected to all other zone members, however this is not enforced.

Zones are uniquely identifiable, and two disparate zones with the same name can co-exist.

Interest Propagation

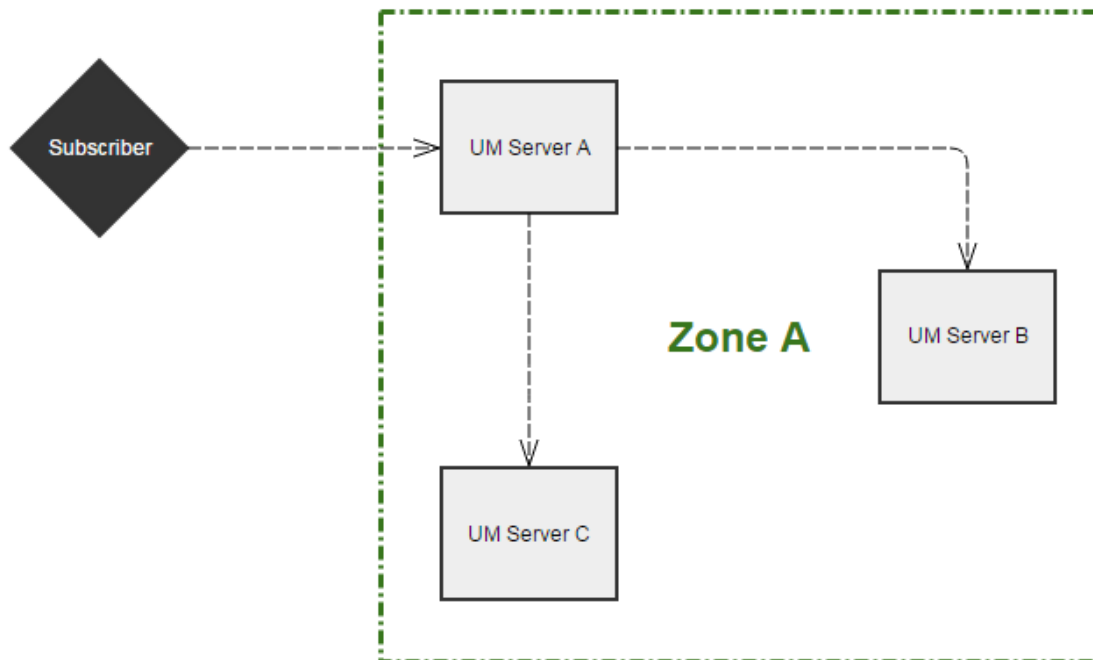
Using zones, Universal Messaging can deliver channel messages between realms and clusters when interest from one realm or cluster has been registered on another realm or cluster in the same zone. Here the term "interest" means that the realm wishes to receive a copy of messages that are sent originally to a channel of the same name on another realm within the same zone. This "interest" only arises on a realm when a client subscribes to a channel on the realm; if there is no such client on the realm, there is no "interested" client, so no interest propagation occurs.

A realm will forward its own interest in a particular channel to all realms directly connected to it within the same zone. When a realm receives a notification of interest from another realm within the same zone it will add an entry within its interest manager and begin to directly forward messages to this realm.

The mechanism within any given zone is as follows. The description assumes that there are several realms in the zone; we'll refer to the realm you are working on as the source realm and all the other realms as remote realms:

- When a client subscribes to a channel on the source realm, the source realm indicates to the remote realms that it wishes to receive any messages sent to the channel on the remote realms. In other words, the source realm propagates interest for this channel to the remote realms.
- Once interest from the source is registered on the remote realms, messages published to the channel on any of the remote realms in the zone will be propagated to the source realm.

Example: Propagating Interest within a Zone

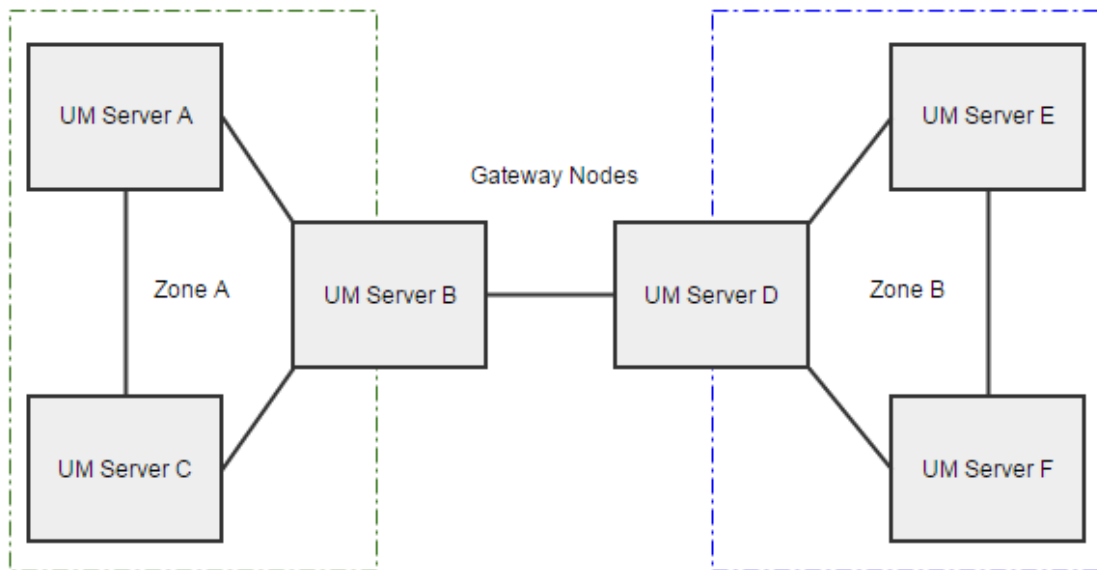


In the example above we see that a client has subscribed to a channel on realm A. The Interest Manager for realm A forwards a notification to all other realms within the same zone (realm B and realm C). Realm B and realm C now know that they should begin to forward messages they receive for this channel to realm A.

As both realm B and realm C are members of the same zone they do not propagate this interest any further. Therefore realm B knows that realm C has no local subscriber and will not forward messages to realm C. Similarly, realm C will not forward messages to realm B.

Propagating Interest between Zones

You can propagate interest between zones using by manually creating static joins, for example:



Note: Currently, a realm cannot be a member of more than one zone, so you cannot define a realm that is common to two zones. Therefore interest propagation between two zones using a common realm is not supported.

Administration of Zones

You can perform the following functions using both the Administration API and the Enterprise Manager:

- Create a new zone
- Add a realm to a zone
- Remove a realm from a zone
- List the zone which a realm currently belongs to

Administration of Zones using the Administrative API

The Administration API provides classes for performing the required administrative functions. These classes are summarized in the following outline example:

```
public class nRealmNode {
    ...
    public Zone createZone(String zoneName);
    public void joinZone(Zone zone);
    public void leaveZone(Zone zone);
    public Zone getZone();
    ...
}
```

Administration of Zones using the Enterprise Manager

The Enterprise Manager provides menu items for performing the required administrative functions.

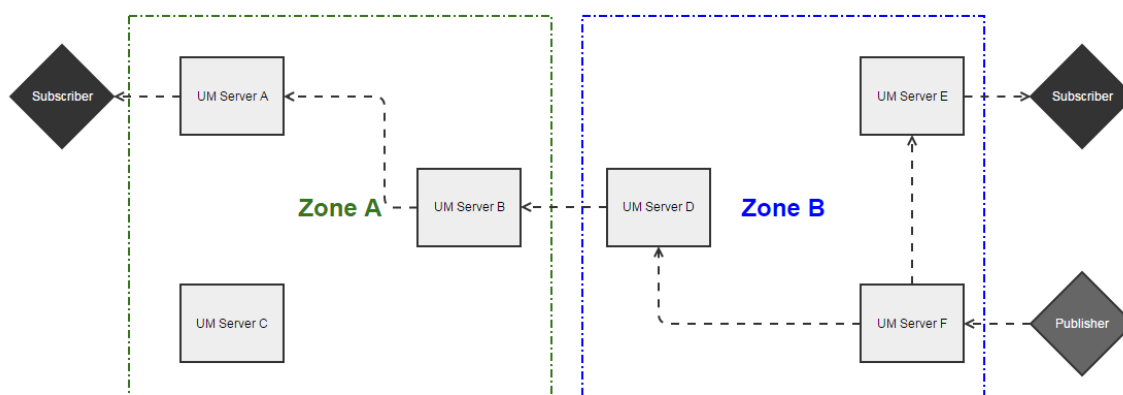
When you select a realm in the Enterprise Manager, you can perform the following functions within the context of the realm:

- Specify that the realm is a member of the given zone. If the zone does not exist already, it is created as part of this process. Currently, a realm cannot be a member of more than one zone.
- Remove the realm from a zone. If there are no more realms in the zone, the zone continues to exist, but its functionality is deactivated.
- Check which zone the realm belongs to.

See the documentation of the Enterprise Manager for details.

Examples of Message Propagation

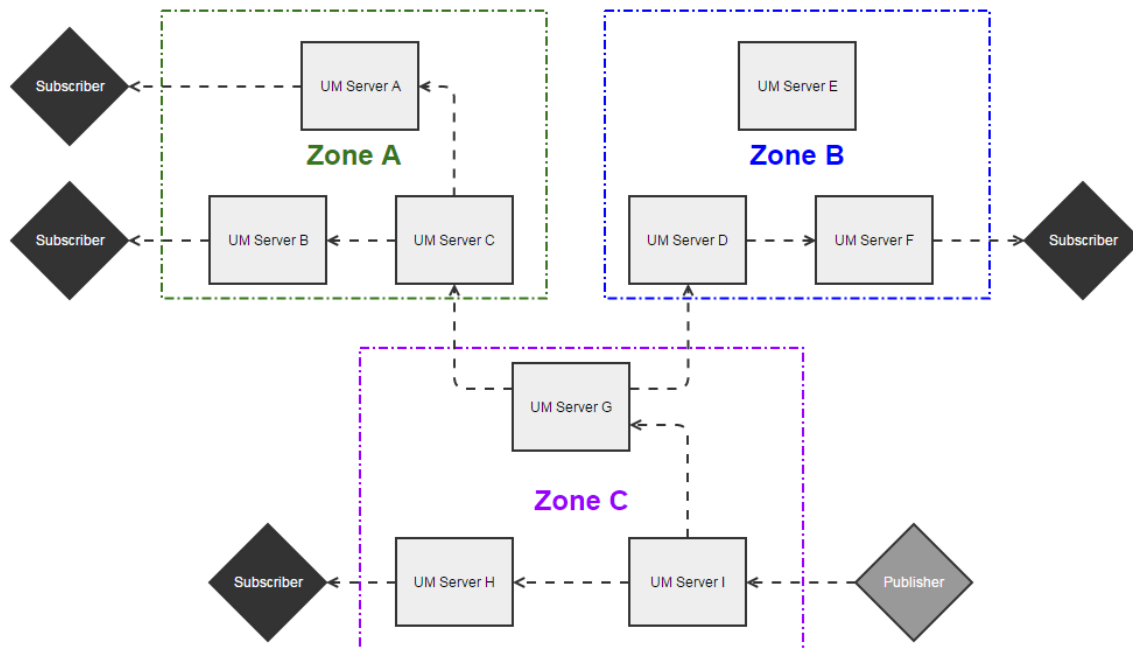
Example 1



1. A message on a given channel is published to Realm F in Zone B.
2. Servers D and E in the same zone have registered interest in the channel, so Server F forwards the message to these servers.
3. Realm E has a local subscriber, so server E delivers the message to the local subscriber. Realm E has no static joins so will not forward the message any further.
4. Realm D has a static join to Realm B. It will forward this message over the static join to Realm B.
5. Realm B receives the message through a static join and therefore forwards the message to all members of its own zone that have registered interest.

6. Realm A has registered interest as it has a local subscriber. It therefore receives this message and delivers it to the local subscriber. There are no other realms with interest so it does not forward the message further.

Example 2



1. A message is published to a channel on Realm I in Zone C.
2. Realm I forwards the message to all realms in the same zone where interest has been registered (Realm H, Realm G), as it is the initial receiver of the message.
3. Realm H forwards the received message to its subscriber. Server H has no statically configured joins so does not forward the message any further.
4. Realm G receives the message and forwards it over the static joins to Realm C in Zone A and Realm D in Zone B.
5. Realm C receives the message and forwards it to all members in Zone A that have interest.
6. Realm A and Realm B deliver the message to their subscribers.
7. Realm D receives the message and forwards it to all members of Zone B that have interest. Therefore it sends the message to Realm F. Realm E has no interest and therefore it does not forward the message to this server.
8. Realm F delivers the message to its subscriber.

Messaging Paradigms supported

Universal Messaging supports two broad messaging paradigms - Publish/Subscribe and Message Queues. Universal Messaging clients can use a mixture of these paradigms from a single session. In addition to this it is possible for clients to further control the conversation that takes place with the server by choosing particular styles of interaction. These styles are available to both readers and writers of messages and include asynchronous, synchronous, transactional and non-transactional.

Publish/Subscribe (using Channels/Topics)

Publish/Subscribe is an asynchronous messaging model where the sender (publisher) of a message and the consumer (subscriber) of a message are decoupled. When using the Channels/Topics, readers and writers of events are both connected to a common topic or channel. The publisher publishes data to the channel. The channel exists within the Universal Messaging realm server. As messages arrive on a channel, the server automatically sends them on to all consumers subscribed to the channel. Universal Messaging supports multiple publishers and subscribers on a single channel.

Publish/Subscribe (using DataStreams and DataGroups)

Universal Messaging DataGroups provide an alternative to Channels/Topics for Publish/Subscribe. Using DataGroups it is possible for remote processes to manage subscriptions on behalf of subscribers. The process of managing subscribers can be carried out by publishers themselves or by some other process. DataGroups are lightweight in nature and are designed to support large numbers of users whose subscriptions are fluid in nature. The addition / removal of users from DataGroups can be entirely transparent from the user perspective.

Message Queues


































Like pub/sub, message queues decouple the publisher or sender of data from the consumer of data. The Universal Messaging realm server manages the fan out of messages to consumers. Unlike pub/sub with channels, however, only one consumer can read a message from a queue. If more than one consumer is subscribed to a queue then the messages are distributed in a round-robin fashion.

Details of Usage for Software Developers

Details of how to develop applications in various languages for the paradigms described above, including code examples, are provided in the appropriate language-specific guides.

Messaging Paradigms Comparison

	DataGroups	Channels	Queues
Subscriptions Managed Remotely			
Nestable			
Clusterable			
Persistence (data)			
Persistence (state)			
Message Replay			
Synchronous Consumer			
Asynchronous Consumer			
Synchronous Producer			
Asynchronous Producer			
Transactional Consumer			
Non- Transactional Consumer			
Transactional Producer			


	DataGroups	Channels	Queues
Non-Transactional Producer			
Destructive Read			
Delta Delivery			
Conflation (Event Merge)			
Conflation (Event Overwrite)			
Conflation (Throttled Delivery)			
User Specified Filters			
Addressable Messages			
User Access Controlled via ACLs			
Microsoft Reactive Extensions			
Accessible via JMS			

umTransport API

Universal Messaging offers, in addition to its standard full-featured client-server API, an extremely lightweight client-client communication API known as the umTransport API.

Broker-based Model


Historically, messaging architecture has predominantly been based on a 'broker in the middle' approach. This is often referred to as 'hub and spoke'. The broker acts as the communications hub, routing messages between logically decoupled peers:



The pub-sub model is a common paradigm for broker based architecture, where one or more publishers send messages to the broker, which then distributes the messages to interested consumers.

umTransport Model

The umTransport model is a peer to peer model that allows peers to be aware of how to communicate directly with one another rather than through a broker. In effect, each publisher peer acts like a server, and each consumer can communicate directly with the publishers:



While this model bypasses broker messaging functionality such as persistence or transactional semantics, it results in a considerably lower latency delivery of information from a publisher to a consumer. By halving of the number of "hops" between client and publisher, latency too is effectively halved. This is especially useful when ultra-low latency message delivery is paramount (in, for example, the links between pricing, quant and risk engines in FX trading platforms).

The umTransport API is currently available for Java and C++. For Java, it is located in the `com.softwareag.um.modules.transport` package. For C++, it is located in `com::softwareag::umtransport`.

The Java API

The Java API is very simple, allowing each client to accept connections from other clients, and to receive arbitrary binary data from these clients synchronously or asynchronously. In many ways the API is similar to a standard TCP socket API, but offers the additional benefit of being able to use not just TCP sockets as a communication transport, but any of the following Universal Messaging communication technologies:

- TCP Sockets: data is transmitted directly over TCP Sockets
- SSL: data is SSL encrypted then transmitted over TCP Sockets
- SHM: data is transmitted via Shared Memory (for near-instant access by processes on the same machine)
- RDMA: data is transmitted via Remote Direct Memory Access (for access by processes on a remote machine; requires network adapters that support RDMA)

The C++ API

The C++ API provides a subset of the functionality available in the Java API, with the following restrictions:

- The C++ API does not support asynchronous communication between clients.
- The C++ API does not support RDMA as a communication transport.

Using the Java API

Let's take a quick look at how to use this API. Here is an example "echo" Java client and server; the EchoClient will write a string to the EchoServer; the EchoServer will respond to the EchoClient.

Here's the Java EchoClient:

```
package com.softwareag.um.modules.examples.transport.echo;
import com.softwareag.um.modules.transport.ClientContextBuilderFactory;
import com.softwareag.um.modules.transport.ClientTransportContext;
import com.softwareag.um.modules.transport.SynchronousTransport;
import com.softwareag.um.modules.transport.TransportFactory;
import com.softwareag.um.modules.examples.transport.SimpleMessage;
import com.softwareag.um.modules.examples.transport.SynchronousClient;
import java.io.BufferedReader;
import java.io.IOException;
import java.io.InputStreamReader;
/**
 * This sample app simply writes a string entered into the console to an EchoServer
 * The EchoServer will respond and this response will be output on the console.
 */
public class EchoClient {
    public EchoClient(String url) throws IOException {
        //Use the factory to generate the required builder based on the protocol
        //in the url string
        ClientTransportContext context = ClientContextBuilderFactory.getBuilder(url).build();
        //We do not pass any handlers to the connect method
        //because we want a synchronous transport
        SynchronousTransport transport = TransportFactory.connect(context);
        //This is just a basic wrapper for the client transport
        //so it is easier to read/write messages
        SynchronousClient<SimpleMessage> client =
            new SynchronousClient<SimpleMessage>(transport);
        BufferedReader br = new BufferedReader(new InputStreamReader(System.in));
        //Start a new thread to read from the client transport
        //because read is a blocking call
        new ReadThread(client);
        //Now continue to write messages to the EchoServer until the user enter 'quit'
        while(true){
            System.out.println("Enter a message or type 'quit' to exit >");
            String line = br.readLine();
            if(line.equalsIgnoreCase("quit")){
                break;
            }
            else{
                client.write(new SimpleMessage(line));
            }
        }
    }
    private static class ReadThread extends Thread{
        SynchronousClient<SimpleMessage> client;
        public ReadThread(SynchronousClient<SimpleMessage> _client){
            client = _client;
            start();
        }
        @Override
        public void run(){
            try{
                while(true){
                    SimpleMessage mess = client.read(new SimpleMessage());
                    System.out.println(mess.toString());
                }
            }
        }
    }
}
```

```

    }
    }
    catch(Exception e){
        e.printStackTrace();
    }
}
}
public static void main(String[] args) throws IOException {
    if(args.length == 0){
        usage();
        System.exit(1);
    }
    new EchoClient(args[0]);
}
public static void usage(){
    System.out.println("EchoClient <URL>");
    System.out.println("<Required parameters>");
    System.out.println("\tURL - protocol://host:port for the server to connect to e.g. "
        +TransportFactory.SOCKET+"://localhost:9000");
}
}
}

```

And, the EchoServer:

```

package com.softwareag.um.modules.examples.transport.echo;
import com.softwareag.um.modules.transport.ServerContextBuilderFactory;
import com.softwareag.um.modules.transport.ServerTransportContext;
import com.softwareag.um.modules.transport.SynchronousServerTransport;
import com.softwareag.um.modules.transport.SynchronousTransport;
import com.softwareag.um.modules.transport.TransportFactory;
import com.softwareag.um.modules.examples.transport.SimpleMessage;
import com.softwareag.um.modules.examples.transport.SynchronousClient;
import java.io.IOException;
/**
 * This sample will only handle one client connection at a time. When a client connects,
 * the EchoServer will immediately respond to any messages with exactly the same message.
 */
public class EchoServer implements Runnable{
    private volatile SynchronousClient<SimpleMessage> client;
    private final SynchronousServerTransport transport;
    private volatile boolean stopped = false;
    public EchoServer(String url) throws IOException {
        //The factory will create the correct context based on the protocol in the url
        ServerTransportContext context = ServerContextBuilderFactory.getBuilder(url).build();
        //Because we have not passed an AcceptHandler into the bind method, we are returned
        //a SynchronousServerTransport. This means we have to call accept on the transport
        //to accept new client transports.
        transport = TransportFactory.bind(context);
    }
    public static void main(String[] args) throws IOException {
        if(args.length == 0){
            usage();
            System.exit(1);
        }
        EchoServer echoServer = new EchoServer(args[0]);
        Thread t = new Thread(echoServer);
        t.start();
        System.out.println("Press enter to quit.");
        System.in.read();
        echoServer.close();
    }
    public static void usage(){
        System.out.println("EchoServer <URL>");
        System.out.println("<Required parameters>");
    }
}

```

```

        System.out.println(
            "\tURL - protocol://host:port to bind the server transport to e.g. "
            +TransportFactory.SOCKET+"://localhost:9000");
    }
    protected void close(){
        stopped = true;
        client.close();
        transport.close();
    }
    public void run() {
        try{
            while(true){
                System.out.println("Waiting for client");
                //accept() will block until a client makes a connection to our server
                SynchronousTransport clientTransport = transport.accept();
                System.out.println("Client connected. Echo service started.");
                //The SynchronousClient is simply a wrapper to make reading/writing easier
                client = new SynchronousClient<SimpleMessage>(clientTransport);
                try{
                    while(!stopped){
                        client.write(client.read(new SimpleMessage()));
                    }
                }
                catch (IOException e){
                    System.out.println("Connection closed");
                }
            }
        }
        catch(IOException e){
            e.printStackTrace();
        }
    }
}

```

Using the C++ API

Here's the C++ EchoClient:

```

#include "EchoClient.h"
#include "ClientTransportContextFactory.h"
#include "TransportFactory.h"
#include <utility>
#include <iostream>
com::softwareag::umtransport::samples::echo::EchoClient::EchoClient(std::string url)
{
    m_url = url;
}
com::softwareag::umtransport::samples::echo::EchoClient::~EchoClient()
{
}
void com::softwareag::umtransport::samples::echo::EchoClient::run() {
    try{
        //Use the factory to generate the required builder based on the protocol
        //in the url string
        auto context = ClientTransportContextFactory::build(m_url);
        //We do not pass any handlers to the connect method
        //because we want a synchronous transport
        auto transport = TransportFactory::connect(std::move(context));
        //This is just a basic wrapper for the client transport
        //so it is easier to read/write messages
        std::shared_ptr<SynchronousTransport> transportShared(std::move(transport));
        SynchronousClient<SimpleMessage> client(std::move(transportShared));
        //Start a new thread to read from the client transport because read is a
    }
}

```



```

        //blocking call
        ReadThread readThread(client);
        Poco::Thread th;
        th.start(readThread);
        bool canRun = true;
        //Now continue to write messages to the EchoServer until the user enter 'quit'
        while (canRun){
            std::cout << "Enter a message or type 'quit' to exit >" << std::endl;
            std::string input;
            std::getline(std::cin, input);
            if (input == "quit"){
                canRun = false;
            }
            else{
                SimpleMessage sm(input);
                client.write(sm);
            }
        }
        readThread.shutdown();
        client.close();
        th.tryJoin(10000);
    }
    catch (Poco::Exception &ex){
        std::cout << ex.displayText();
    }
}

int com::softwareag::umtransport::samples::echo::EchoClient::main(int argc,
    char** argv){
    if (argc < 2){
        std::cout <<
            "EchoClient <URL>\n<Required parameters>\n\tURL - " <<
            "protocol://host:port for the server to connect to e.g. tcp://localhost:9000" <<
            std::endl;
    }
    EchoClient client(argv[1]);
    client.run();
    return 0;
}

com::softwareag::umtransport::samples::echo::EchoClient::ReadThread::ReadThread(
    SynchronousClient<SimpleMessage> &client) : m_client(client){
}

com::softwareag::umtransport::samples::echo::EchoClient::ReadThread::~~ReadThread(){
}

void com::softwareag::umtransport::samples::echo::EchoClient::ReadThread::shutdown(){
    canRun = false;
}

void com::softwareag::umtransport::samples::echo::EchoClient::ReadThread::run(){
    try {
        while (canRun){
            SimpleMessage message;
            m_client.read(message);
            std::cout << "Message Content: " << message << std::endl;
        }
    }
    catch (Poco::Exception& e) {
        std::cout << "Connection Closed " << e.displayText();
    }
}
}

```

And here's the EchoServer:

```

#include "EchoServer.h"
#include "ServerTransportContextFactory.h"
#include "TransportFactory.h"

```

```

#include "SynchronousClient.h"
#include "SimpleMessage.h"
using namespace com::softwareag::umtransport;
com::softwareag::umtransport::samples::echo::EchoServer::EchoServer(std::string url) :
stopped(false){
    //The factory will create the correct context based on the protocol in the url
    //Because we have not passed an AcceptHandler into the bind method, we are returned
    //a SynchronousServerTransport. This means we have to call accept on the transport
    //to accept new client transports.
    m_transport = TransportFactory::bind(ServerTransportContextFactory::build(url));
}
com::softwareag::umtransport::samples::echo::EchoServer::~EchoServer(){
}
int com::softwareag::umtransport::samples::echo::EchoServer::main(int argc,
    char** argv){
    if (argc < 2){
        std::cout << "EchoServer <URL>" << std::endl << "EchoServer <URL>" << std::endl <<
            "\tURL - protocol://host:port to bind the server transport
            to e.g. tcp://localhost:9000" << std::endl;
        exit(1);
    }
    EchoServer echoServer(argv[1]);
    Poco::Thread th;
    th.start(echoServer);
    std::cout << "Press any key to finish" << std::endl;
    std::cin.ignore();
    echoServer.close();
    th.tryJoin(10000);
    return 0;
}
void com::softwareag::umtransport::samples::echo::EchoServer::close(){
    stopped = true;
    m_transport->close();
}
void com::softwareag::umtransport::samples::echo::EchoServer::run(){
    try{
        while (!stopped){
            std::cout << "Waiting for a client" << std::endl;
            //accept() will block until a client makes a connection to our server
            std::shared_ptr<SynchronousTransport> syncTrShared(
                std::move(m_transport->accept()));
            SynchronousClient<SimpleMessage> client(syncTrShared);
            //Client connected echo service started
            try{
                while (!stopped){
                    SimpleMessage msg;
                    client.read(msg);
                    client.write(msg);
                }
            }
            catch (Poco::Exception & ex){
                std::cout << "Connection Closed" << std::endl;
            }
        }
    }
    catch (Poco::Exception & ex){
        std::cout << ex.displayText() << std::endl;
    }
}

```

Communication Protocols and RNAMEs

Universal Messaging supports several Native Communication Protocols (see "[Native Communication Protocols](#)" on page 186) and Comet Communication Protocols (see "[Comet Communication Protocols](#)" on page 189).

The following table shows the Communication Protocols supported by each Universal Messaging Client API:

	Native Communication Protocols					Comet Communication Protocols	
	Socket (nsp)	SSL (nsp)	HTTP (nhp)	HTTPS (nhps)	Shared Memory (shm)	HTTPS (https)	HTTP (http)
Java	✓	✓	✓	✓	✓	○	○
C# .NET	✓	✓	✓	✓	○	○	○
C++	✓	✓	✓	✓	✓	○	○
Python	✓	✓	✓	✓	○	○	○
Excel VBA	✓	✓	✓	✓	○	○	○
JavaScript	○	○	✓ via WebSocket	✓ via WebSocket	○	✓	✓
Microsoft Silverlight	✓	○	✓	✓	○	○	○
iPhone	✓	✓	✓	✓	○	○	○
Android	✓	✓	✓	✓	○	○	○

RNAMEs

An RNAME is used by Universal Messaging Clients to specify how a connection should be made to a Universal Messaging Realm Server.

A Native Communication Protocol (see ["Native Communication Protocols" on page 186](#)) RNAME string looks like:

```
<protocol> :// <hostname> :< port>,<protocol> :// <hostname> :< port>
```

or

```
<protocol> :// <hostname> :< port>;<protocol> :// <hostname> :< port>
```

where:

- `<protocol>` can be one of the 4 available native communications protocol identifiers `nsp` (socket), `nhp` (HTTP), `nsps` (SSL) and `nhps` (HTTPS).
- `<hostname>` is the hostname or IP address that the Universal Messaging Realm is running on.
- `<port>` is the TCP port on that hostname that the Universal Messaging Realm is bound to using the same wire protocol.

The RNAME entry can contain a comma or semicolon separated list of values, each one representing the communications protocol, host and port currently running on a Universal Messaging Realm.

Using a comma-separated list indicates the traversal of the RNAME values in the order given, while using a semicolon-separated list indicates the random traversal of the RNAME values.

If a list of RNAMEs is used and the Universal Messaging session becomes disconnected and cannot reconnect, the API will traverse through the RNAME list until it manages to reconnect. This functionality is particularly useful within the contexts of both clustering (see ["Clusters: An Overview" on page 111](#)) and failover (see ["Failover" on page 147](#)).

Native Communications Protocol Client Extensions

In addition to the supported protocols shown above, Universal Messaging clients implemented with APIs that support Native Communication Protocols have the following extension available to them:

- *nhpsc*

This mode of https extracts any configured proxy from within the JVM settings and issues a PROXY CONNECT command via this proxy to establish a connection with the required Universal Messaging realm. The established connection then becomes an SSL encrypted socket connection mode and no longer uses http/https connections for each server request. If the proxy uses authentication then authentication parameters are also extracted from the JVM settings.

2 Management

■ Administration and Management	30
■ JMX Administration and Management	31

Administration and Management

In addition to its communications APIs and features Universal Messaging provides a sophisticated collection of management tools and APIs. These tools and APIs are designed exclusively for:

- ["Collection of Statistical Data from Universal Messaging" on page 30](#)
- ["Monitoring of Events" on page 30](#)
- ["Creation of Universal Messaging Resources, ACLs and Clusters" on page 30](#)
- ["Management of Configuration Parameters" on page 30](#)
- ["Seamless Integration with Third Party Enterprise Systems Management Tools" on page 31](#)

Universal Messaging's management client, the Enterprise Manager, is written using the same management APIs thus demonstrating the powerful features of these features.

Statistical Data

Through the use of the Universal Messaging management API clients can access a very detailed range of performance related data. Performance metrics can be gathered at many levels ranging from the realm throughput statistics to individual client connection round trip latency details. See the code example "Monitor the Remote Realm Log and Audit File" for an illustration of usage.

Management Event Monitoring

Most client and server induced actions in Universal Messaging result in a management event being created. Asynchronous listeners can be created using the management API that enables management clients to capture these events. As an example consider a client connection to a Universal Messaging realm server. This creates a client connection event. A management client at this point might dynamically create channel resources for said client and programmatically set ACLs. See the code example "Monitor Realms for Client Connections Coming and Going" for an illustration of usage.

Resource Creation

Resources (see ["Messaging Paradigms supported" on page 17](#)) can all be created programmatically using the Universal Messaging Administration API. Coupled with statistical data and event monitoring resources can be created on the fly to support users in specific operational configurations. For example, create channel x when user x logs in OR change channel ACL when realm throughput exceeds a specific value. An ACL creation example can be found in the add queue acl sample.

Configuration Management

Every Universal Messaging Server has a number of configurable parameters. In addition specific interfaces supporting specific protocols and plugins can be added to Universal


Messaging realms. Universal Messaging's configuration management feature allow clients to snapshot configurations and generate configuration XML files. New realms can be very quickly configured with the XML files enabling the very fast bootstrapping of new environments. The Enterprise Manager documentation has an XML sample.

3rd Party Integration

While Universal Messaging's Administration API can be using directly to integrate with 3rd party products Universal Messaging servers also support JMX (see "[JMX Administration and Management](#)" on page 31) and (security permitting) can be queried by any JMX management tool.

JMX Administration and Management

In addition to Universal Messaging's Administration API, a series of JMX beans enable monitoring and management via any JMX container. The following image illustrates a JConsole view running on Mac OSX:



In order to connect to Universal Messaging over the network using JConsole or another tool that supports JMX, you must first enable the realm configuration property EnableJMX. For information on setting this property, see the section *Realm Configuration* of the Enterprise Manager documentation.

Additionally, the following JVM system properties need to be added to the nserver.conf or nserverdaemon.conf file:

```
wrapper.java.additional.991=-Dcom.sun.management.jmxremote
wrapper.java.additional.992=-Dcom.sun.management.jmxremote.port=9010
wrapper.java.additional.994=-Dcom.sun.management.jmxremote.authenticate=false
wrapper.java.additional.995=-Dcom.sun.management.jmxremote.ssl=false
```

Note: nserver.conf is used if you start the realm server manually, e.g. from the command line. nserverdaemon.conf is used if you start the realm server as a Windows service or as a UNIX daemon.

Note:

- The port can be freely chosen, as long as it is not already occupied (otherwise Universal Messaging will not start).
- If you start JMX by passing JVM system properties that are different from the Universal Messaging configuration shown above, Universal Messaging will not be aware of this change and the Enterprise Manager will not show the proper URL.

Start JMC (Java Mission Control), or whatever tool is your favorite, and connect with the following JMX server URL:

```
service:jmx:rmi:///jndi/rmi://localhost:9010/jmxrmi
```

If you wish to enable authentication or SSL, please consult the JMX documentation or contact our support team.

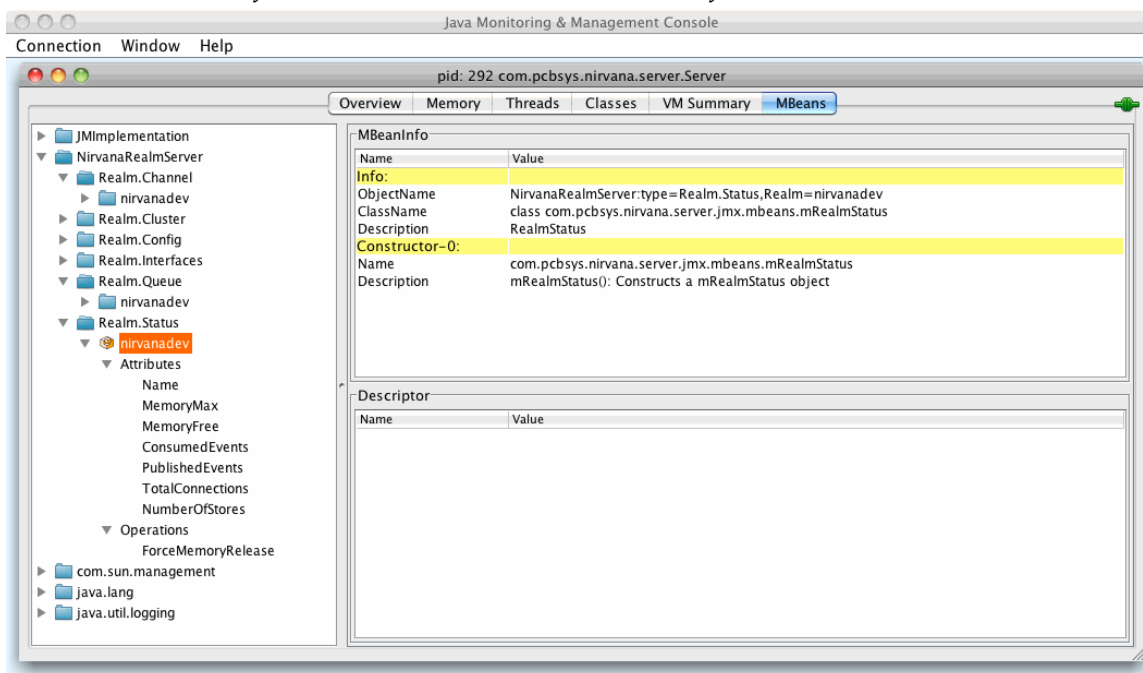
Universal Messaging offers the following JMX beans:

- [" Universal Messaging Realm Status JMX bean" on page 32](#)
- [" Universal Messaging Realm Cluster JMX bean" on page 33](#)
- [" Universal Messaging Realm Configuration JMX bean" on page 34](#)
- [" Universal Messaging Realm Interfaces JMX bean" on page 34](#)
- [" Universal Messaging Channel JMX bean" on page 35](#)
- [" Universal Messaging Queue JMX bean" on page 36](#)

Universal Messaging Realm Status JMX bean

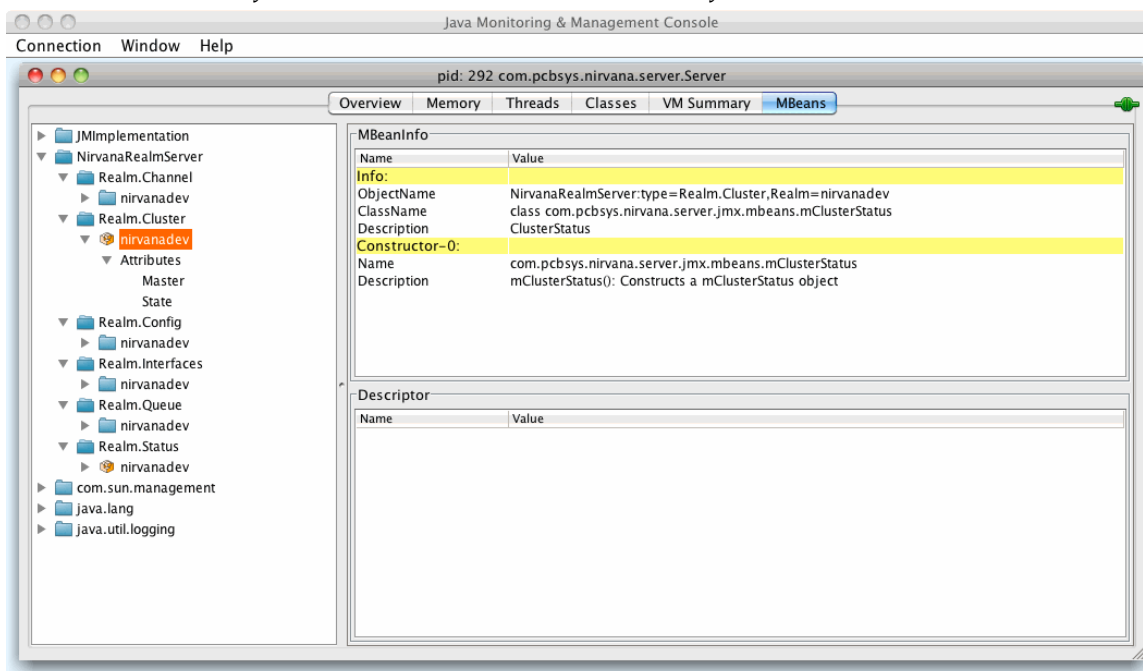
The Universal Messaging Realm Status JMX bean enables access to data visible in the Enterprise Manager Realm Status view. The following image

illustrates how the JMX bean looks when accessed via JConsole on Mac OSX:



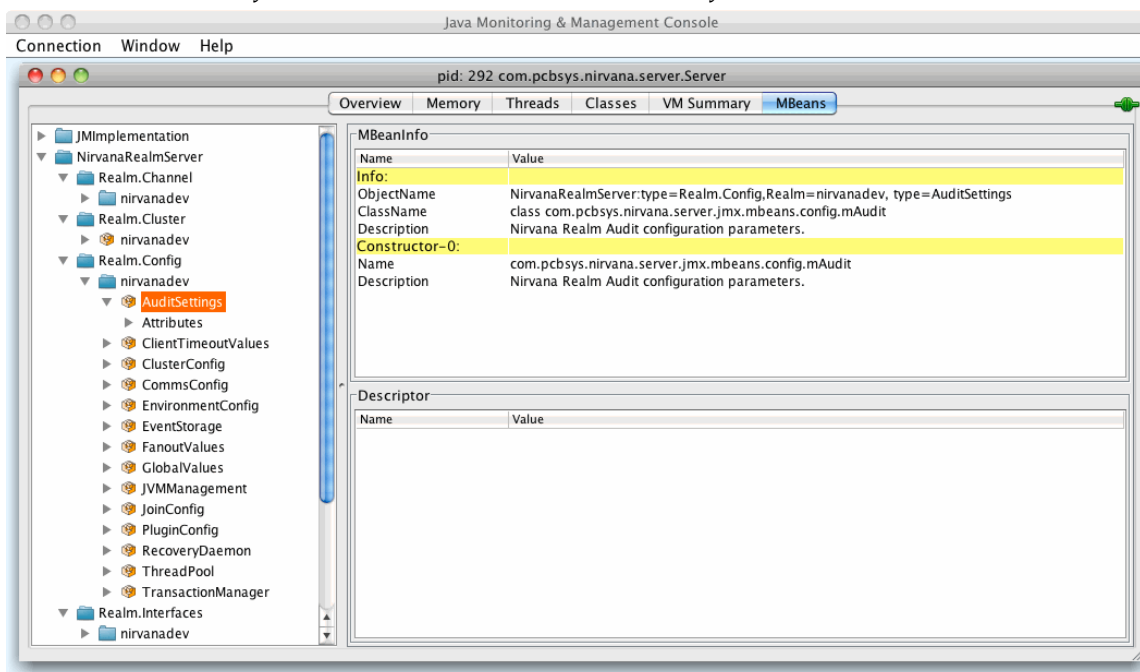
Universal Messaging Realm Cluster JMX bean

The Universal Messaging Realm Cluster JMX bean enables access to data visible in the Enterprise Manager Cluster Status view. The following image illustrates how the JMX bean looks when accessed via JConsole on Mac OSX:



Universal Messaging Realm Configuration JMX bean

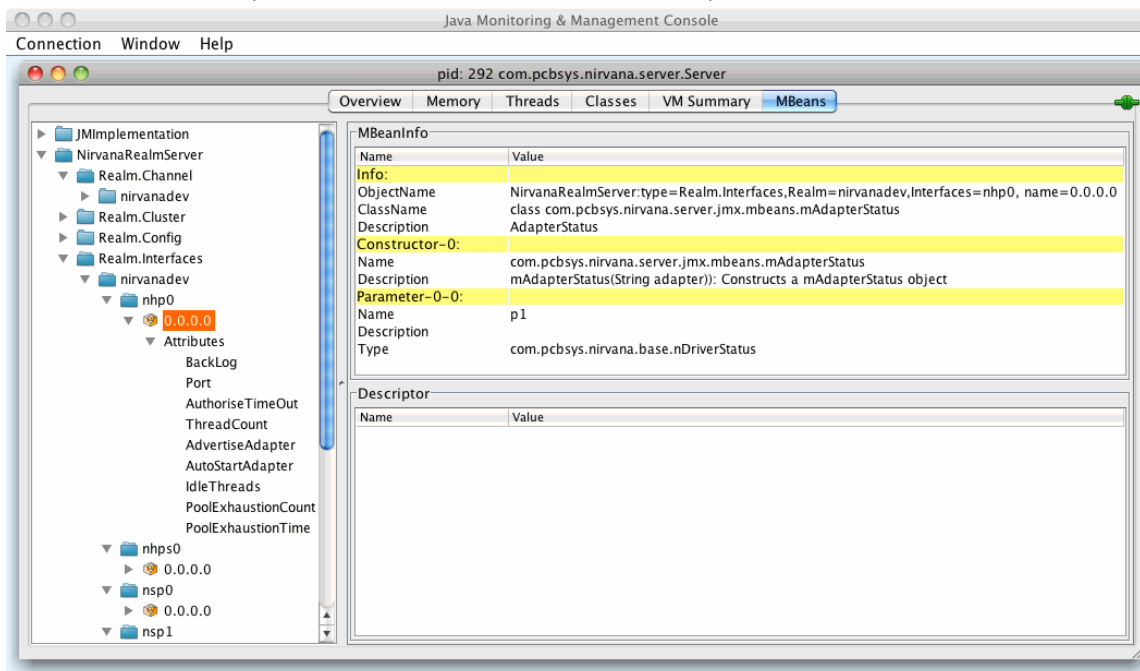
The Universal Messaging Realm Configuration JMX bean enables access to data visible in the Enterprise Manager Realm Configuration view. The following image illustrates how the JMX bean looks when accessed via JConsole on Mac OSX:



Universal Messaging Realm Interfaces JMX bean

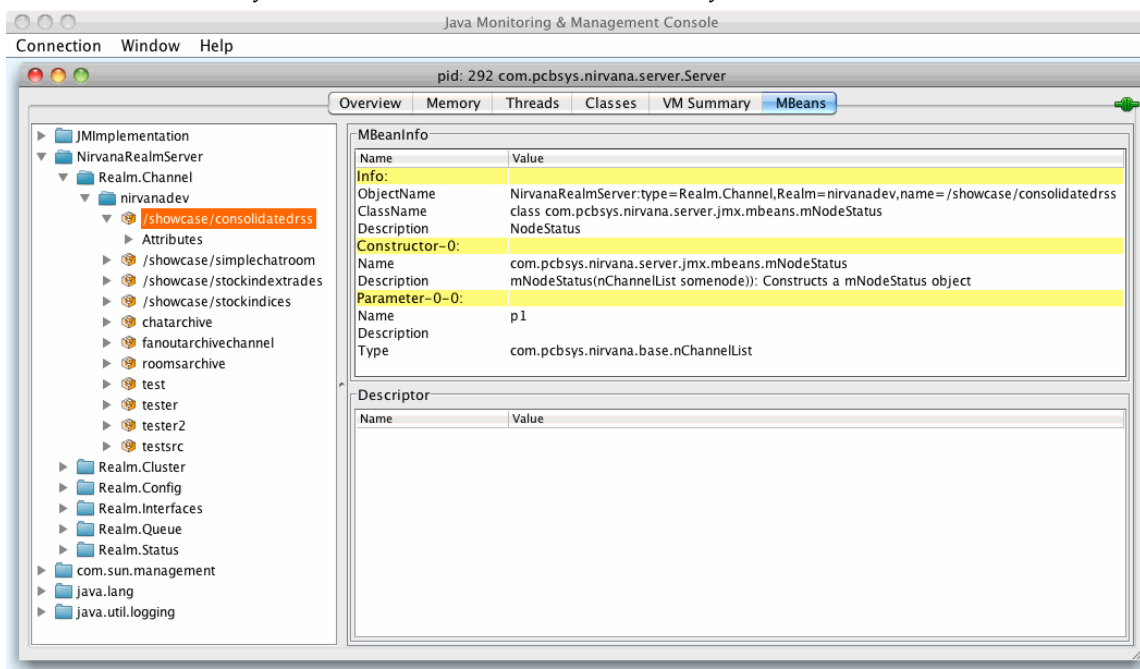
The Universal Messaging Realm Interfaces JMX bean enables access to data visible in the Enterprise Manager Realm Interfaces view. The following image

illustrates how the JMX bean looks when accessed via JConsole on Mac OSX:



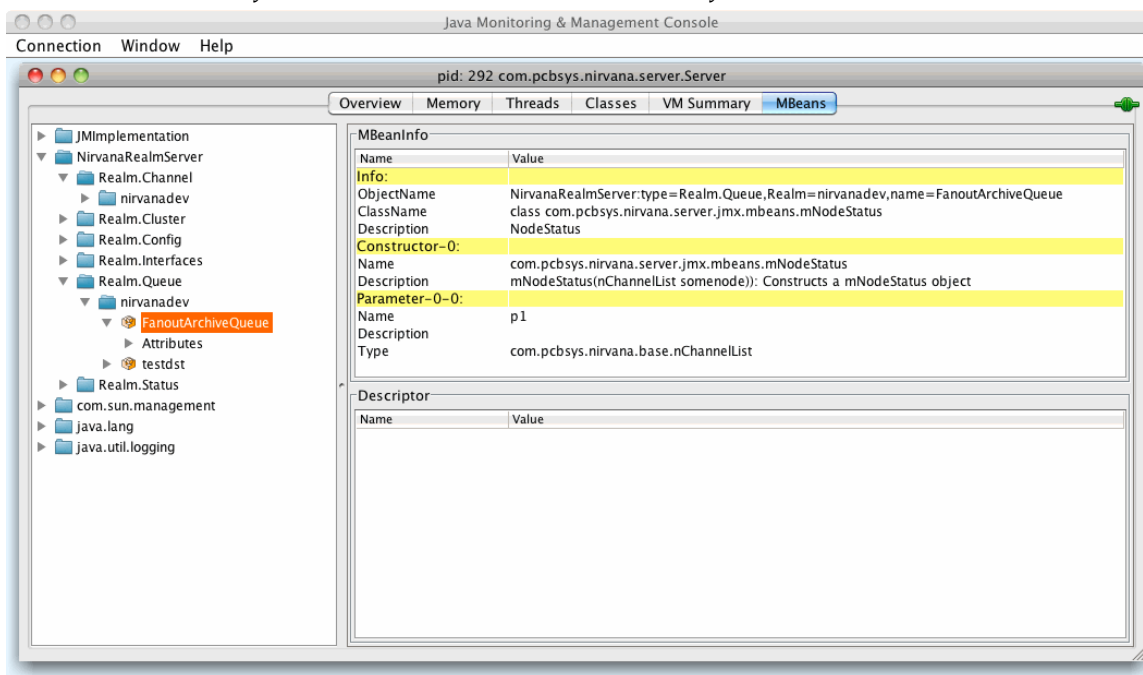
Universal Messaging Channel JMX bean

The Universal Messaging Channel JMX bean enables access to data visible in the Enterprise Manager Channel Statu view. The following image illustrates how the JMX bean looks when accessed via JConsole on Mac OSX:



Universal Messaging Queue JMX bean

The Universal Messaging Queue JMX bean enables access to data visible in the Enterprise Manager Queue Status view. The following image illustrates how the JMX bean looks when accessed via JConsole on Mac OSX:



3 Deployment

■ Deployment	38
■ Server	39
■ Client	63
■ Language Deployment Tips	69

Deployment

The structure and target audience for any Enterprise Application determines the deployment topology for the underlying infrastructure that supports it. Universal Messaging provides a wide degree of flexibility when faced with stringent deployment requirements. Key features are discussed below

Persistence and Configuration

Universal Messaging uses its own persistent stores that remain relative to its installation location on your file system. Multiple realms can be configured from a single installation, each with their own configuration files and persistent stores for event storage.

Configuration Snapshots

All aspects of a Universal Messaging realms configuration can be stored in an XML file. Channels, ACL's, Interface configuration, Plugins etc. can all be included. New realms can quickly be bootstrapped from existing configuration files making the deployment to new environments quick and simple.

Interfaces

Universal Messaging Realms can support multiple communications interfaces, each one defined by a protocol a port. Universal Messaging Realms can be configured to bind to all network interfaces on a machine or specific ones depending on configuration requirements. Specific SSL certificate chains can be bound to specific interfaces thus insuring clients always authenticate and connect to specific interfaces

Web Deployment

As well as providing a wide range of client web technology support Universal Messaging's realm server provides useful features to aid web deployment. In addition to providing a messaging backbone for external clients Universal Messaging can also act as a web server delivering static and server-generated content to clients. This resolves security sandbox problems and port use.

Forward and reverse proxy server functionality is available for those that wish to host web content on a different server but stream real time data from Universal Messaging.

Server

Performance and Tuning

This section will provide some initial information and guidance on how to get the best out of Universal Messaging, as well as provide an explanation to understand the significance of certain steps when tuning applications for low latency.

General Advice

Much of the information given in this section of the Website is related to tuning a specific element of your system. In the list below are more general pieces of advice which may help improve performance. Links to more in depth tuning articles can be found further down the page.

- Ensure you are running the latest Universal Messaging release. We strive to enhance performance between releases, upgrading will ensure you are able to leverage the newest improvements
- Use the latest version of the Java Virtual Machine. JVM vendors often improve the performance of the virtual machine, or its garbage collector between releases.
- Collect monitoring information which will allow you to make informed decisions based on the origin of performance bottlenecks. Operating System provide statistics on memory consumption, processor and network utilization. Java Virtual Machines can output Garbage Collection statistics which can be a key part of diagnosing why an application may not be performing.

Detailed Topics

The links below provide additional information on tuning important aspects of a deployed application.

- ["Realm Configuration" on page 59](#)
- ["JVM Tuning" on page 52](#)
- ["OS Tuning" on page 58](#)
- ["Network Tuning" on page 56](#)

Validating Results

Much of the advice given here is based on our own observation by running our internal benchmarking suite. Your environment and needs may differ from those we model, so we would encourage that you validate any changes you make to your environment.

Many parameters, usually kernel parameters, are specific to an individual machine. Furthermore, it can be dangerous to change them without proper knowledge. It is encouraged to exercise caution when changing such settings.

Server Failover / High Availability

In order to provide your clients with a service that is highly available, clustering is recommended. Clusters enable transparency across your clients. If one server becomes unavailable, the client will automatically reconnect to another realm within the cluster. All cluster objects within the realm are replicated among all cluster realms and their state is maintained exactly the same across all realm members. Therefore whenever a client disconnects from one realm and reconnects to another, they will resume from the same position on the newly connected realm.

When a client provides a list of RNAMEs as a comma separated list, if each entry in the list corresponds to realm that is a member of the cluster, then the client will reconnect to the next realm in the cluster list.

For more information on clustering, please see the clustering section in the Administrators Guide.

Data Routing

Joining a channel to another channel or queue allows you to set up content routing so that events published to the source channel will be passed on to the destination channel/queue automatically. Joins also support the use of filters, thus enabling dynamic content routing.

Please note that while channels can be joined to both channels and queues, queues cannot be used as the source of a join.

Channels can be joined using the Universal Messaging Enterprise Manager GUI or programmatically.

When creating a join there is one compulsory option and two optional ones. The compulsory option is the destination channel. The optional parameters are the maximum join hops and a filter to be applied to the join.

Note: For a description of the general principles involved in creating channel joins, see the section *Creating Channel Joins* in the *Administration Guide*. The description details the usage based on the Enterprise Manager, but the same general principles apply if you are using the API.

Hop Count

Joins have an associated hop-count, which can optionally be defined when the join is created. The hop count allows a limit to be put on the number of subsequent joins an event can pass through if published over this join. If a hop count is not defined for a join, it will default to 10.

The hop count is the number of *intermediate* stores between the source channel and the final destination. As an example, imagine we have 10 channels named "channel0" to "channel9" and all these channels are joined sequentially. When we publish to channel

0, if the join from channel0 to channel1 has a hop count of 5 then the event will be found on channel0 (the source channel), channels 1 to 5 (the intermediate channels) and channel6 (the endpoint).

Loop Detection

Joins allow the possibility of defining a loop of joined channels. To prevent channels receiving multiple copies of the same event, Universal Messaging implements loop detection on incoming events. To illustrate this, imagine a simple example with two channels (channel0 and channel1) and we create a loop by joining channel0 to channel1 and channel1 to channel0. If we publish to channel0 the event will also be published to channel1 over the join. But channel1 is joined to channel0 too, so now the event would get published to channel0 again. Without Universal Messaging's loop detection, this cycle would repeat until the maximum hopcount has been reached.

To prevent this, Universal Messaging detects when a message which has already been published to a channel or queue and will not publish it a second time.

Multiple Path Delivery

Universal Messaging users can define multiple paths over different network protocols between the same places in Universal Messaging. Universal Messaging guarantees that the data always gets delivered once and once only.

Federation Of Servers

Universal Messaging supports the concept of a federated namespace, where realm servers may be located in different geographical locations but form part of the same logical namespace. A Universal Messaging name space can contain one or more Universal Messaging message servers, each one containing many topics or queues.

Each Universal Messaging server is aware of others that have been added to the namespace and each one can redirect clients automatically to the required resource thus providing alternative routes when network outages occur. There is no single point of entry to a federated Universal Messaging namespace and it can be traversed in any direction from any point.

The entry into a Universal Messaging name space or server is via a custom URL called an RNAME (see "[Communication Protocols and RNAMEs](#)" on page 27). The RNAME provides the protocol, host and port required to access the Universal Messaging server. Universal Messaging clients can be passed an array of RNAME's. Should a connection fail to one of the realms the Universal Messaging client automatically moves onto the next.

The remote management of either clustered or federated realm servers is enabled via the Universal Messaging Enterprise Manager or administration API. There is no limit placed on the number of Universal Messaging Realms that can be managed from the Universal Messaging Enterprise Manager or using the Universal Messaging Administration API.

Proxy Servers and Firewalls

Universal Messaging transparently traverses modern proxy servers and firewall technology. Universal Messaging's HTTP and HTTPS drivers (see ["Using HTTP/HTTPS" on page 63](#)) support straight proxy servers as well as user authenticated proxy servers.

Universal Messaging self contained HTTP/HTTPS implementation ensures that if a remote client can access your web site the same client can access a Universal Messaging realm.

Server Memory for Deployment

Universal Messaging servers provide 3 different memory modes. Typically, the Universal Messaging Realm server will be deployed using the large memory mode. When deploying a Universal Messaging server, one of the considerations for memory consumption concerns the volatility of your data, and specifically the types of channels and queues you are using.

The channels that consume the most memory are those channels that keep the events in memory and do not write events to persistent store. These channels are known as *Simple* and *Reliable*.

If you have a simple channel with a TTL of say 1 day (86400000 milliseconds), and you expect to publish a 1k event per second, this channel alone will consume approximately 86.4MB of memory. However if your data has a very short lifespan defined by setting a low TTL, then the memory consumption would be much less than it would be with a 1 day TTL.

This kind of calculation will indicate to you how much maximum memory the Realm Server JVM needs to be allocated to avoid running out of memory.

If you follow these simple guidelines, you should be able to estimate the memory required for your channels.

Server Parameters

Introduction

When a Universal Messaging Realm Server is started, there are a number of parameters used in its initial startup sequence.

If you start the realm server manually, e.g. from a shortcut or from the command line, these parameters are in the form of -D options specified in the nserver.conf configuration file. This file is located in the server/<realm>/bin directory of your installation, where <realm> is the name of the Universal Messaging realm. If you start the realm server as a Windows service or as a UNIX daemon, the parameters are taken from nserverdaemon.conf instead.

The -D options are specified in the configuration file in the following format:

```
wrapper.java.additional.<n>=-D<parameter>=<value>
```

where <parameter> is the name of the parameter whose value you wish to set, <value> is the value that you want to assign to the parameter, and <n> is any positive integer. If you have several -D options in the file, ensure that <n> values are unique. If you supply several -D options with the same <n> value, only the last of the specified options with this value of <n> will be effective and the earlier ones with this value of <n> will be ignored.

Example:

```
wrapper.java.additional.25=-DMaxFileSize=4000000000
wrapper.java.additional.26=-DLOGLEVEL=3
```

The Universal Messaging client API also supports several parameters that can be specified in the command line of any Universal Messaging Client application.

Description of the Server Parameters

This section describes the parameters, what they are used for and their typical values.

Name	Required	Default	Description
ADAPTER	N		Specifies an interface to use, for example nsp://0.0.0.0:9000/
ADAPTER_x	N		Specifies an interface to use, for example nsp://0.0.0.0:9000/ where x = 0 -> 9
CAKEystore	N		Shorthand for javax.net.ssl.trustStore
CAKEystorePASSWD	N		Shorthand for javax.net.ssl.trustStorePassword
CHANNELUMASK	N		Specifies the default channel protection mask
CKEystore	N		Shorthand for javax.net.ssl.keyStore
CKEystorePASSWD	N		Shorthand for javax.net.ssl.keyStorePassword

Name	Required	Default	Description
DATADIR	Y		The path that Universal Messaging uses to store its internal status and configuration information about the realm. See the section "Notes on DATADIR" on page 46 below for related information.
javax.net.ssl.debug	N		Useful to debug SSL issues, see www.javasoft.com for more information
javax.net.ssl.keyStore	N		Used to set the default KeyStore the server will use. If not supplied the client MUST set one when configuring an SSL interface
javax.net.ssl.keyStorePassword	N		Used to set the default password for the keystore. If not supplied the client must set one when configuring an SSL interface
javax.net.ssl.trustStore	N		Used to set the default trust store the server will use. If not supplied the client MUST set one when configuring an SSL interface
javax.net.ssl.trustStorePassword	N		Used to set the default Truststore password the server will use. If not supplied the client MUST set one when configuring an SSL interface

Name	Required	Default	Description
LOG_FRAMEWORK	N		Specifies a third part logging framework to use. Can be set to "LOGBACK" or "LOG4J2". The default is "fLogger".
LOGLEVEL	N	5	Specifies the current log level to use
LOGFILE	N	System.out	Used to specify a log file to write the log entries to
LOGSIZE	N	100000	Specified in bytes before the log file is rolled
MaintenanceFileSizeThreshold			Specifies the upper size limit of the file (for the store's associated store on disk) at which the store will perform maintenance operations (compaction on disk) to attempt to free up disk space.
MaintenanceMemorythreshold			Specifies the memory limit for a store, after which the store will attempt to free memory by clearing caches and performing any other maintenance operations.
MaxFileSize	N	1000000000 (1GB)	Specifies the maximum size (in bytes) of a channel/queue file on disk, before automatic maintenance is performed to re-order the file to remove purged events. The default is 1GB, although maintenance is usually

Name	Required	Default	Description
			performed well before this limit is reached.
MaxMemory	N	Uses the -Xms (minimum heap size) value	Specifies a target value for the maximum memory usage. As this value is approached, the realm server will attempt to free caches in order to release memory.
mode	N		If set to IPAQ forces a small memory mode for the server
REALM	Y		Specifies the name of the Realm Server
SECURITYFILE	N		Used to specify the Super Users for this realm. Format is user@host (one per line). Note that this is only a bootstrap method on startup of a realm. If you had previously started the realm before specifying a SECURITYFILE, you will need to remove the files realms.nst and realms.nst_old from the RealmSpecific directory, then restart the realm with the -DSECURITYFILE setting in the nserver.conf or nserverdaemon.conf file for the super user entries in the file to be added to the realm ACL.

Notes on DATADIR

The server parameter DATADIR defines a location where Universal Messaging stores its internal status and configuration information for the realm. This location is set

automatically when you create a realm. Normally you do not need to change this location.

DATADIR points to the location `server/<realmname>/data`. This location contains various files and subdirectories. If you define any plugins, Universal Messaging stores the related internal configuration information in the parallel location `server/<realmname>/plugins`.

Note: The `data` and `plugins` locations contain various files and subdirectories that are essential for the smooth running of the product, and any attempt by the customer to modify these files directly may make the system unusable.

The value of the DATADIR parameter is set in the following files that are located in the realm-specific location `server/<realmname>/bin`:

- `CertificateGenerator.conf`
- `env.bat` on Windows, or `env.sh` on UNIX-based systems
- `nstopserver.bat` on Windows, or `nstopserver` on UNIX-based systems
- `Server_Common.conf`

Note: On UNIX-based systems, the environment variable in the `nstopserver` file is named `data` instead of `DATADIR`.

There are some rare cases when you might want to change the value of DATADIR to point to a new location. Such cases are:

- You want to point DATADIR to a new location but you want the realm server to continue using all of its existing status and configuration data.
- You want to point DATADIR to a new location but you do not want the realm server to use any existing status or configuration data.

The steps required to implement these cases are described in the following sections.

Case 1: Define a new DATADIR, retaining the existing realm status and configuration data

1. If the realm server is running, stop the realm server.
2. In a temporary location, make a copy of the contents of the `server/<realmname>/data`. Also copy `server/<realmname>/plugins` if you have defined any plugins.
3. Change the value of DATADIR in the files listed above. Ensure there are no trailing spaces in the new value. In the `nstopserver` file on UNIX-based systems, change the value of the variable `data` instead of `DATADIR`.

If you specify a value for DATADIR that does not end with `"/data"`, the realm server will assume a value for DATADIR as if you had appended `"/data"` to the value.

For example, if you define DATADIR to be `"/folder1/folder2/data"`, then the realm server uses exactly this value as the new DATADIR. If however you define DATADIR to be `"/dir1/dir2"`, then the realm server will behave as if you had specified `"/dir1/dir2/data"`.

4. The following steps 5-7 describe how to create the structure for the new `data` directory and copy in the required files from the old `data` directory. Note that there is an alternative method, described at the end of this procedure.
5. Start the realm server. This automatically creates the structure for the `data` directory at the location defined by the new DATADIR.
6. Stop the realm server.
7. Overwrite the contents of the new `data` directory and all subdirectories with the temporary copy of the `data` directory and all subdirectories you stored earlier.
8. If you had defined plugins, copy in the `plugins` directory and contents from your temporary copy.
9. Start the realm server.

As an alternative to steps 5-7 above, you can create the new `data` directory manually using standard operating system commands while the realm server is stopped, and copy in the contents of the old `data` directory and all subdirectories.

Case 2: Define a new DATADIR, without using any existing realm status or configuration data

If you want to define a new value for DATADIR but you do not want to copy in the existing status and configuration data from the old DATADIR location, follow the instructions in Case 1 above but omit any steps that involve manually copying data to the new DATADIR location. This means that when the realm server starts for the first time after DATADIR has been changed, the realm server has no knowledge of any previously stored data.

Template files

Note that there are also template files in the location `server/templates`. These template files are used to automatically create the above mentioned realm-specific files when you create a realm. If you want to create new realms and ensure that they also use the new value of DATADIR, you need to change the contents of these template files before you create the new realms.

Server Security for Deployment

Universal Messaging provides configurable security for authentication and entitlements. When a user connects using SSL, the server must have an SSL enabled interface configured. Once the interface is configured correctly, clients can connect to a realm using an SSL encrypted session.

Before clients can use the realm correctly, the correct permissions must be granted to each user within the ACLs for the realm, resources (see "[Messaging Paradigms supported](#)" on page 17) and services. For more information on this please see the security section.

Deployment

The structure and target audience for any Enterprise Application determines the deployment topology for the underlying infrastructure that supports it. Universal Messaging provides a wide degree of flexibility when faced with stringent deployment requirements. Key features are discussed below

Persistence and Configuration

Universal Messaging uses its own persistent stores that remain relative to its installation location on your file system. Multiple realms can be configured from a single installation, each with their own configuration files and persistent stores for event storage.

Configuration Snapshots

All aspects of a Universal Messaging realms configuration can be stored in an XML file. Channels, ACL's, Interface configuration, Plugins etc. can all be included. New realms can quickly be bootstrapped from existing configuration files making the deployment to new environments quick and simple.

Interfaces

Universal Messaging Realms can support multiple communications interfaces, each one defined by a protocol a port. Universal Messaging Realms can be configured to bind to all network interfaces on a machine or specific ones depending on configuration requirements. Specific SSL certificate chains can be bound to specific interfaces thus insuring clients always authenticate and connect to specific interfaces

Web Deployment

As well as providing a wide range of client web technology support Universal Messaging's realm server provides useful features to aid web deployment. In addition to providing a messaging backbone for external clients Universal Messaging can also act as a web server delivering static and server-generated content to clients. This resolves security sandbox problems and port use.

Forward and reverse proxy server functionality is available for those that wish to host web content on a different server but stream real time data from Universal Messaging.

Connecting to multiple realms using SSL

This Section describes how to connect to multiple Universal Messaging realms using SSL when different certificate hierarchies are used on each respective realm. The information below applies to any of the various wire protocols (see "[Communication Protocols and RNAMEs](#)" on page 27) that Universal Messaging supports, such as SSL enabled sockets (nsp) and HTTPS (nhps). Please note that the example programs contained in the Universal Messaging package will all work with SSL enabled on the realm server.

The certificate requirements differ depending on whether the realms require client certificate authentication or not. Let us assume that we want to connect to 2 realms over nsps, realmA and realmB. RealmA has interface nsps0 which uses a certificate signed by CA1, while RealmB has interface nsps0 which uses a certificate signed by CA2. The next few paragraphs describe what needs to be done for each possible configuration.

The Universal Messaging Client API supports both custom SSL properties that can be set per UM session and JSSE properties that can be set via system properties.

When using custom SSL properties that are set on the session level, you can set keystore/trustore and other SSL properties for each of your connections.

When using JSSE, only 1 keystore file/keystore password and 1 truststore file/truststore password can be used. In order to achieve your goal you will then have to create a combined keystore and / or a combined truststore depending on your configuration.

Setting Custom SSL Properties on a Session

Setting SSL properties directly is done using the session attributes. Before initialising the session ensure, that your session attributes for each session are set with the correct keystore and truststore information. You can also set the alias of the certificate to use when client certificate authentication is enabled on the interface you are connecting to (see ["Client SSL Configuration" on page 100](#) for more info). If you are using JMS then you will need to set the SSL properties directly on the connection factory or on the create connection overload method provided as part of the connection factory implementation (see ["JMS Client SSL Configuration" on page 102](#) for more info).

Setting SSL Properties using JSSE

Client certificate authentication NOT required

In the case where client certificate authentication is not required by both realms, your application needs to use a combined truststore / truststore password only using the -DCAKEystore and -DCAKEystorePASSWORD parameters.

1. Both CA1 and CA2 are well known Root Certificate Authorities

All well known Root CAs are already included in the JRE cacerts file which can be found in jre\lib\security\cacerts . Unless you have manually changed that keystore's password the default password is changeit. You have to use these values for your -DCAKEystore and -DCAKEystorePASSWORD parameters.

2. CA1 is a well known Root Certificate Authority but CA2 is not (or vice versa)

Two choices are available for this configuration. Either you add CA2's certificate to the JRE cacerts file or you create a combined keystore with CA2's certificate and CA1's certificate. You have to use these values for your -DCAKEystore and -DCAKEystorePASSWORD parameters.

3. CA1 and CA2 are not well known Root Certificate Authorities

In this instance you have to create a combined truststore file that contains both the CA1 and CA2 certificates. In order to do this export your CA certificates from their current JKS store files then create a new JKS file and import them. You can do this

using the JDK keytool command line utility. Finally you have to use these values for your -DCAKEystore and -DCAKEystorePASSWD parameters.

Client certificate authentication required

In the case where client certificate authentication is required by both realms, your application needs to use a combined keystore / keystore password and a combined truststore / truststore password using the -DCKEystore, -DCKEystorePASSWD, -DCAKEystore and -DCAKEystorePASSWD parameters respectively.

1. Both CA1 and CA2 are well known Root Certificate Authorities

With regards to the truststore, all well known Root CAs are already included in the JRE cacerts file which can be found in `jre\lib\security\cacerts`. Unless you have manually changed that keystore's password the default password is `changeit`. You have to use these values for your -DCAKEystore and -DCAKEystorePASSWD parameters.

With regards to the keystore, you need to create a combined keystore that contains both client certificates and then point the -DCKEystore parameter to its path as well as set the -DCKEystorePASSWD to the password of that combined keystore. In order to create a combined keystore, export the certificates and private keys in PKCS#12 format and then import them as trusted certificates in the same keystore file. You can do this using the JDK keytool command line utility.

2. CA1 is a well known Root Certificate Authority but CA2 is not (or vice versa)

The easiest way for this configuration option is to create a single JKS file that contains the CA1 certificate, the CA1 signed client certificate, the CA2 certificate and the CA2 client certificate. You then have to use the same values for CKEystore, CAKEystore and CKEystorePASSWD, CAKEystorePASSWD respectively.

3. CA1 and CA2 are not well known Root Certificate Authorities

Again the easiest way for this configuration option is to create a single JKS file that contains the CA1 certificate, the CA1 signed client certificate, the CA2 certificate and the CA2 client certificate. You then have to use the same values for CKEystore, CAKEystore and CKEystorePASSWD, CAKEystorePASSWD respectively.

Environment Settings

The CKEystore, CKEystorePASSWD, CAKEystore and CAKEystorePASSWD system properties are used by the Universal Messaging sample apps, but are mapped to system properties required by a jsse enabled JVM by the utility program 'com.pcbsys.foundation.utils.fEnvironment', which all sample applications use. If you do not wish to use this program to perform the mapping between Universal Messaging system properties and those required by the JVM, you can specify the SSL properties directly. To do this in your own applications, the following system properties must be set:

```
-Djavax.net.ssl.keyStore=%INSTALLDIR%\client\Universal Messaging\bin\client.jks
-Djavax.net.ssl.keyStorePassword=password
-Djavax.net.ssl.trustStore=%INSTALLDIR%\client\Universal Messaging\bin\nirvanacacerts.jks
-Djavax.net.ssl.trustStorePassword=password
```

where :

javax.net.ssl.keyStore is the client keystore location

javax.net.ssl.keyStorePassword is the password for the client keystore

javax.net.ssl.trustStore is the CA keystore file location

javax.net.ssl.trustStorePassword is password for the CA keystore

As well as the above system properties, if you are intending to use https, both the Universal Messaging sample apps and your own applications will require the following system property to be passed in the command line:

`-Djava.protocol.handler.pkgs="com.sun.net.ssl.internal.www.protocol"`

As well as the above, the RNAME (see ["Communication Protocols and RNAMEs" on page 27](#)) used by your client application must correspond to the correct type of SSL interface, and the correct hostname and port that was configured earlier.

The Java Virtual Machine

Selecting and Tuning a JVM is an important part in running any Java application smoothly. Applications with low latency requirements often require more attention paid to the JVM, as the JVM is often a big factor in performance.

This section outlines JVM selection, and advice on tuning for low latency applications on these JVMs. There are many different JVM vendors available and each JVM has slightly different configurable parameters. This section outlines a few key vendors and important configuration parameters.

Selecting a Java Virtual Machine

As mentioned above, there are a variety of JVMs to choose from that come from different vendors. Some of these are free, some require a license to use. This section will outline the standard Oracle HotSpot VM, the Oracle JRockit VM and the Azul Zing VM.

Oracle HotSpot VM

The Oracle HotSpot VM is offered by Oracle, the JDK and JRE freely available from the Oracle Website. This JVM is suitable to fulfil most users needs for Universal Messaging.

Oracle JRockit VM

The Oracle JRockit VM is another offering from Oracle. This JVM was made free and publicly available in May 2011. It contains many of the same assets from the Oracle HotSpot VM. The performance capabilities of this VM are often advertised as greater than that of the HotSpot VM, however this varies depending on the usage scenario.

Oracle plan to eventually merge the code for the HotSpot and JRockit VMs in the future. This is not expected to happen until JDK8 at the earliest however.

Azul Zing VM

The Azul Zing VM is a commercial offering from Azul. Its primary feature is a 'Pauseless Garbage Collection'. This VM is well suited to applications which require the absolute lowest latency requirements. Applications which experience higher garbage collection pause times may also benefit from using this VM.

Configuring the Java Virtual Machine - Oracle HotSpot

This section covers parameters for the Oracle HotSpot VM which may help improve application performance. These settings can be applied to a Universal Messaging Realm Server by editing the `nserver.conf` or `nserverdaemon.conf` file found under the `server/<realm>/bin` directory of your installation, where `<realm>` is the name of the Universal Messaging realm.

Note: `nserver.conf` is used if you start the realm server manually, e.g. from the command line. `nserverdaemon.conf` is used if you start the realm server as a Windows service or as a UNIX daemon.

General Tuning Parameters

Below are some suggestions of general tuning parameters which can be applied to the HotSpot VM.

<code>-Xmx</code>	The maximum heap size of the JVM.
<code>-Xms</code>	The minimum heap size of the JVM. Set this as equal to the maximum heap size
<code>-XX:+UseLargePages</code>	Allows the JVM to use large pages. This may improve memory access performance. The system must be configured to use large pages.
<code>-XX:+UseNUMA</code>	Allows the JVM to use non uniform memory access. This may improve memory performance

Monitoring Garbage Collection Pause Times

It is important to collect proper monitoring information when tuning an application. This will allow you to quantify the results of changes made to the environment. Monitoring information about the Garbage collection can be collected from a JVM without any significant performance penalty.

We recommend using the most verbose monitoring settings. These can be activated by adding the following commands to the `nserver.conf` or `nserverdaemon.conf` file:

```
-verbose:gc
-XX:+PrintGCDetails
-XX:+PrintGCDateStamps
-XX:+PrintGCApplicationStoppedTime
```

This will produce output similar to the following:

```
2012-07-06T11:42:37.439+0100:
    [GC
      [ParNew: 17024K->1416K(19136K), 0.0090341 secs]
        17024K->1416K(260032K), 0.0090968 secs]
      [Times: user=0.02 sys=0.01, real=0.01 secs]
```

The line starts by printing the time of the garbage collection. If Date Stamps are enabled, this will be the absolute time, otherwise it will be the uptime of the process. Printing the full date is useful for correlating information taken from the nirvana logs or other application logs.

The next line shows if this is a full collection. If the log prints GC, then this is a young generation collection. Full garbage collections are denoted by the output *Full GC (System)*. Full garbage collections are often orders of magnitude longer than young garbage collections, hence for low latency systems they should be avoided. Applications which produce lots of full garbage collections may need to undergo analysis to reduce the stress placed on the JVMs memory management.

The next line displays the garbage collectors type. In this example *ParNew* is the Parallel Scavenge collector. Detailed explanation of garbage collectors are provided elsewhere on this page. Next to the type, it displays the amount of memory this collector reclaimed, as well as the amount of time it took to do so. Young garbage collections will only produce one line like this, full garbage collections will produce one line for the young generation collection and another for the old generation collection.

The last line in this example shows the total garbage collection time in milliseconds. The user time is the total amount of processor time taken by the garbage collector in user mode. The system time is the total amount of processor time taken by the garbage collector running in privileged mode. The real time is the wall clock time that the garbage collection has taken, in single core systems this will be the user + system time. In multiprocessor systems this time is often less as the garbage collector utilizes multiple cores.

The last flag will also cause the explicit application pause time to be printed out to the console. This output will usually look like the following:

```
Total time for which application threads were stopped: 0.0001163 seconds
```

If you observe high client latencies as well as long application pause times, it is likely that the garbage collection mechanism is having an adverse affect on the performance of your application.

Tuning Garbage Collection

The Garbage Collector can be one of the most important aspects of Java Virtual Machine tuning. Large pause times have the capability to negatively impact an applications performance by a noticeable degree. Below are some suggestions of ways to combat specific problems observed by monitoring garbage collection pause times.

Frequent Full Garbage Collections

Full garbage collections are expensive, and often take an order of magnitude longer than a young generation garbage collection to complete. This kind of collection occurs when the old generation is full, and the JVM attempts to promote objects from the younger generation to the older generation. There are two scenarios where this can happen on a regular basis:

1. There are many live objects on the heap, which are unable to be cleaned up by the JVM.
2. The allocation rate of objects with medium-long lifespans is exceptionally high

If the information from garbage collection monitoring shows that full garbage collections are removing very few objects from the old generation, and that the old generation remains nearly full after a old generation collection, it is the case that there are many objects on the heap that cannot be cleaned up.

In the case of a Universal Messaging Realm Server exhibiting this symptom, it would be prudent to do an audit of data stored on the server. Stored Events, ACL Entries, DataGroups, Channels and Queues all contribute to the memory footprint of Universal Messaging. Reducing this footprint by removing unused or unnecessary objects will reduce the frequency of full collections.

If the information from garbage collection monitoring shows that young garbage collection results in many promotions on a consistent basis, then the JVM is likely to have to perform full garbage collections frequently to free space for further promotions.

This kind of heap behaviour is caused by objects which remain live for more than a short amount of time. After this short amount of time they are promoted from the young generation into the old generation. These objects pollute the old generation, increasing the frequency of old generation collections. As promotion is an expensive operation, this behaviour often also causes longer young generation pause times.

Universal Messaging will mitigate this kind of problem by employing a caching mechanism on objects. To further decrease the amount of objects with this lifespan it is important that the administrator perform an audit of creation of resources, such as events, acl entries, channels, datagroups or queues. Heavy dynamic creation and removal of ACL Entries, Channels, DataGroups and Queues may induce this kind of behaviour.

If an administrator has done everything possible to reduce the static application memory footprint, as well as the allocation rate of objects in the realm server then changing some JVM settings may help achieve better results.

Increasing the maximum heap size will reduce the frequency of garbage collections. In general however larger heap sizes will increase the average pause time for garbage collections. Therefore it is important that pause times are measured to ensure they stay within an acceptable limit.

Long Young Generation Collection Pause Times

As mentioned above the primary cause of long young generation pauses is large amounts of object promotion. These objects often take the form of events, ACL entries, channels, datagroups and queues being created.

To minimise the amount of object creation during normal operating hours it is suggested to employ static creation of many channels, datagroups and queues at start up time. This will result in these objects being promoted once at the beginning of operation, remaining in the old generation. Analysing where possible events can be given short lifespans (possibly even made transient) will also reduce the amount of promotion, as these objects will become dereferenced before they are eligible to be moved to the old generation.

It is important to remember that the Java Virtual Machine's memory subsystem performs best when long living objects are created in the initialisation stage, while objects created afterwards die young. Therefore designing your system to create long lived objects like channels at startup and objects like events to be short lived allows Universal Messaging to harmoniously work with the underlying JVM.

Long Full Collection Pause Times

Full Garbage collections which take long periods of time can often be remedied by proper tuning of the underlying JVM. The two recommended approaches to reducing the amount of time spent in full garbage collections is detailed below.

The first approach would be to reduce the overall heap size of the application. Larger heaps often increase the amount of time for a garbage collection cycle to finish. Reducing the heap will lower the average time that a garbage collection cycle takes to complete. Smaller heap sizes will require garbage collecting more often however, so it is important to ensure that you balance the need for lower collection times with collection frequency.

If you are not able to reduce the heap size any further, because garbage collection frequency is increasing, it may be beneficial to change the type of garbage collector used. If you are experiencing high maximum latencies correlated with long GC times it may be beneficial to switch to using the CMS collector.

The Concurrent Mark Sweep (CMS) collector aims to minimize the amount of time an application is paused by doing many of its operations in parallel with the application. This collector can be enabled by adding the following parameter to `nserver.conf` or `nserverdaemon.conf`:

```
-XX:+UseConcMarkSweepGC
```

CMS Collections will usually take more time overall than those done with the Parallel Collector. Only a small fraction of the work done by the CMS collector requires the application to pause however, which will generally result in improved response times.

The Network

This page details important network settings that may improve network performance when using Universal Messaging. Many of the commands detailed on this page are

specific to Red Hat Linux, though many of the concepts apply globally to all operating systems.

Stop the iptables Service

The iptables service is used to control packet filtering and NAT. In many cases it is not necessary to run this service and a minor performance gain can be seen by disabling this service. To disable this service use the following command:

```
service iptables stop
service ip6tables stop
```

Disable Adaptive Interrupts on Network Interfaces

Interrupts on a network interface notify the system that some network task is required to be run, for example reading some data from the network. Adaptive Interrupts control the rate at which interrupts are generated for these tasks. The delay in processing subsequent interrupts from interrupt coalescing may degrade performance.

```
ethtool -C eth0 adaptive-rx off
```

Disabling adaptive interrupts on an interface will make that interface use the set interrupt rate. This rate will delay interrupts by a set number of microseconds. The minimum value that this delay can be is 0 (immediate). To set this value on an interface use the command:

```
ethtool -C eth0 rx-usecs-irq 0
```

Kernel Settings

The Kernel has many network settings for TCP which can provide performance improvements when correctly tweaked. This section will outline a few suggestions, however care should be taken when changing these parameters. It is also important to validate results as your mileage may vary. These settings should be added to the sysctl.conf file.

Increase Socket Memory

The settings below will increase the amount of memory allocated by the kernel to tcp sockets.

Important: It is important to set these limits to a reasonable level for the amount of memory available on your machine.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216
net.ipv4.tcp_mem = 50576 64768 98152
```

Increase Backlog Queue Size

The command below will increase the queue size in packets waiting to be processed by the kernel. This queue fills up when the interface receives packets faster than the kernel can process them.

```
net.core.netdev_max_backlog = 2500
```

Increase the local Port Range

Applications which have to manage large numbers of current connections may find that they will run out of ports under the default settings. This default can be increased by using the following command:

```
net.ipv4.ip_local_port_range = 1024 65535
```

The maximum number of allocated ports are 65535, 1024 of these are reserved. Applications which manage extremely high numbers of connections will require more ports than this. One way to get around these limits would be to create multiple virtual network interfaces.

The Operating System

This page details important operating system and kernel settings that should be considered when optimising the server. The focus of this page is geared towards Red Hat Linux. Many of the suggestions here have synonymous commands under Solaris or Windows, which can be applied to have a similar effect.

Configuring User Limits

Unix has a configurable limit on the number of processes, file descriptors and threads available per user. This functionality is aimed to prevent a user from consuming all of the resources on a machine. These limits are often set to a reasonably low level, as a general purpose user will not consume many of these objects at any one time.

Application Servers like Universal Messaging may, if under considerable load, wish to consume a large number of these resources. Each open connection to a client for example consumes a file descriptor, and application servers which can support tens of thousands of concurrent connections will thus require as many file descriptors. It is therefore important to increase these limits for Universal Messaging.

Temporarily Increasing limits using the ulimit command

`ulimit` is a unix command which can be used to alter user limits. To increase the user limits which Universal Messaging consumes the following commands are recommended:

```
ulimit -n 250000
ulimit -u 10000
```

This will increase the number of file descriptors and the number of user processes allowed. Any processes spawned from the terminal this was entered on will inherit these limits.

Permanently Increasing User Limits

It is also possible to permanently increase the user limits by editing the relevant configuration file. This configuration file can usually be found in `/etc/security/limits.conf`.

```
user    soft    nofile   250000
```

user	hard	nofile	250000
user	soft	nproc	10000
user	hard	nproc	10000

Disabling Processor Power Saving States

Many new processors have mechanisms which allow them to dynamically turn individual cores on and off to save power. These mechanisms may sometimes degrade processor and memory performance. In applications that require consistent low latency performance it is recommended to disable this feature.

Many processors manage this by using the `cpuspeed` service. This service can be disabled, which on many machines and architectures will turn this functionality off.

```
service cpuspeed stop
```

Some processors however will require further work to disable power saving states. Whether or not your processor will require extra configuration and what those configuration steps are vary from processor to processor. Many Intel processors for example may require the following command to be appended to the boot options of your operating system

```
intel_idle.max_cstate=0
```

As mentioned above however, this will not be necessary for all processors. Consult with your processor specific documentation for information on disabling power saving states.

Stop the Interrupt Request Balance Service

Interrupts are signals generated, generally by devices, to notify a CPU that there is processing which needs to be done. Interrupt Request (IRQ) Balancing is the act of dividing these processes up between cores on a CPU. In some situations this may harm performance of applications running on the CPU, as these interrupts consume processor cycles and loads information into memory.

Disabling IRQ balancing will assign all interrupts to a single core by default. It is possible to assign interrupts to certain cores, but that is beyond the scope of this section. To disable IRQ balance, use the following command.

```
service irqbalance stop
```

The Realm Server

This page details important configuration options that can be changed on the server to improve performance. Important monitoring information that can be collected using the Admin API is also mentioned here. This monitoring information can be used to diagnose common problems.

Lowering the Log Level

Logging information can be useful for debugging a variety of problems that may occur. However, particularly verbose logging levels can negatively affect the performance of your application. Logging creates extra objects which increases memory usage, and also promotes contention between threads which wish to print to the logger.

The verbosity of the logging can be changed by using the Enterprise Manager. The most verbose log level is 0 and produces the most output. The least verbose log level is 6, which produces very little output. Unless you are attempting to gather logging information related to a particular issue, it is recommended to use a the log level no more verbose than 5 (Quiet). Particularly demanding applications may wish to reduce the verbosity to logging level 6.

For client side logging please see the relevant API documentation.

Increasing the Peak Watermark

The server is configured to enter a peak operating mode when a certain number of messages are being delivered through the server per second. Peak mode will batch messages in an effort to keep server load at an optimal level. This batching may increase average latencies for clients.

It is possible to raise the peak mode threshold so that the server does not utilise peak mode until a much higher load is reached. It is important to stress that beyond a certain point the non batching performance will suffer as machine limitations are reached.

Machines with good hardware will benefit from having this threshold raised, but slower machines may function better in batching mode after a certain message rate is reached.

Enable Low Latency Fanout Mechanism

The most aggressive fanout mechanism Universal Messaging provides is called *spin locking*. By default, this is disabled. This particular mechanism is capable of meeting extremely demanding latency and message rate requirements, however is very demanding on the resources of a system. It is disabled to prevent it consuming resources on machines with less resources (for example development machines).

If the hardware which the Universal Messaging server runs has greater than 8 cores it is recommended that you enable this fanout mechanism to produce the best latencies. This fanout mechanism will consume multiple cores entirely, so will therefore increase the load average of the machine. It is important that you have sufficient free cores available, as otherwise it is possible that this mode will cause Universal Messaging to starve other threads/processes running on the system.

The mechanism can be enabled by adding the following flag to the `nserver.conf` or `nserverdaemon.conf` file under the `server/<realm>/bin` directory of your installation, where `<realm>` is the name of the Universal Messaging realm:

```
-DCORE_SPIN=true
```

Note: `nserver.conf` is used if you start the realm server manually, e.g. from the command line. `nserverdaemon.conf` is used if you start the realm server as a Windows service or as a UNIX daemon.

There are further flags that can be applied to the `nserver.conf` or `nserverdaemon.conf` file to customize the behaviour of this fanout setting. The first of these flags can be used to adjust the number of times that this fanout mode will spin on a core attempting to do work before switching to a less aggressive fanout mechanism.

```
-DSPIN_COUNT=1000000000
```

The default value for this spin count is one billion. Reducing this value will generally encourage the server to switch to a CPU intensive fanout mechanism, if the server is not busy. Reducing this value may result in a performance penalty which occurs as a result of using the less intensive fanout mechanism. The maximum value of this parameter is the same as Long.MAX_VALUE.

The less aggressive fanout mechanism employs a wait as opposed to spinning mechanism, the second flag can be used to specify the wait time between checking if work is available.

```
-DSPIN_WAIT=1
```

This parameter will alter the number of nanoseconds which the fanout mechanism will wait for between checking if it has tasks to complete. Increasing this number will decrease the CPU consumption of Universal Messaging, but at a cost to latency.

Periodic Logging of Server Status

The Universal Messaging server writes status information to the log file at regular intervals. The default interval can be configured using the `StatusBroadcast` realm configuration property, and the default value is 5 seconds).

For information on realm configuration properties, see the section *Realm Configuration* in the *Enterprise Manager* part of the *Administration Guide*.

Here is a sample status log message:

```
ServerStatusLog> Memory=3577, Direct=3925, EventMemory=0,
Disk=277070, CPU=0.2, Scheduled=29, Queued=0, Connections=5,
BytesIn=12315, BytesOut=19876, Published=413, Consumed=1254,
QueueSize=0, ClientsSize=0, CommQueueSize=0
```

The following table describes the status log metrics:

Metric	Description	Unit
Memory	Free JVM memory	Megabyte
Direct	Free direct memory	Megabyte
EventMemory	JVM memory consumed by events	Megabyte
Disk	Free disk space in UM data directory	Megabyte
CPU	Recent CPU usage for the JVM process	A floating point number in the [0.0,1.0] interval. A value of 0.0 means that none of the CPUs were running threads from the

Metric	Description	Unit
		JVM process during the recent period of time observed, while a value of 1.0 means that all CPUs were actively running threads from the JVM 100% of the time during the recent period being observed. All values between 0.0 and 1.0 are possible depending of the activities going on in the JVM process and the whole system.
Scheduled	Number of tasks in the internal scheduler	
Queued	Number of total tasks currently queued in all thread pools	
Connections	Active Universal Messaging connections	
BytesIn	Total bytes received for currently active Universal Messaging connections	byte
BytesOut	Total bytes sent for currently active Universal Messaging connections	byte
Published	Total number of consumed events	
Consumed	Total number of published events	
QueueSize	Cluster queue size (present only if the realm is part of a cluster)	

Metric	Description	Unit
ClientsSize	Cluster client request queue size (present only if the realm is part of a cluster)	
CommQueueSize	Cluster communication queue size (present only if the realm is part of a cluster)	

During server startup and each time the log file is rolled, a status log legend is printed in the log file:

```
ServerStatusLog> Activating Realm status log with format:
Memory          - free memory (MB)
Direct           - free direct memory (MB)
EventMemory      - used event memory (MB)
Disk             - free disk space (MB) in server data directory
CPU              - cpu load average
Scheduled        - scheduled tasks
Queued           - total queued tasks in all thread pools
Connections      - active connections count
BytesIn          - total bytes received (for all active connections)
BytesOut         - total bytes send (for all active connections)
Published        - total count of published events
Consumed         - total count of consumed events
QueueSize        - cluster queue size
ClientsSize      - cluster client request queue size
CommQueueSize    - cluster communication queue size
```

Periodic logging of the server status can be disabled using the `EnableStatusLog` realm configuration property.

For information on realm configuration properties, see the section *Realm Configuration* in the *Enterprise Manager* part of the *Administration Guide*.

For information on log rolling, see the section *Universal Messaging Enterprise Manager : Logs Panel* in the *Enterprise Manager* part of the *Administration Guide*.

Client

Using HTTP/HTTPS

The Universal Messaging messaging APIs provides a rich set of functionality that can be used over sockets, SSL, HTTP and HTTPS. The code used to connect to the Universal Messaging server is the same regardless of which network protocol you are using to connect.

Under the Universal Messaging programming model there are a number of logical steps that need to be followed in order to establish a connection to a Universal Messaging

sever (Realm). These involve establishing a session, obtaining a reference to a channel or a transaction, or registering an object as a subscriber.

Universal Messaging fully supports HTTP and HTTPS. Rather than tunnel an existing protocol through HTTP Universal Messaging has a pluggable set of communications drivers supporting TCP/IP Sockets, SSL enabled TCP/IP sockets, HTTP and HTTPS. Both the client and server make use of these pluggable drivers. From the server perspective different driver types can be assigned to specific Universal Messaging interfaces. From a client perspective a Universal Messaging session can be built on any one of the available drivers dynamically.

Please note that before making an HTTP/HTTPS connection to a Universal Messaging realm server you will first need to add a HTTP/HTTPS interface to the realm. See the Enterprise Manager documentation for details.

To create a connect to a Universal Messaging Realm over HTTPS you would use an RNAME (see "[Communication Protocols and RNAMEs](#)" on page 27) that specific the Universal Messaging HTTPS protocol (nhps) as follows:

1. Create a nSessionAttrib object with the RNAME value of your choice

```
//use an RNAME indicating the wire protocol you are using (HTTPS in this case)
//you can pass an array of up to four values for RNAME for added robustness
String[] RNAME= ( {"nhps://remoteHost:443" } );
nSessionAttrib nsa = new nSessionAttrib( RNAME );
```

2. Call the create method on nSessionFactory to create your session

```
nSession mySession = nSessionFactory.create( nsa );
```

Alternatively, if you require the use of a session reconnect handler to intercept the automatic reconnection attempts, pass an instance of that class too in the create method:

```
Public class myReconnectHandler implements nReconnectHandler {
myReconnectHandler rhandler = new myReconnectHandler( );
nSession mySession = nSessionFactory.create( nsa, rhandler );
```

3. Initialise the session object to open the connection to the Universal Messaging Realm

```
mySession.init();
```

After initialising your Universal Messaging session, you will be connected to the Universal Messaging Realm using HTTPS. From that point, all functionality is subject to a Realm ACL check. If you call a method that requires a permission your credential does not have, you will receive an nSecurityException.

For detailed information including code samples for connecting to Universal Messaging over HTTP/HTTPS please see our developer guides for the language you require.

Browser / Applet Deployment

Introduction

Universal Messaging client applications can run within stand alone applications as well as within Java applets loaded via a web browser such as Chrome, Firefox, MS Internet Explorer, Microsoft Edge, Netscape, etc.

The Universal Messaging client APIs can be used with most Java Plugin versions.

Applet Sandbox / Host Machine Limitation

Applets run within a client's browser, and are subject to strict security limitations as defined by the Applet Model. These limitations need to be considered when deploying applets. One such limitation is that the applet is only allowed to communicate with the host machine from which the applet was downloaded. This restricts the applet to only being permitted to make connections to the applet host machine. This has a number of implications for an applet that uses Universal Messaging's APIs.

Universal Messaging's APIs communicate with a Realm Server (or potentially multiple servers in a cluster). This limitation means that the applet source host must be the same hostname as each Universal Messaging Realm in use by the applet. If the applet is served from a web server, such as Apache, and it is assumed the communication protocol required for Universal Messaging communication is nhp/nhps (http/https). The usual ports used by web servers running http and https are 80 and 443 respectively. Since the web server uses these ports and the realm servers need to run on the same machine with these ports there is obviously a problem since these ports are in use.

However, Universal Messaging provides 2 different methods for ensuring this is not a problem. The first is Universal Messaging's ability to act as a web server through its file plugin. Running a file plugin on an nhp or nhps interface enables the realm server to deliver the applet to the client browser, this removing the need for the web server and of course freeing up the ports for use by the realm server interfaces.

The second method can be used when the web server is apache. We can provide an apache module that acts similarly to mod.proxy for apache. This apache module called mod.Universal Messaging allows the web server to proxy all requests for a specific URL to another host. This host can be the realm server running on any other port on the same machine or any other machine, and hence once again fixes this issue.

Another way to circumvent this restriction is to digitally sign the applet and thus allowing the applet to communicate with any host.

Browser Plugins

Universal Messaging can either run within a 4.0 browsers own Java Virtual Machine or run within a Java Virtual machine started using the Java plugin.

Universal Messaging does not require installation of the Java Plugin.

Client Jars

Depending on the functionality used by your Universal Messaging application, different jar files are required. This following table illustrates the deployment dependencies between the jar libraries installed by the Universal Messaging installer.

JAR File	Description	Dependency
nClient.jar	Provides Universal Messaging Client functionality (Pub/Sub & Queues)	(none)
nJMS.jar	Provides Universal Messaging Provider to support JMS functionality	nClient.jar, gf.javax.jms.jar
nJ2EE.jar	Provides Universal Messaging support for interacting with Application servers that support J2EE	nClient.jar, JMS.jar, gf.javax.jms.jar
nAdminAPI.jar	Provides Universal Messaging Administration & Monitoring functionality	nClient.jar
nAdminXMLAPI.jar	Provides Universal Messaging Configuration XML Import / Export functionality	nClient.jar, nAdminAPI.jar
nEnterpriseManager.jar	Contains the Enterprise Manager tool	nClient.jar, nAdminAPI.jar, nAdminXMLAPI.jar (Optional), gf.javax.jms.jar
nServer.jar	Contains the Universal Messaging Realm Server	(none)
nPlugin.jar	Contains the Universal Messaging Server plugins	nServer.jar

JAR File	Description	Dependency
gf.javax.jms.jar	JMS reference jar	(none)

Client Security

Universal Messaging makes use of JSSE for SSL enabled communication. Clients are able to connect using standard sockets (nsp), http (nhp), SSL enabled sockets (nsps) or https (nhps). Universal Messaging's client SSL communication uses the JVM's own SSL implementation.

Clients connecting using SSL (see ["Client SSL Configuration" on page 100](#)) will connect to a realm server that has an SSL enabled interface with either client authentication on or off.

Once authenticated using SSL, the client must have the desired permissions on the realm and its objects in order to perform the operations required. The entitlements are defined within the ACLs for the realm, channels, queues and services. The ACLs must contain the correct level of permissions for clients connecting to the realm.

Please also see the description of managing realm security ACLs in the documentation of the Enterprise Manager.

Client Parameters

The Universal Messaging client API supports a variety of different parameters that can be specified in the command line of any Universal Messaging Client application.

This section describes those -D parameters, what they are used for and their typical values.

Name	Required	Default	Description
LOGLEVEL	N	5	Specifies the current log level to use
LOG_FRAMEWORK	N		Specifies a third part logging framework to use. Can be set to "LOGBACK" or "LOG4J2". The default is "fLogger".
LOGFILE	N	System.out	Used to specify a log file to write the log entries to
LOGSIZE	N	100000	Specified in bytes before the log file is rolled

Name	Required	Default	Description
user.name	N	Signed on name	Used to override the current username without coding it.
HPROXY	N		Used as a short hand to set http.proxyHost and http.proxyPort
http.proxySet	N		Set to true if the url handler is to use a proxy
http.proxyHost	N		Sets the Proxy host name to use
http.proxyPort	N		The proxy port to use
CKEYSTORE	N		Short hand for javax.net.ssl.keyStore
CKEYSTOREPASSWD	N		Short hand for javax.net.ssl.keyStorePassword
CAKEYSTORE	N		Short hand for javax.net.ssl.trustStore
CAKEYSTOREPASSWD	N		Short hand for javax.net.ssl.trustStorePassword
javax.net.ssl.debug	N		Useful to debug SSL issues, see www.javasoft.com for more information
SSLClientProtocol	N		Sets the default SSL Protocol to use. If not supplied then client SSL communication will attempt to use the JVM default protocol, or TLS when that is not possible.

Multiplexing Sessions

Universal Messaging supports the multiplexing of sessions to a specific host in Java and C#. This allows the circumvention of connection limit issues by packing multiple Universal Messaging sessions into one connection, and can be used to allow the same client to set up multiple subscriptions to a given channel or queue if required.

Multiplexing Sessions

To multiplex two sessions, first construct one session, and then create a new session by multiplexing the original session. These two sessions will now appear to act as normal sessions, but will, in fact, share a single connection.

This can be accomplished either by using the `nSession` object associated with the original session, or by using the `nSessionAttributes` used to create this original session. Below are examples of how to multiplex sessions via both methods:

```
//Construct first session
nsa = new nSessionAttributes(realmDetails, 2);
mySession = nSessionFactory.create(nsa, this);
mySession.init();
//Construct second session by multiplexing the first session.
otherSession = nSessionFactory.createMultiplexed(mySession);
otherSession.init();
//Construct a third session by multiplexing the first session's nSessionAttributes.
thirdSession = nSessionFactory.createMultiplexed(nsa);
thirdSession.init();
```

Examples of multiplexing a session are available for Java and C# .NET.

Language Deployment Tips

Silverlight Application Deployment

Silverlight and the Universal Messaging Policy File Server

Universal Messaging has a built in socket policy server for Silverlight applications that require tcp socket or ssl socket (NSP or NSPS) protocol.

Universal Messaging's Silverlight API enables clients to specify the communication protocol as nsp (ie tcp sockets). This means that when a Silverlight application is loaded and tries to make a socket request, Silverlight will first of all attempt to make a policy file request on port 943 to the host the original socket request was being made to. For example, if i specify an RNAME of nsp://myhost.mydomain.com:4502, when the Universal Messaging Silverlight API attempts to construct a socket on port 4502, it will first of all make a socket connect request to myhost.mydomain.com:943, and issue a policy file request.

With this in mind, Universal Messaging enables you to run a policy file server on a socket interface that will automatically handle these requests. You would first need

to create a socket interface on port 943, and select the "Enable Policy Server" check box under the Basic tab for an nsp interface. Once this is setup, you will also need a `clientaccesspolicy.xml` file in the `/install/server/name/plugins/htdocs` directory of the server. This policy file might contain something like the following:

```
<?xml version="1.0" encoding="utf-8"?>
<access-policy>
  <cross-domain-access>
    <policy>
      <allow-from http-request-headers="*">
        <domain uri="*" />
      </allow-from>
      <grant-to>
        <resource path="/" include-subpaths="true" />
        <socket-resource port='4502-4534' protocol='tcp' />
      </grant-to>
    </policy>
  </cross-domain-access>
</access-policy>
```

For further information on socket policies see Microsoft's Silverlight Policy Documentation at <http://msdn.microsoft.com/en-us/library/cc645032%28v=vs.95%29.aspx>.

JavaScript Application Deployment

JavaScript applications can be served directly from a Universal Messaging realm server, using a file plugin, or via a third party web server of your choice.

Serving Applications via a Universal Messaging File Plugin

For performance and security, we strongly recommend that applications are served from an SSL-encrypted file plugin. You may however choose to serve applications from a non-encrypted file plugin. See the description of using JavaScript for HTTP/HTTPS delivery in the Developer Guide.

Serving Applications via a third party Web Server

Most components of your JavaScript application can be served from any web server. A Universal Messaging File plugin is still required however, to serve certain parts of the JavaScript libraries. This is necessary to permit secure cross domain communication.

4 Security

■ Security Overview	72
■ Authentication	72
■ Access Control Lists	95
■ Using SSL	98
■ Setting up and using FIPS	104

Security Overview

Universal Messaging provides a wide range of features to ensure that user access and data transmission is handled in a secure manner.

Universal Messaging includes built in authentication and entitlements functionality. Additionally Universal Messaging can drive 3rd party authentication and entitlements systems or be driven by organizations existing authentication and entitlements systems.

Universal Messaging makes use of standards based cryptography to provide encryption and the signing of events with digital signatures if required. Further information on Universal Messaging's security features can be found below.

Authentication

Authentication Overview

While distributed applications offer many benefits to their users the development of such applications can be a complex process. The ability to correctly authenticate users has been a complex issue and has lead to the emergence of standard Authentication and Authorization frameworks, frameworks such as JAAS.

JAAS authentication is performed in a pluggable fashion. This permits applications to remain independent from underlying authentication technologies. New or updated authentication technologies can be plugged under an application without requiring modifications to the application itself.

Universal Messaging provides a wide variety of client APIs to develop enterprise, web and mobile applications. On the enterprise application front, Universal Messaging offers a transport protocol dependent authentication scheme while on the web and mobile application front a pluggable authentication framework is offered. The end result is that all applications can share the same Universal Messaging authorization scheme which requires a token@host based subject that access control lists can be defined upon.

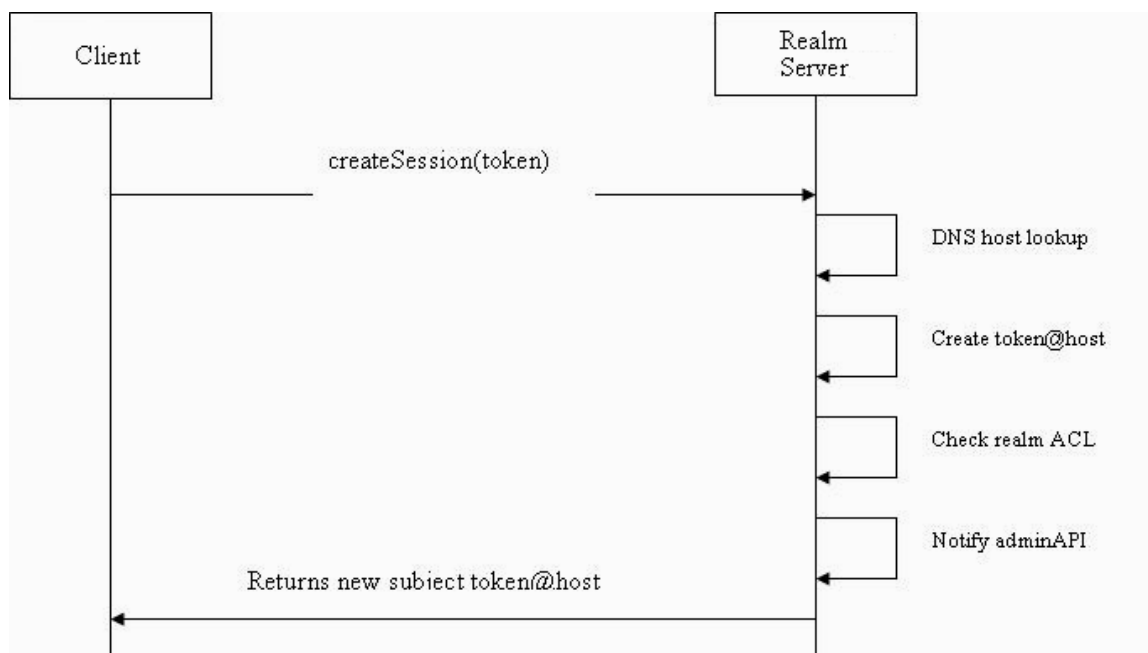
Enterprise Application Authentication

Universal Messaging enterprise applications can be written in a variety of programming languages. Each one of these client APIs offers connectivity using one of the 4 available transport protocols, namely nsp (TCP Sockets), nhp (HTTP), nsps (SSL Sockets) and nhps (HTTPS). The authentication scheme is transport protocol dependent therefore providing a basic authentication scheme for TCP based transport protocols (nsp, nhp) and an SSL authentication scheme for SSL based transport protocols (nsps, nhps).

Basic Authentication Scheme

Under this mode of authentication the client passes the username to the server as part of the initial connection handshake. The server then extracts the remote host name and creates the subject to be used by this connection.

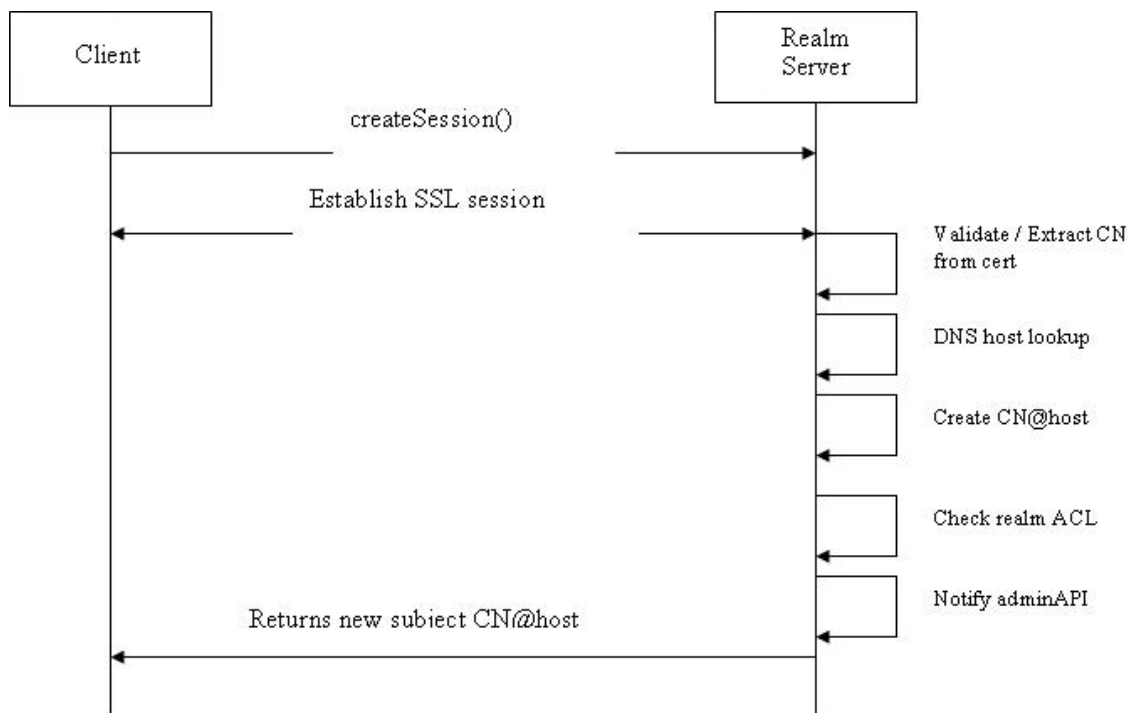
The client API can set the username component, however, the remote host is always set on the server. This stops clients from impersonating users from other hosts. The following diagram illustrates the basic authentication scheme's operation:



SSL Authentication Scheme

The Universal Messaging Realm server can be configured to perform Client Certificate authorization or to allow anonymous SSL clients to connect. When the server is configured to allow anonymous clients to connect the subject is built up based on the previous authentication method. That is the username portion is passed to it from the client.

When the server is configured for client certificate processing the subject is constructed with the Common Name (CN) of the certificate and the remote host name. This allows the ACLs to be configured such that not only is the certificate valid but it can only access the Realm Server from a specific host. The following diagram illustrates the SSL authentication scheme's operation when using client certificates:



Web Application Authentication

Universal Messaging web applications can use a pluggable authentication framework that presents its self as basic http authentication as defined by RFC 1945. Basic authentication is supported by all popular web browsers and users have to enter a username and password in a browser provided login dialog before proceeding. The web browser then automatically includes the token in the Authorization HTTP header for all subsequent requests to the server's authentication realm, for the lifetime of the browser process. Please note that although Universal Messaging supports basic authentication on both nhp (HTTP) and nhps (HTTPS) interfaces, it is only advised to use it over HTTPS connections to secure your web application against man in the middle attacks and network sniffing tools.

In order to host your web application on Universal Messaging, a number of server side plugins are provided that you can configure and mount on the various URLs that your application expects connections on. These are the XML plugin, the Servlet plugin, the SOAP plugin, the File plugin, the REST plugin and the Proxy Pass Through plugin.

Plugin Authentication Parameters

Each one of these plugins contains an identical set of configuration parameters that control its behavior towards authentication. These are described below:

- **Security Realm:** Name of the authentication realm
- **AddUserAsCookie:** Specifies if the authenticated username should be added as a cookie.

- **Authenticator:** Fully qualified class name of authenticator to use, or blank to use the default implementation provided.
- **AuthParameters:** A space delimited list of key=value definitions which are passed to the authenticator instance to initialize and configure it. These are passed to the Authenticator's init method.
- **GroupNames:** An optional comma separated list of groups. The user must be a member of at least one group to be granted access even if a valid username/password is provided. The groups are dependent on the authenticator implementation.
- **RoleNames:** An optional comma separated list of roles. The user must have at least one role to be granted access even if a valid username/password is provided. The roles are dependent on the authenticator implementation and are effectively the permissions defined.
- **ReloadUserFileDynamically:** If set to true, the reload method of the authenticator implementation will be called prior to serving each http request. If set to false, the reload will only be called once when the Universal Messaging interface starts.

Common AuthParameters

Irrespective of the authenticator implementation you use in your Universal Messaging server plugins, there are some AuthParameters that are also used by the server. These are:

- **NamedInstance:** This parameter requests that this authenticator configuration is bound to the specified named instance which will be shared across all plugins on this server that are configured to do so. Please note that the first plugin that accepts a connection will bind the name to the server together with the remaining configuration parameters. For this reason please make sure that configuration is always the same on all plugins that share the same instance.

Default Authenticator Implementation

Universal Messaging comes with a default authenticator implementation that uses a properties file to define users, groups and permissions (roles). In order to enable it on a Universal Messaging plugin, the Authenticator parameter needs to be left empty (this implies using the Default), the Authentication Realm set and one parameter needs to be set in AuthParameters.

The necessary parameter is called UserFile and should point to the full path of a java properties file, e.g. c:\users.txt. In order to get the Universal Messaging realm server to encrypt your user passwords, you need to add a property called initialize as shown below. This notifies the default authenticator that passwords are not encrypted so on the first load it will encrypt them, remove the initialize property and save your user file.

An example of a UserFile defining 3 permissions (roles), 3 groups and 3 users is shown below:

```
#Request password initialisation
initialise=true
#Permissions (Roles) Definition
```

```

perm_name_1=Guest
perm_name_2=User
perm_name_3=Admin
#Guests Group Definition
group_ID_Guests=1
group_desc_Guests=Guests Group
group_perm_Guests={1}
#Users Group Definition
group_ID_Users=2
group_desc_Users=Users Group
group_perm_Users={2}
#Admins Group Definition
group_ID_Admins=3
group_desc_Admins=Admins Group
group_perm_Admins={3}
#Example Guest User Definition
user_desc_somewhat=Some Guest User
user_pass_somewhat=password
user_perm_somewhat={1}
user_home_id_somewhat=Guests
user_group_somewhat=Guests
#Example Regular User Definition
user_desc_someuser=Some User
user_pass_someuser=password
user_perm_someuser={1,2}
user_home_id_someuser=Users
user_group_someuser=Users
user_group_0_someuser=Guests
#Example Admin User Definition
user_desc_someadmin=Some Admin User
user_pass_someadmin=password
user_perm_someadmin={1,2,3}
user_home_id_someadmin=Admins
user_group_someadmin=Admins
user_group_0_someadmin=Guests
user_group_1_someadmin=Users

```

Custom Authenticator Implementations

The interface for creation of custom authenticator implementations is defined in the following 3 classes of the `com.pcbsys.foundation.authentication` package:

```

fAuthenticator: Represents the Authenticator Implementation and
                 has the following methods
public void init(Hashtable initParams);
public String getName();
public synchronized void close();
public void reload();
public fPermission addPermission(int permNo, String name) ;
public fUser addUser(String username, String description, String plainPassword,
String groupName);
public fUser copyUser(fUser user) ;
public fUser getUser(String username);
public void delUser(fUser user);
public fGroup addGroup(int id, String name, String description);
public fPermission getPermission(int id);
public fPermission getPermission(String name);
public fGroup getGroup(String name);
public void delPermission(int id);
public void delGroup(fGroup group);
public void saveState() throws IOException
fGroup: Represents the user groups and contains the following methods:
public void reload(int id, String name, String description);

```

```

public boolean isModified();
public void setModified(boolean flag);
public String getName();
public int getId();
public String getDescription();
public BitSet getPermissions();
public void addUser(fUser aUser);
public Enumeration getUsers();
public Hashtable getUserHash();
public void setUserHash(Hashtable newhash);
public void delUser(fUser aUser);
public int getNoUsers();
public void setPermission(fPermission perm);
public void clearPermission(fPermission perm);
public void resetPermission();
public BitSet getPermissionBitSet();
public boolean can(fPermission perm);
fUser : Represents the authentication users and has the following methods:
public void reload(String name, String description, String password,
    fGroup group);
public void createUser(String name, String description, String password,
    fGroup group) ;
public void setPassword(String pass);
public BitSet getPermissions();
public BitSet getTotalPermissions();
public boolean can(fPermission perm);
public String login(byte[] password, boolean requestToken, Hashtable params);
public String login(String password, boolean requestToken, Hashtable params);
public String getHomeId();
public void setHomeId(String myHomeId);
public void setGroup(fGroup group);
public void delGroup(fGroup group);
public String getName();
public String getDescription();
public String getPassword();
public fGroup getGroup();
public Enumeration getGroups();
public Hashtable getGroupHash();
public void setGroupHash(Hashtable newhash);
public int getNumGroups();
public void setPermission(fPermission perm);
public void setDescription(String desc);
public void clearPermission(fPermission perm);
public void setPermissionBitSet(BitSet newperms);
public BitSet getPermissionBitSet();
public void resetPermission();
public boolean isModified();
public void setModified(boolean flag);

```

Example Database Authenticator

As discussed in the previous section the default implementation is based on an optionally encrypted text file, with passwords being MD5 digested. It is however possible to use different storage mechanisms for users, groups and permissions such as a relational database. There are no restrictions on the design of the database schema as Universal Messaging simply needs a set of classes that comply to the `fAuthenticator`, `fGroup` and `fUser` interfaces. Please note that not all classes need to be subclassed but only the ones that you need to modify the default behaviour.

In the context of this example we are going to use a mysql database running on localhost and containing a users table with the following columns:

- "Name": varchar
- "Password": varchar
- "Rights": int
- "Home": varchar

In order to keep the example simple we are going to statically define the groups and permissions within the authenticator source code. We will use the group functionality on the base fGroup class and therefore will only subclass fAuthenticator and fUser as shown below:

DBAuthenticator

```
package com.myapp;
import com.pcbsys.foundation.authentication.*;
import com.pcbsys.nirvana.client.*;
import com.mysql.jdbc.Driver;
import java.io.*;
import java.sql.Connection;
import java.sql.SQLException;
import java.sql.DriverManager;
import java.util.Date;
import java.util.Enumeration;
import java.util.Hashtable;
import java.util.Vector;
public class DBAuthenticator extends fAuthenticator {
    private static fGroup MYAPP_GROUP =null;
    private static fGroup MYCOMPANY_GROUP =null;
    protected static fPermission CLIENT_PERMISSION=null;
    protected static fPermission ADMIN_PERMISSION=null;
    private static fPermission GUEST_PERMISSION=null;
    private boolean initialised=false;
    private String myName="DBAuthenticator";
    private static int myUniqueID=0;
    private static Connection myConnection;
    private static String jdbcurl = "jdbc:mysql://localhost:3306/test";
    private static String myDBUser="root";
    private static String myDBPassword="";
    //Lets statically define the groups and permissions
    static {
        //Company Group
        MYCOMPANY_GROUP =new fGroup();
        MYCOMPANY_GROUP.reload(2,"mycompany", "MyCompany Group");
        //Application Group
        MYAPP_GROUP =new fGroup();
        MYAPP_GROUP.reload(0,"mycompany/myapp", "MyApp Group");
        GUEST_PERMISSION=new fPermission();
        GUEST_PERMISSION.reload(0,"Guest");
        CLIENT_PERMISSION=new fPermission();
        CLIENT_PERMISSION.reload(1,"Client");
        ADMIN_PERMISSION=new fPermission();
        ADMIN_PERMISSION.reload(4,"Admin");
    }
    public void close(){
        super.close();
        if(getUsageCount() == 0){
```

```

        fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
            "Closing Authenticator [" +getUsageCount()+"]");
        //release connection pool
        if (myConnection!=null){
            try {
                myConnection.close();
            } catch (SQLException e) {}
            myConnection=null;
        }
        initialised=false;
    }
    else {
        fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
            "Closing Authenticator [" +getUsageCount()+"]");
    }
}
public DBAuthenticator() {
    super();
    addGroup(MYCOMPANY_GROUP);
    addGroup(MYAPP_GROUP);
    getPermissionsCollection().put("Client", CLIENT_PERMISSION);
    getPermissionsCollection().put("Admin", ADMIN_PERMISSION);
    getPermissionsCollection().put("Guest", GUEST_PERMISSION);
}
protected static Connection getDBConnection() throws SQLException{
    if (myConnection==null){
        try {
            Class.forName("com.mysql.jdbc.Driver").newInstance();
            myConnection = DriverManager.getConnection(jdbcurl,myDBUser,
                myDBPassword);
        } catch (InstantiationException e) {
            e.printStackTrace();
            //To change body of catch statement use
            // File | Settings | File Templates.
        } catch (IllegalAccessException e) {
            e.printStackTrace();
            //To change body of catch statement use
            //File | Settings | File Templates.
        } catch (ClassNotFoundException e) {
            e.printStackTrace();
            //To change body of catch statement use
            // File | Settings | File Templates.
        }
    }
    return myConnection;
}
public String getName() {
    return myName;
}
private static String getNextId() {
    return ""+myUniqueID++;
}
public void init(Hashtable initParams) throws IOException {
    if (!initialised){
        if (initParams.containsKey("NamedInstance")){
            myName=(String)initParams.get("NamedInstance");
        }
        else {
            myName=myName+" "+getNextId();
            fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
                "Default Instance Requested ");
        }
        if (initParams.get("DBUser")!=null){

```

```

        myDBUser=(String)initParams.get("DBUser");
    }
    if (initParams.get("DBPassword")!=null){
        myDBPassword=(String)initParams.get("DBPassword");
    }
    if (initParams.get("JDBCURL")!=null){
        jdbcurl=(String)initParams.get("JDBCURL");
    }
    initialised=true;
}
}
private DBUser createDBUserInstance(String username, String description,
    String password, int permissions, int home){
    DBUser someuser=new DBUser();
    someuser.setAuthenticator(this);
    if (home==1){ //MyCompany/MyApp Users
        someuser.reload(username,description,password, MYAPP_GROUP);
        //Set home
        someuser.setHomeId(MYAPP_GROUP.getName());
        //Group association logic
        someuser.setGroup(MYAPP_GROUP);
        //Set outer group
        someuser.setGroup(MYCOMPANY_GROUP);
        MYCOMPANY_GROUP.addUser(someuser);
        //Set inner group
        MYAPP_GROUP.addUser(someuser);
        //Permissions association logic
        switch (permissions){
            case 1:
                someuser.setPermission(CLIENT_PERMISSION);
                break;
            case 10:
                someuser.setPermission(ADMIN_PERMISSION);
                break;
            default:{
                someuser.setPermission(GUEST_PERMISSION);
                break;
            }
        }
    }
}
else {
    fAuthenticator.logAuthenticatorMessage("WARNING: User "+username+
        " has a home value of "+home+". User will be ignored!");
    return null;
}
return someuser;
}
private DBUser LoadUserFromDatabase(String username){
    DBUser someuser=null;
    Connection conn =null;
    java.sql.Statement stmt = null;
    java.sql.ResultSet rset=null;
    String name=null;
    try{
        conn=getDBConnection();
        stmt = conn.createStatement();
        rset = stmt.executeQuery(
            "select name,password, rights,home from USERS where name='"+
            username.toLowerCase()+"'");
        while(rset.next())
        {
            int permissions= rset.getInt("RIGHTS");
            int home = rset.getInt("HOME"); // home desk association

```



```

        name=rset.getString("name").toLowerCase();
        String password=rset.getString("password");
        if (password==null || password.trim().length()==0 ||
            password.equals("null")) password="nopassword";
        String description="A "+name+" user";
        //safeguard for users without a description!
        someuser=createDBUserInstance(name,description,password,
            permissions,home);
        if (someuser==null) continue;
        //In case we have invalid data to create this user object
        //Cache instance
        getUsersCollection().put(someuser.getName(),someuser);
    }
    rset.close();
    rset=null;
    stmt.close();
    stmt=null;
}
catch (Throwable t){
    logAuthenticatorException(
        "DBAuthenticator: Error obtaining details for user "+name);
    logAuthenticatorException(t);
    t.printStackTrace();
}
finally {
    if (rset!=null){
        try {
            rset.close();
        } catch (SQLException e) {
        }
        rset=null;
    }
    if (stmt!=null){
        try {
            stmt.close();
        } catch (SQLException e) {
        }
        stmt=null;
    }
}
return someuser;
}
private void LoadUsersFromDatabase(){
    Connection conn =null;
    java.sql.Statement stmt=null;
    java.sql.ResultSet rset=null;
    String name=null;
    try{
        conn=getDBConnection();
        stmt = conn.createStatement();
        rset = stmt.executeQuery(
            "select name,password, rights,home from USERS order by name" );
        while(rset.next())
        {
            int rights= rset.getInt("RIGHTS");
            int home = rset.getInt("HOME"); // home desk association
            name=rset.getString("name").toLowerCase();
            String password=rset.getString("password");
            if (password==null || password.trim().length()==0 ||
                password.equals("null")) password="nopassword";
            String description="A "+name+" user";
            DBUser someuser=createDBUserInstance(name,description,password,
                rights,home);

```

```

        if (someuser==null) continue;
        getUsersCollection().put(someuser.getName(),someuser);
    }
    rset.close();
    rset=null;
    stmt.close();
    stmt=null;
}
catch (Throwable t){
    logAuthenticatorException("Error obtaining details for user "+name);
    logAuthenticatorException(t);
    t.printStackTrace();
}
finally {
    if (rset!=null){
        try {
            rset.close();
        } catch (SQLException e) {
        }
        rset=null;
    }
    if (stmt!=null){
        try {
            stmt.close();
        } catch (SQLException e) {
        }
        stmt=null;
    }
}
}

public void reload() throws IOException {
    LoadUsersFromDatabase();
    fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
        "Reload called");
}

/**
 * Creates a new fPermission with the unique ID and name supplied.
 *
 * The implementation should save the new permission to the relevant
 * technology used.
 *
 * @param permNo Unique ID from 0 to 63.
 * @param name Name describing this new permission.
 * @return the new fPermission.
 * @throws java.io.IOException If unable to create the new fPermission.
 */
public fPermission addPermission(int permNo, String name) throws
    IOException {
    if(getPermissionsCollection().get(""+permNo) == null){
        fPermission perm = new fPermission();
        perm.reload(permNo,name);
        getPermissionsCollection().put(""+permNo, perm);
        return perm;
    }
    fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
        "Added Permission "+name+"("+permNo+")");
    return (fPermission) super.getPermissionsCollection().get(""+permNo);
}

public fUser addUser(String username, String description,
    String plainPassword, String groupName) throws IOException {
    fGroup group = null;
    if (groupName != null) {
        group = (fGroup) getGroupsCollection().get(groupName);
    }
}

```

```

        if (group == null) throw new IOException("No known group " + groupName);
    }
    fUser user = createUser(username, description, plainPassword, group);
    getUsersCollection().put(user.getName(), user);
    if(group != null) group.addUser(user);
    fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+"Added User "+
        username+" NOTE: This is not currently persisted in the database!");
    return user;
}
/**
 * Creates a new fUser with the supplied values.
 * The password field is passed as plain text but it is up to the
 * implementation to ensure the password
 * is secure.
 *
 * The implementation should save the new user to the relevant
 * technology used.
 *
 * @param user The user to copy.
 * @return The new fUser created.
 * @throws java.io.IOException If there where any errors during the
 * construction of the user.
 */
public fUser copyUser(fUser user) throws IOException {
    fGroup group = null;
    group = (fGroup) getGroupsCollection().get(user.getGroup().getName());
    fUser aUser = createUser(user.getName(),user.getDescription(),
        user.getPassword(),user.getGroup());
    getUsersCollection().put(aUser.getName(), aUser);
    if(group != null) group.addUser(aUser);
    fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+"
        "Copied User "+user.getName());
    return aUser;
}
/**
 * Adds a new group with the supplied values.
 *
 * The implementation should save the new group to the relevant
 * technology used.
 *
 * @param id Unique ID for the group.
 * @param name Name of the new group.
 * @param description Description of the new group.
 * @return The new fGroup object.
 * @throws java.io.IOException If unable to create the new fGroup object.
 */
public fGroup addGroup(int id, String name, String description) throws
    IOException {
    fGroup group = new fGroup();
    group.reload(id, name, description);
    addGroup(group);
    fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+"
        "Added Group "+group.getName());
    return group;
}
/**
 * Returns the permission with the ID supplied or null if not found.
 *
 * @param id fPermission Id to search for.
 * @return the fPermission or null if not found.
 */
public fPermission getPermission(int id) {
    Enumeration perms = getPermissionsCollection().elements();

```

```

        while (perms.hasMoreElements()) {
            fPermission fPermission = (fPermission) perms.nextElement();
            if (fPermission.getId() == id) return fPermission;
        }
        return null;
    }
    /**
     * Returns the permission with the name supplied or null if not found.
     *
     * @param name fPermission name to search for.
     * @return the fPermission or null if not found.
     */
    public fPermission getPermission(String name) {
        return (fPermission) getPermissionsCollection().get(name);
    }
    public Enumeration getUsers(){
        return getUsersCollection().elements();
    }
    public fUser getUser(String username) {
        return (fUser) LoadUserFromDatabase(username.toLowerCase());
    }
    public fGroup getGroup(String name) {
        return (fGroup) getGroupsCollection().get(name);
    }
    /**
     * Removes the permission with the ID supplied.
     *
     * The implementation should remove the permission from the relevant
     * technology used.
     *
     * @param id of the permission to delete.
     * @throws java.io.IOException if unable to delete the permission.
     */
    public void delPermission(int id) throws IOException {
        getPermissionsCollection().remove(""+id);
        fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
            "Deleted permission (" +id+")");
    }
    /**
     * Removes the user supplied.
     *
     * The implementation should remove the user from the relevant
     * technology used.
     *
     * @param user fUser object to remove.
     * @throws java.io.IOException If unable to remove the user.
     */
    public void delUser(fUser user) throws IOException {
        if (user.getGroup() != null) {
            user.getGroup().delUser(user);
        }
        getUsersCollection().remove(user.getName());
        fAuthenticator.logAuthenticatorMessage("{ "+getName()+" } "+
            "Deleted User " +user.getName());
    }
    /**
     * Removes the supplied fGroup object.
     *
     * Any user currently a member of this group will have the group reset
     * to null meaning no group membership.
     *
     * The implementation should remove the group from the relevant
     * technology used.

```

```

    *
    * @param group Group to remove.
    * @throws java.io.IOException If unable to remove the group.
    */
    public void delGroup(fGroup group) throws IOException {
        Enumeration enm = group.getUsers();
        while (enm.hasMoreElements()) {
            fUser user = (fUser) enm.nextElement();
            user.delGroup(group);
        }
        getGroupsCollection().remove(group.getName());
        fAuthenticator.logAuthenticatorMessage("{}"+getName()+" "+
            "Deleted Group "+group.getName());
    }
    /**
    * Requests that the implementation save the current state.
    *
    * This should include all users, groups and permissions.
    *
    * @throws java.io.IOException if the save failed.
    */
    public void saveState() throws IOException {
        //TODO: Implement saving of data to the database
    }
    public void roll() throws IOException {
        //TODO: Implement any log file rolling
    }
    protected fUser createUser(String name, String desc, String password,
        fGroup group){
        // System.out.println("CreateUser being called");
        DBUser usr = new DBUser();
        usr.reload(name, desc, password, group);
        usr.setHomeId(group.getName());
        usr.setAuthenticator(this);
        return usr;
    }
}

```

DBUser

```

package com.myapp;
import com.pcbsys.foundation.authentication.fUser;
import com.pcbsys.foundation.authentication.fGroup;
import com.pcbsys.foundation.authentication.fAuthenticator;
import java.util.Hashtable;
import java.sql.SQLException;
public class DBUser extends fUser {
    private static fAuthenticator myAuthenticator;
    protected DBUser(){
        super();
        super.setGroupHash(new Hashtable());
    }
    //Allow setting a reference to the authenticator instance so that we can
    // obtain its DB connection for
    // user authentication purposes
    public static void setAuthenticator (fAuthenticator authenticator){
        myAuthenticator=authenticator;
    }
    protected DBUser(String name, String desc, String password){
        this(name, desc, password, null);
    }
    protected DBUser(String name, String desc, String password, fGroup group){
        super(name,desc,password,group);
    }
}

```

```

    }
    public String login(byte[] password, boolean requestToken){
        return login(password,requestToken,null);
    }
    public String login(String password, boolean requestToken){
        return login(password,requestToken,null);
    }
    public String login(byte[] password, boolean requestToken, Hashtable params){
        return login(new String(password), requestToken, params);
    }
    public String login(String password, boolean requestToken, Hashtable params){
        java.sql.Connection conn =null;
        java.sql.Statement stmt=null;
        java.sql.ResultSet rset=null;
        String name=null;
        try{
            conn=((DBAuthenticator)myAuthenticator).getDBConnection();
            stmt = conn.createStatement();
            rset = stmt.executeQuery(
                "select password, rights from USERS where name='"+
                this.getName()+"' " );
            while(rset.next())
            {
                int rights= rset.getInt("RIGHTS");
                String thepassword=rset.getString("password");
                if (thepassword.equals(password)){
                    if (rights > 0) return "true";
                }
            }
            rset.close();
            rset=null;
            stmt.close();
            stmt=null;
        }
        catch (Throwable t){
            myAuthenticator.logAuthenticatorException(
                "DBAuthenticator: Error obtaining details for user "+name);
            myAuthenticator.logAuthenticatorException(t);
        }
        finally {
            if (rset!=null){
                try {
                    rset.close();
                } catch (SQLException e) {
                }
                rset=null;
            }
            if (stmt!=null){
                try {
                    stmt.close();
                } catch (SQLException e) {
                }
                stmt=null;
            }
        }
        return null;
    }
}

```

Web Application Single Sign On

Single sign-on (SSO) is a method of access control that enables a user to log in once and gain access to the required application and its resources without being prompted to log in again. When developing multi node Universal Messaging web applications, you have to take into consideration that incidents such a network failure, could cause the user's browser to fail over to a different node than the one initially authenticated with. In order to prevent the application user from authenticating again, or to integrate your Universal Messaging web application to a 3d party authentication mechanism and provide alternative authentication user interfaces, you can use Single Sign On Interceptors (SSI).

A Single Sign On Interceptor (SSI) is a class that conforms to a specific interface and gets invoked by the Universal Messaging realm prior to authentication in order to decide which of the following 3 outcomes should occur:

- If the user meets the criteria required, allow access to the plugin content as if they where normally authenticated, optionally generating a unique session id.
- If the user does not meet the criteria required, but a redirection is configured, redirect their browsers to the specified URL in order to authenticate.
- If the user does not meet the criteria required, and no redirection is configured, then fall back to the regular authenticator configured.

The interface for creation of a nirvana SSI implementations is defined in the following 2 classes of the com.pcbsys.foundation.authentication package:

fSSIInterceptor

```
//Return an fSSUser with a null username to fall back to authenticator
// (or redirect if a URL is set)
public abstract fSSUser getSSUser(Hashtable httpHeaders,
    Hashtable urlParameters);
public abstract void setParameters(Hashtable params);
public abstract void clear();
public abstract String getName();
```

fSSUser

```
public fSSUser(String username, String redirectURL);
public fSSUser(String username, String redirectURL, String token);
public String getUsername();
public String getRedirectURL();
public String getToken();
```

Plugin Single Sign On Interceptor Parameters

- SSIInterceptor: Fully qualified class name of SSI to use. If not specified, no interceptor will be used
- SSOAppendToken: Setting this parameter to true instructs the SSI object to generate and return a unique session ID when an affirmative single sign on decision is reached. Please note that the absence of a session ID is irrelevant to the single sign on decision.

Common Single Sign On AuthParameters

Irrespective of the Single Sign On implementation you use in your Universal Messaging server plugins, there are some AuthParameters that are also used by the server. These are:

- `SSONamedInstance`: This optional parameter requests that this SSI object is bound to the specified named instance which will be shared across all plugins on this server that are configured to do so. Please note that the first plugin that accepts a connection will bind the name to the server together with the remaining configuration parameters. For this reason please make sure that configuration is always the same on all plugins that share the same SSI instance.
- `REDIRECT_URL`: This optional parameter specifies the URL that a web client should be redirected to should the interceptor's criteria are not met. This allows the creation of alternative authentication methods such as form based authentication or others.

Mobile Application Authentication

When developing Universal Messaging based mobile applications, authentication is dependent on your mobile technology of choice. This is because the Universal Messaging Mobile APIs work exactly like any Universal Messaging enterprise or web API (with the exception of SSL client certificates).

Using SASL

Overview

The entire set of session creation methods of the Universal Messaging client and admin APIs for Java (`nsp/nsps/nhp/nhps`, native and JMS) have overloaded variants that accept username/password credentials which are then supplied to the Universal Messaging server.

The Universal Messaging server enables those credentials to be authenticated against pluggable Directory backends (ranging from LDAP to flat files) or by means of JAAS-based methods (based on user-configurable pluggable modules).

Note that authentication does not supplant the traditional Universal Messaging ACLs and is merely an additional security step performed before the relevant ACLs are evaluated and applied.

The exchange of user credentials can be performed by means of either SASL (embedded within the Universal Messaging SDK) or JAAS (controlled by user-configurable pluggable modules).

The configuration is determined by a set of Java system properties on both the client and server side, the latter typically centralized in the `nserver.conf` or `nserverdaemon.conf` configuration file. `nserver.conf` is used if you start the realm server manually, e.g. from the command line. `nserverdaemon.conf` is used if you start the realm server as a Windows service or as a UNIX daemon.

Note: Basic authentication, supplying a username and password, is only supported for Java and .NET clients.

Client-side Authentication

If the pre-existing session connection methods with no username/password parameters are used, then the client will continue to use unauthenticated sessions as before (assuming the server is configured to allow that), i.e. by defaulting the user identity to the username under whose identity the client process is running (as reflected in the Java `user.name` system property for Java clients).

The client API is controlled by three main environment variables/system properties:

- *Nirvana.auth.client.jaaskey*

If set, this means that any authentication should be performed via JAAS and its value specifies the name of the entry to use in the JAAS login configuration file. JAAS is an established Java standard which is beyond the scope of the UM documentation.

The login configuration file's pathname is specified by the usual JAAS system property, `java.security.auth.login.config`.

The Universal Messaging client SDK supplies the username and password to the JDK's built-in JAAS subsystem, and the JAAS login config file specifies one or more pluggable JAAS modules that will perform the authentication. The precise modules to use are a matter of site-specific policies determined by the Nirvana administrators, and the JAAS modules configured into a client should obviously be aligned with those configured on the server.

If `Nirvana.auth.client.jaaskey` is not set, then the authentication mechanism defaults to SASL.

- *Nirvana.sasl.client.mech*

This specifies which SASL mechanism to use, and the supported options are PLAIN, CRAM-MD5 and DIGEST-MD5.

The mechanism defaults to PLAIN if this system property is not set, and the usual SASL trade-offs apply. PLAIN transmits the user password in plain text, so it is advisable to only use it over an SSL connection. On the other hand, CRAM-MD5 and DIGEST-MD5 do not transmit the password in plain text so are more appropriate for general connections.

- *Nirvana.sasl.client.enablePrehash*

This specifies whether to prehash the supplied password when using the CRAM-MD5 or DIGEST-MD5 mechanisms. It may be set to "true" or "false". This should be set to "true" only when the server is using the `fSAGInternalUserRepositoryAdapter` to store client credentials, otherwise CRAM-MD5 and DIGEST-MD5 authentication will fail. If `Nirvana.sasl.client.enablePrehash` is not set, then the value defaults to "false" and prehashing is not enabled.

One of the JAAS modules available is the Nirvana class, `com.pcbsys.foundation.security.sasl.fSaslClientLoginModule`, which will result in the authentication being performed via SASL after all, despite initially being routed via JAAS. From the server's perspective, the authentication negotiation is conducted entirely in SASL. The `fSaslClientLoginModule` class is integrated with the Software AG family of JAAS modules, and one reason you might opt for this JAAS-SASL hybrid is to chain it with other SAG modules.

Note: Basic authentication, supplying a username and password, is only supported for Java and .NET clients.

Server-side Authentication

There is a much broader range of configuration options on the server, controlling every aspect of authentication from whether it's enabled in the first place, to how to look up user credentials.

Client Negotiation

Authentication is disabled by default on the server for backward compatibility, meaning that even if clients do supply user credentials, they will be accepted without verification. This is controlled by the `Nirvana.auth.enabled` system property, which must be explicitly set to "Y" or "y" to enable authentication. System admins have to set up various other config options when enabling authentication, so they would set `Nirvana.auth.enabled` as part of that effort. Even when authentication is enabled, authenticating clients can exist side-by-side with non-authenticating ones, meaning it is entirely optional for clients to supply any user credentials, and if they don't they will be handled in the traditional Universal Messaging manner.

The `Nirvana.auth.mandatory` system property controls this behaviour, and should be explicitly set to "Y" or "y" to make authentication mandatory, meaning clients that don't supply a username and password will be rejected (with the exception of the super-user on localhost, so that Enterprise Manager doesn't get locked out). When a client does authenticate, the UM client-server protocol automatically signals the server whether they're using SASL or JAAS and if JAAS, then the `Nirvana.auth.server.jaaskey` system property must be set on the server, and it specifies the name of the entry to use in the JAAS login configuration file. As in the client case, the pathname of this file is specified by the standard JAAS `java.security.auth.login.config` property. If the `Nirvana.auth.server.jaaskey` system property is not set, then all attempts to authenticate via JAAS will be rejected.

Directory Backend

The UM server can make use of a variety of backend Directory servers or mechanisms, as controlled by the `Nirvana.directory.provider` system property, which specifies the pluggable Java class representing the Directory.

Username are case-sensitive and are used in the form supplied to do the Directory lookup. This is the authentication step, and is followed by an authorization step in which the username is normalized to lowercase to match against Nirvana ACLs.

Nirvana ACLs are case-insensitive but expressed in lower-case and any ACLs created via the Enterprise Manager will be forced to lower case.

Internal User Repository

If the `Nirvana.directory.provider` system property is set to `com.pcbsys.foundation.security.auth.fSAGInternalUserRepositoryAdapter`, then usernames will be looked up in a standard Software AG store called the 'Internal User Repository', which is a flat file maintained by the SAG command-line utility `internaluserrepo.sh` (found in `INSTALLROOT/common/bin`).

This mechanism is the default user repository if the `Nirvana.directory.provider` property is not set.

The location of the file containing the user repository is given by the system property, `Nirvana.auth.sagrepo.path`, and would default to `./users.txt` (relative to runtime directory of UM server), but the `nserver.conf` shipped with UM overrides this as `./users.txt`, locating it in the same `INSTALLROOT/nirvana/server/umserver` directory as the `licence.xml` file. The `nserver.conf` file may of course be edited as usual to move the `users.txt` file into a location that is shared by all the realms of an installed UM instance.

For more information, see the description of the Internal User Repository in the Universal Messaging Developer Guide.

LDAP

If the `Nirvana.directory.provider` system property is set to `com.pcbsys.foundation.security.auth.fLDAPAdapter`, then LDAP will be used as the source of user information.

Interaction with the LDAP server is configured via the following Java system properties:

- `Nirvana.ldap.provider`: The LDAP client class - defaults to the JDK's built-in provider, `com.sun.jndi.ldap.LdapCtxFactory`
- `Nirvana.ldap.url`: The address of the LDAP server. This has no default and must be specified, using syntax such as `ldap://localhost:389/dc=sag,dc=com`
- `Nirvana.ldap.suffix`: The suffix to apply to LDAP queries. This has no default and may be null, but if non-null it qualifies the URL above. Eg. `Nirvana.ldap.url=ldap://localhost:389/dc=sag` and `Nirvana.ldap.suffix=dc=com` will result in the same effective query root as `Nirvana.ldap.suffix=ldap://localhost:389/dc=sag,dc=com` when the `Nirvana.ldap.suffix` property is not set.
- `Nirvana.ldap.rootcreds`: The privileged-admin login credentials to use on the LDAP server, in order to perform user queries. There is no default and if not set it means there is no need to specify any such credentials, but if present the format must be `username:password`.

The remaining system properties relate to the LDAP schema and default to the standard COSINE schema:

- `Nirvana.ldap.attribute.username`: This specifies the LDAP attribute which represents the username, and defaults to the standard schema convention of `"cn"`.

- `Nirvana.ldap.attribute.password`: This specifies the LDAP attribute which represents the password, and defaults to the standard schema convention of "userPassword".
- `Nirvana.ldap.search.username`: This specifies the search expression to use for a given username, and defaults to `cn=%U%`, where `%U%` gets substituted by the username.

Server JAAS Authentication with Software AG Security infrastructure component

Universal Messaging can use the Software AG Security Infrastructure component (SIN) to provide server JAAS authentication capabilities. The SIN component provides a variety of options for using different authentication back-ends and implementing flexible authentication scenarios.

Full details of the Software AG Security Infrastructure component are available in the Software AG documentation web site in the document *Software AG Infrastructure Administrator's Guide*. See the description of the predefined login modules in that guide for details of setting up authentication using JAAS login modules and a description of the parameters available for configuring the login modules.

Enabling JAAS authentication

The SIN modules are delivered with the Universal Messaging distribution and are available on the Universal Messaging server classpath. By default the server is configured to use a JAAS context named `UM-Default` which is defined in a file `jaas.conf`, that resides under the server `bin` directory. The server configuration file `Server_Common.conf` (again under the server `bin` directory) defines the path to the JAAS configuration file as the value of the system property `java.security.auth.login.config`. The name of the JAAS context to be used for authentication is configured through the system property `Nirvana.auth.server.jaaskey`. Note that when this property is present, the server will use JAAS authentication, overriding any configured legacy directory-based authenticators. Finally, in order to enable the server authentication one needs to switch the flag `Nirvana.auth.enabled` to `Y` as well as modify the flag `Nirvana.auth.mandatory` to `Y` in order to make authentication mandatory.

The default values that come with a Universal Messaging server instance are:

```
-DNirvana.auth.enabled=N
-DNirvana.auth.mandatory=N
-DNirvana.auth.server.jaaskey=UM-Default
-Djava.security.auth.login.config=jaas.conf
```

The `UM-Default` JAAS context is pre-configured to use the Software AG default internal user repository as an authentication back-end.

SIN authentication templates for JAAS

Universal Messaging installation delivers a JAAS configuration file that comes with a pre-configured `UM-Default` JAAS context and templates for a few other authentication scenarios.

Besides using the Software AG internal user repository, SIN offers a variety of options for LDAP authentication:

1. LDAP Lookup

The LDAP look up method involves logging on to the LDAP server either anonymously (if allowed) or by using a predefined principal, and performing a lookup for a user entry matching a certain attribute. Finally the SIN-specific LDAP login module will attempt to bind with the retrieved user's distinguished name and the supplied password. The following template configuration is provided where you should fill in the required properties:

```
/*
 * Authentication with an LDAP server by looking up users based on an attribute.
 * The login module will use the configured principal (parameter "prin") to bind
 * and perform a lookup.
 * The principal can be omitted if the LDAP server is configured to allow
 * anonymous bind and lookup.
 * Afterwards the login module will search for an entry under
 * the specified root DN with the attribute
 * configured through the "uidprop" matching the supplied username.
 * The "uidprop" parameter can be omitted, its default value is "cn".
 */
UM-LDAP-Lookup {
    com.softwareag.security.sin.is.ldap.lm.LDAPLoginModule sufficient
        url="<LDAP URL>"
        prin="<principal complete DN>"
        cred="<principal password>"
        userrootdn="<root DN for searching users>"
        uidprop="<user ID property name>";
};
```

2. Active Directory direct bind

With active directory the SIN-specific LDAP login module offers the possibility for a direct bind with the supplied username and password, using the following configuration template where the user should fill in the active directory URL.

```
/*
 * Authentication with an Active Directory server by binding directly with
 * the supplied credentials.
 * Active Directory allows for a direct bind with a principal in the
 * format "domainPrefix\user".
 */
UM-AD-DirectBind {
    com.softwareag.security.sin.is.ldap.lm.LDAPLoginModule sufficient
        url="<LDAP URL>"
        noPrinIsAnonymous=false
        useFQDNForAuth=true;
};
```

3. LDAP direct bind

It is also possible to attempt a direct bind against an LDAP server where the bind DN will be constructed by the SIN-specific JAAS login module using a prefix and a suffix, with the supplied username in-between. In this case the user is required to configure the URL of the LDAP server, as well as the prefix and the suffix to be used for the construction of the bind DN:

```
/*
 * Authentication with an LDAP server by building a complete DN and binding directly.
 * The bind DN will be composed by prepending the supplied username with the
 * value of the "dnprefix" parameter and appending the value
 * of the "dnsuffix" parameter.
 */
UM-LDAP-DirectBind-Affixes {
    com.softwareag.security.sin.is.ldap.lm.LDAPLoginModule sufficient
        url="<LDAP URL>"
        noPrinIsAnonymous=false
        useaf=true
        dnprefix="<Bind DN prefix>"
        dnsuffix="<Bind DN suffix>";
};
```

4. Combination of any of the above

JAAS is extremely flexible and allows for using a combination of different back-ends, such as, several LDAP servers and/or repositories. For example it is possible to configure LDAP authentication with an internal user repository as a fallback mechanism:

```
/*
 * Authentication with an LDAP server through a user lookup, with an
 * internal user repository as a fallback mechanism.
 * This configuration would first try to log in against the LDAP if that fails
 * would try authentication against the internal user repository.
 * Similarly one could use several LDAP servers or
 * combine more authentication backends.
 */
UM-Combined {
    com.softwareag.security.sin.is.ldap.lm.LDAPLoginModule sufficient
        url="<LDAP URL>"
        prin="<principal complete DN>"
        cred="<principal password>"
        userrootdn="<root DN for searching users>"
        uidprop="<user ID property name>";
    com.softwareag.security.jaas.login.internal.InternalLoginModule sufficient
        template_section=INTERNAL
        internalRepository="../../../../common/conf/users.txt";
};
```

To use any custom JAAS context, or any of the predefined templates, you would need to modify the value of the `Nirvana.auth.server.jaaskey` system property inside the `Server_Common.conf` file in the server bin directory, pointing it to the context to be used. At the moment the context name is globally defined for the entire server.

Note: Since the `jaas.conf` file is edited manually and requires certain knowledge of the default JAAS configuration file format, there is always the possibility of a syntactic error. Such errors during parsing of the file will be printed to the `UMRealmService.log` file which resides under the server bin directory.

Converting a .jks Key Store to a .pem Key Store

In order to convert a Java key store into a Privacy Enhanced Mail Certificate, you will need to use two tools :

1. keytool.exe - to import the keystore from JKS to PKCS12 (supplied with Java)
2. openssl.exe - to convert the PCKS12 to PEM (supplied with OpenSSL)

Neither keytool or openssl can be used to convert a jks directly into a pem. First we must use keytool to convert the JKS into PKCS:

```
keytool -importkeystore -srckeystore client.jks -destkeystore client.pkcs
-srcstoretype JKS -deststoretype PKCS12
```

You will be prompted to enter passwords for the key stores when each of these programs are run. The password used for the keystores created using the generator is "nirvana". Next you need to use openssl.exe to convert the PKCS into PEM.

```
openssl pkcs12 -in client.pkcs -out client.pem
```

Repeat the above code for the any other jks key stores. After this you will have the required key stores in pem format.

Access Control Lists

Security Policies

Universal Messaging offers complete control over security policies. Universal Messaging can either store security policies locally or be driven by an external entitlements service.

Universal Messaging's rich set of entitlements ensure that everything from a network connection through to topic and queue creation can be controlled on a per user basis.

Every component of a Universal Messaging server has a set of entitlements associated with it. These entitlements can be set programmatically or through the Universal Messaging Enterprise Manager.

For more information on the components that entitlements can be set against please refer to the Universal Messaging ACL Guide (see ["Access Control Lists \(ACLs\)" on page 95](#)).

Access Control Lists (ACLs)

Note: ACLs provide authorization for your application, but authentication is also required for robust security. To achieve robust security you should use either LDAP or SSL client-side certificate authentication in addition to ACLs. If authentication is a concern, refer to the LDAP description in the section

["Using SASL " on page 88](#), or the SSL client certificate authentication description in the section ["Using SSL" on page 98](#).

Universal Messaging's Access Control List (ACL) controls client connection requests and subsequent Universal Messaging operations. By default, access control checks are performed within a realm.

The Universal Messaging Administration API exposes the complete security model of the Universal Messaging Realm Server, remotely allowing customer specific security models to be created. This means that it is easy to integrate Universal Messaging into an existing authentication and entitlement service.

It is also possible to manage Universal Messaging ACLs using the Enterprise Manager GUI.

The Universal Messaging realm has an ACL associated with it. The ACL contains a list of subjects and the operations that each subject can perform on the realm.

Users are given entitlements based on their subject. A subject is made up of a username and a host.

The username part of the subject is the name of the user taken from either the operating system of the machine they are connecting from or the certificate name if they are using an SSL protocol.

The host part of the subject is either the IP address or the hostname of the machine they are connecting from.

The subject takes the form of :

```
username@host
```

For example:

```
johnsmith@192.168.1.2
```

Each channel, queue and service also has an associated ACL that defines subjects and the operations the subjects can perform.

A subject corresponds to the user information for a realm connection

Each type of ACL entry has a number of flags that can be set to true or false in order to specify whether the subject can or can't perform the operation.

General ACL permissions

The following flags apply to every ACL:

- Modify - Allows the subject to add/remove ACL entries
- List - Allows the subject to get a list of ACL entries
- Full Privileges - Has complete access to the secured object

Universal Messaging Realm Server ACL permissions

The Realm Access Control Entry has the following controllable flags:

- Use Admin API - Can use the nAdminAPI package
- Manage Realm - Can add / remove realms from this realm
- Manage Joins - Can add/delete channel joins
- Manage Channels - Can add/delete channels on this realm
- Access The Realm - Can currently connect to this realm
- Override Connection Count - Can bypass the connection count on the realm
- Configure Realm - Can set run time parameters on the realm
- Cluster - Can perform cluster operations, such as create, delete or modify cluster information
- Manage DataGroups - Can add, delete and manage any DataGroup on the realm
- Publish to DataGroups - Can publish to any DataGroup on the realm
- Own DataGroups - Can take ownership of any DataGroup on the realm

Channel ACL permissions

The Channel Access Control Entry has the following controllable flags:

- Write - Can publish events to this channel
- Read - Can subscribe to events on this channel
- Purge - Can delete events on this channel
- Get Last EID - Can get the last event ID on this channel
- Named - Can connect using a named (durable) subscriber

Queue ACL permissions

The Queue Access Control Entry has the following controllable flags:

- Write - Can push events to this queue
- Read - Can peek the events on this queue
- Purge - Can delete events on this queue
- Pop - Can pop events from the queue

Wildcard Support

As well as being able to specify an access control entry for a specific subject, the subject itself can contain wildcards. In this way you can specify access control based on hostname or on username.

Example Wildcard ACLs :

ACL Entry	Description
@	Represents all users from all nodes
*@client1.com	Represents all users from the node client1.com
username@nodename	Represents the user "username" on the node "nodename"
username@*	Represents the user "username" on all nodes

If a user is matched by more than one wildcard ACL, the user receives the cumulative permissions of all of these wildcard ACLs. If a user is matched by a combination of wildcard ACLs and non-wildcard ACLs, the effect of the permissions is also cumulative.

The subject *@* is a special case. It is provided in all ACL objects by default, and corresponds to the default permission that all subjects inherit who connect but do not individually appear within the ACL. If the realm configuration parameter `OverrideEveryoneUser` is set to false (this is the default), then *@* will override any more specific ACL entries which match a given user. If this parameter is set to true, then *@* will be overridden by any more specific ACL entries which match a given user. `OverrideEveryoneUser` is one of the parameters in the Global Values realm configuration group; see the relevant section of the Enterprise Manager documentation for further information.

Using SSL

SSL Concepts

SSL With Client Certificate Validation

When a client requests a secure connection to a server, both the client and the server must verify each other's certificate to ensure that the source is trusted. A certificate will generally contain the source's name, the certificate authority that signed it and the public key for that source. The certificate can be authenticated by validating the signature.

■ *Certificate signed by well known certification authority*

Several certification authorities exist (such as Verisign) that will review applications for certificates and digitally sign them if they have sufficient proof that the company is what it claims to be. These authorities are trusted as default by web browsers and applications supporting SSL.

When a connection attempt is made, the certificate is checked with these well known authorities. The signature is created using the private key of the certificate authority, therefore it is known to have been encrypted by that authority only if it can be decrypted by its public key. If this is the case then it is known that the certificate is valid.

■ *Self signed certificates*

To acquire a certificate from a well known certification authority is expensive and time consuming as the company will have to prove itself by providing certain documentation. It is possible however to create and sign your own certificates. This essentially makes the company that created the certificate also act as the certification authority for that certificate. This means that the certificate will not be validated by checking with the well known authorities.

If the certificate is not validated by the well known authorities then it is checked against the user created trust store. This store contains certificates for companies that are not registered with the well known authorities but the user has chosen to trust.

Because the sample certificates created are self signed, it is important that the trust store of the client/server contains the certificates of the server/client respectively otherwise an exception will be thrown as the system cannot verify the source.

Once the certificates have been validated, the client and server will have each others public keys. When the client communicates with the server, it will encrypt the data with the server's public key. Public keys are asymmetric which means that it cannot be decrypted using the same key, instead it requires the server's private key which only the server holds. This means that only the server can read the data sent. Similarly, any response from the server will be encrypted with the client's public key.

Anonymous SSL

Universal Messaging also supports anonymous (server-side) SSL. Anonymous SSL does not validate the client. This means that the client does not need to have a certificate as it is never checked by the server. Instead, the client sends a request for the server's certificate as usual but instead of the server encrypting data sent to the client with the client's public key, a session is created. To create a session, the client generates a random number and sends this number to the server encrypted with the server's public key. Only the server can decrypt the random number. So only the server and the client know this number. The random number is used along with the server's public key to create the session.

SSL Encryption

Universal Messaging supports SSL and offers support for SSL-enabled TCP/IP sockets as well as HTTPS. When SSL is used, the subject used for entitlements can be extracted from the client's certificate CN attribute.

Universal Messaging's support for SSL provides both client side and anonymous (Server side) SSL. Different SSL certificate chains can be assigned to different Universal Messaging interfaces, each one supporting its own set of cryptographic algorithms.

There is no limit to the number of interfaces and therefore different SSL certificate chains that can be supported on a single Universal Messaging realm server.

For more information on configuring Universal Messaging interfaces to use SSL encryption with Universal Messaging please see the Enterprise Manager guide.

To learn more about SSL please see the SSL Concepts section ("[SSL Concepts](#)" on page 98).

Client SSL Configuration

This section describes how to use SSL in your Universal Messaging client applications. Universal Messaging supports various wire protocols (see "[Communication Protocols and RNAMEs](#)" on page 27) including SSL enabled sockets and HTTPS. The example programs contained in the Universal Messaging package will all work with SSL enabled on the realm server.

This guide explains client SSL configuration on Java. Guides for C++ and C# are also available.

Once you have created an SSL enabled interface for your realm (see the description of how to do this in the section "Creating an SSL network interface" of the Enterprise Manager documentation), you need to ensure that your client application passes the required SSL properties either via custom setters on the session attributes or via system properties used by your JSSE-enabled JVM. The Universal Messaging download contains some sample Java keystore files that will be used in this example.

The first such keystore is the client keystore, called `client.jks`, which can be found in your installation directory, under the `client/<umserver>/bin` directory. The second is the client truststore called `nirvanacacerts.jks`, which is also located in the `client/<umserver>/bin` directory. On Windows, these files exist once you run the '**Server/<umserver>/Create Demo SSL Certificates**' shortcut.

If you would like to add your own client certificates please see the section "How to generate certificates" of the Enterprise Manager documentation.

Custom SSL Properties

Using the sample keystores, you can use the following steps to set the SSL properties on a session.

1. Construct your session attributes object:

```
nSessionAttributes sessionAttributes = new nSessionAttributes(rname);
```

2. Set your required SSL attributes using these methods:

```
attrs.setKeystore(String keyStorePath, String keyStorePassword)
attrs.setKeystore(String keyStorePath, String keyStorePassword,
    String certificateAlias);
    // Certificate alias specifies which certificate the client can use
    // if client validation is enabled on the interface being connected to
attrs.setTruststore(String trustStorePath, String trustStorePassword);
```

So your session attributes should look like this:

```
attribs.setKeystore(%INSTALLDIR%\client\<umserver>\bin\client.jks, password);
attribs.setTruststore(%INSTALLDIR%\client\<umserver>\bin\nirvanacacerts.jks,
password);
```

In addition to this you are also able to set a list of enabled cipher strings the client can use and the SSL protocol:

```
attribs.setEnabledCiphers(String[] enabledCiphers);
attribs.setSSLProtocol(String protocol);
```

3. Construct your session with the constructed session attributes:

```
nSession session = nSessionFactory.create(attribs);
session.init();
```

JSSE SSL System Properties

Using the example keystores, the following system properties are required by the Universal Messaging sample apps and must be specified in the command line as follows:

```
-DCKEYSTORE=%INSTALLDIR%\client\<umserver>\bin\client.jks
-DCKEYSTOREPASSWD=password
-DCAKEYSTORE=%INSTALLDIR%\client\<umserver>\bin\nirvanacacerts.jks
-DCAKEYSTOREPASSWD=password
```

where:

CKEYSTORE is the client keystore location

CKEYSTOREPASSWD is the password for the client keystore

CAKEYSTORE is the CA keystore file location

CAKEYSTOREPASSWD is password for the CA keystore

The above system properties are used by the Universal Messaging sample apps, but are mapped to system properties required by a JSSE-enabled JVM by the utility program 'com.pcbsys.foundation.utils.fEnvironment', which all sample applications use. If you do not wish to use this program to perform the mapping between Universal Messaging system properties and those required by the JVM, you can specify the SSL properties directly. To do this in your own applications, the following system properties must be set:

```
-Djavax.net.ssl.keyStore=%INSTALLDIR%\client\<umserver>\bin\client.jks
-Djavax.net.ssl.keyStorePassword=password
-Djavax.net.ssl.trustStore=%INSTALLDIR%\client\<umserver>\bin\nirvanacacerts.jks
-Djavax.net.ssl.trustStorePassword=password
```

where :

javax.net.ssl.keyStore is the client keystore location

javax.net.ssl.keyStorePassword is the password for the client keystore

javax.net.ssl.trustStore is the CA keystore file location

javax.net.ssl.trustStorePassword is password for the CA keystore

As well as the above system properties, if you are intending to use https, both the Universal Messaging sample apps and your own applications will require the following system property to be passed in the command line:

```
-Djava.protocol.handler.pkgs="com.sun.net.ssl.internal.www.protocol"
```

As well as the above, the RNAME (see ["Communication Protocols and RNAMEs" on page 27](#)) used by your client application must correspond to the correct type of SSL interface, and the correct hostname and port that was configured earlier.

JMS Clients

In JMS, the RNAME corresponds to a JNDI reference. The example JMSAdmin application can be used to create a sample file based JNDI context, where the RNAME is specified as the content of the TopicConnectionFactoryFactory reference. Once your SSL interface is created you can simply change this value in your JNDI context to be the RNAME you require your JMS applications to use.

JMS Client SSL Configuration

This section describes how to use SSL in your Universal Messaging Provider for JMS applications. Universal Messaging supports various wire protocols including SSL enabled sockets and HTTPS.

Once you have created an SSL enabled interface for your realm you need to ensure that your client application passes the required SSL properties either on the connection factory or via system properties used by your JSSE-enabled JVM. The Universal Messaging download contains some sample Java keystore files that will be used in this example.

The first such keystore is the client keystore, called `client.jks`, which can be found in your installation directory, under the `/server/Universal Messaging/bin` directory. The second is the truststore called `nirvanacacerts.jks`, which is also located in the `/server/Universal Messaging/bin` directory.

Custom SSL Properties

Using the sample keystores, you can set custom SSL attributes on JMS as follows:

Setting the SSL Attributes on the JNDI Context

In your properties object the following properties will set SSL attributes on the JNDI Context.

```
env = new Properties();
env.setProperty("java.naming.factory.initial",
    "com.pcbsys.nirvana.nSpace.NirvanaContextFactory");
env.setProperty("java.naming.provider.url", rname);
env.setProperty("nirvana.ssl.keystore.path",
    "%INSTALLDIR%\client\Universal Messaging\bin\client.jks");
env.setProperty("nirvana.ssl.keystore.pass", password);
env.setProperty("nirvana.ssl.keystore.cert", certAlias);
    // Certificate alias for the client to use when connecting to an interface
    // with client validation enabled
env.setProperty("nirvana.ssl.truststore.path",
    "%INSTALLDIR%\client\Universal Messaging\bin\nirvanacacerts.jks");
env.setProperty("nirvana.ssl.truststore.pass", password);
env.setProperty("nirvana.ssl.protocol", "TLS");
```

Setting the SSL Attributes on the Connection Factory

- You can set the SSL attributes using the same Properties object like this:

```
connectionFactory.setProperties(env);
Connection con = connectionFactory.createConnection();
```

- You can set the SSL attributes using the available setters:

```
connectionFactory.setSSLStores(String keyStorePath, String keyStorePass,
    String trustStorePath, String trustStorePass);
connectionFactory.setSSLStores(String keyStorePath, String keyStorePass,
    String certificateAlias, String trustStorePath, String trustStorePass);
connectionFactory.setSSLProtocol(String protocol);
connectionFactory.setSSLEnabledCiphers(String[] enabledCiphers);
Connection con = connectionFactory.createConnection();
```

Setting the SSL Attributes on the Connection

```
Connection con = connectionFactory.createConnection(keyStorePath, keyStorePass,
    keyStoreCert, trustStorePath, trustStorePass, cipherSuite, protocol)
```

JSSE SSL System Properties

The following system properties are used by the jsse implementation in your JVM. You can specify the SSL properties by passing the following as part of the command line for your JMS application:

```
-Djavax.net.ssl.keyStore=%INSTALLDIR%\client\Universal Messaging\bin\client.jks
-Djavax.net.ssl.keyStorePassword=password
-Djavax.net.ssl.trustStore=%INSTALLDIR%\client\Universal Messaging\bin\nirvacacerts.jks
-Djavax.net.ssl.trustStorePassword=password
```

where :

- javax.net.ssl.keyStore is the client keystore location
- javax.net.ssl.keyStorePassword is the password for the client keystore
- javax.net.ssl.trustStore is the CA keystore file location
- javax.net.ssl.trustStorePassword is password for the CA keystore

As well as the above system properties, if you are intending to use https, your JMS applications will require the following system property to be passed in the command line:

```
-Djava.protocol.handler.pkgs="com.sun.net.ssl.internal.www.protocol"
```

As well as the above, the RNAME used by the JMS application must correspond to the correct type of SSL interface, and the correct hostname and port that was configured earlier.

In JMS, the RNAME corresponds to a JNDI reference. The example JMSAdmin application can be used to create a sample file based JNDI context, where the RNAME is specified as the content of the TopicConnectionFactoryFactory reference. Once your SSL interface is created you can simply change this value in your JNDI context to be the RNAME you require your JMS applications to use.

Setting up and using FIPS

Federal Information Processing Standards (FIPS) provides standards for information processing for use within the Federal government.

Many government and financial organizations require their software to be FIPS 140-2 compliant, which follows the current standards and guidelines for cryptographic information processing.

In Universal Messaging, FIPS capabilities can be enabled through the use of Mozilla NSS. NSS is an open source security library that has received FIPS validation/certification for both its FIPS and cryptographic modes. Further information can be found at the Mozilla website.

Note: Universal Messaging itself is not FIPS 140 certified.

The NSS library is not bundled in the Universal Messaging product. Details of how to download and build your own copy of NSS for your respective platform can be found at the Mozilla website.

Once you have built the library, the JVM will need to know where to load this library from via environment variable (for example, `LD_LIBRARY_PATH`, or equivalent). Only the Oracle JVM is supported at this time.

To enable FIPS in Universal Messaging, the property `enableFIPS` must be set on either the client and/or server before either one is started. For example:

```
-DenableFIPS=<NSS mode>
```

where `<NSS mode>` can be one of the following:

- `nssfips`
- `nsscrypto`

The biggest difference between these modes is that `nssfips` uses an NSS database for all key management purposes, while `nsscrypto` uses just the cryptographic algorithms and functions of NSS, while maintaining native JKS support. And, at the time of this writing, `nssfips` does not yet support TLS 1.2.

Storing certificates into an NSS database requires you to use certain NSS tools; documenting these tools is beyond the scope of this document (see the Mozilla website for details). However, an example shell script is provided below to get you started:

```
#!/bin/csh
setenv DYLD_LIBRARY_PATH /Users/test/Downloads/nss-3.20/dist/Darwin13.4.0_64_OPT.OBJ/lib
setenv PATH /Users/test/test/nss-3.20/dist/Darwin13.4.0_64_OPT.OBJ/bin:$PATH
echo Make nssdb directory
mkdir ~/nssdb
modutil -create -dbdir ~/nssdb
echo Enable FIPS mode on nssdb
modutil -fips true -dbdir ~/nssdb
# Password12345!
echo Set password for nssdb
```



```

modutil -changePW "NSS FIPS 140-2 Certificate DB" -dbdir ~/nssdb
# DER format/PKCS#7
echo Export alias from CA JKS to CRT
# Export only public keys
keytool -export -alias "localhost,127.0.0.1" -keystore ./nirvanacacerts.jks \
    -file public.crt
echo Import CRT into nssdb
certutil -A -d ~/nssdb -t "TCu,,TCu" -n "Test CA 1" -i ./public.crt
echo Convert client JKS to PKCS12
keytool -importkeystore -srckeystore ./client.jks -destkeystore ./client.p12 \
    -deststoretype PKCS12
echo Convert server JKS to PKCS12
keytool -importkeystore -srckeystore ./server.jks -destkeystore ./server.p12 \
    -deststoretype PKCS12
echo Import client PKCS12 into nssdb
pk12util -i ./client.p12 -n client -d ~/nssdb
echo Import server PKCS12 into nssdb
pk12util -i ./server.p12 -n CA -d ~/nssdb

```

The above example assumes self signed certificates. Certificate aliases are NOT supported when NSS is run in FIPS mode, so only 1 certificate/trusted certificate pair should be imported into an NSS database. It is recommended that you create both a client and server NSS database to separate certificates from one another. Each NSS instance/library has a single NSS database available for key management (i.e., the relationship is 1:1).

Both modes use the Oracle JVM support for NSS; specifically, this means using the Oracle PKCS11 provider to access functionality within the NSS library. Again, no other JVM is currently supported. Only 1 NSS instance/library can be loaded per JVM via the PKCS11 provider. Universal Messaging automatically manages the loading of the PKCS11 provider and NSS library for you, so you do not have to do this yourself (assuming that the LD_LIBRARY_PATH is set properly for the NSS library).

NSS uses separate configuration files to determine the mode that it should run in (among other things). For example, to run NSS in FIPS mode, create the following configuration file as a text file:

```

name = <unique name>
nssLibraryDirectory = <path to location of NSS library>
nssSecmodDirectory = <path to location of NSS database>
nssModule = fips

```

Example:

```

name = CNSS
nssLibraryDirectory = /Users/test/Downloads/nss-3.20/dist/Darwin13.4.0_64_OPT.OBJ/lib
nssSecmodDirectory = /Users/test/nssdb
nssModule = fips

```

The key "name" is important, as it signifies the NSS name which uniquely identifies the NSS instance. The rest should be self explanatory. If you set up an NSS configuration file for FIPS mode, then `-DenableFIPS` should be set to `nssfips`.

To run NSS in crypto mode, create a configuration file as follows:

```

name = <unique name>
nssLibraryDirectory = <path to location of NSS library>
nssModule = crypto
attributes = compatibility
nssDbMode = noDb

```

Example:

```
name = CNSS
nssLibraryDirectory = /Users/test/Downloads/nss-3.20/dist/Darwin13.4.0_64_OPT.OBJ/lib
nssModule = crypto
attributes = compatibility
nssDbMode = noDb
```

If you set up an NSS configuration file for crypto mode, then `-DenableFIPS` should be set to `nsscrypto`.

Note that NSS configuration files can only have 1 set of configuration parameters in them; in other words, you cannot have multiple configurations within a single file.

Once the configuration file and name have been set up, you'll need to specify them either via the Universal Messaging APIs, or via the Enterprise Manager (if configuring a Universal Messaging server NHPS or NSPS interface). For the Universal Messaging APIs, see the appropriate Java documentation for:

```
com.pcbsys.nirvana.client.nSessionAttributes.getPKCS11NSSConfigFile
com.pcbsys.nirvana.client.nSessionAttributes.setPKCS11NSSConfigFile
com.pcbsys.nirvana.client.nSessionAttributes.getPKCS11NSSName
com.pcbsys.nirvana.client.nSessionAttributes.setPKCS11NSSName
com.pcbsys.nirvana.nAdminAPI.nSSLInterface.getPKCS11NSSConfigFile
com.pcbsys.nirvana.nAdminAPI.nSSLInterface.setPKCS11NSSConfigFile
com.pcbsys.nirvana.nAdminAPI.nSSLInterface.getPKCS11NSSName
com.pcbsys.nirvana.nAdminAPI.nSSLInterface.setPKCS11NSSName
com.pcbsys.nirvana.nAdminAPI.nHTTPSInterface.getPKCS11NSSConfigFile
com.pcbsys.nirvana.nAdminAPI.nHTTPSInterface.setPKCS11NSSConfigFile
com.pcbsys.nirvana.nAdminAPI.nHTTPSInterface.getPKCS11NSSName
com.pcbsys.nirvana.nAdminAPI.nHTTPSInterface.setPKCS11NSSName
com.pcbsys.foundation.drivers.configuration.fSSLConfig
com.pcbsys.foundation.drivers.configuration.fHTTPSConfig
```

While running NSS in FIPS mode, the password to the NSS database must be specified via `nSessionAttributes.setKeystore` (keystore/alias would be set to null or is unused), or via `nSSLInterface/nHTTPSInterface.setKeyStorePassword`. While running NSS in crypto mode, use the APIs as you normally would for JKS keystores/truststores. Set the PKCS11 NSS configuration file and name using the appropriate `nSessionAttributes/nSSLInterface/nHTTPSInterface.setPKCS11NSS*` methods.

For configuring an SSL/HTTPS interface in the Enterprise Manager, specify a key store password, PKCS11 NSS configuration file and name for NSS FIPS mode; all other fields can be left blank. Any combination of SSL and/or HTTPS interfaces configured while in NSS FIPS mode will share the same NSS database (since only 1 NSS instance/database is available per JVM). For NSS crypto, configure the interfaces as you normally would for JKS keystores/truststores.

To recap, refer to the following NSS checklist when enabling FIPS for UM:

1. Shut down any Universal Messaging servers and/or clients that you want to configure for FIPS.
2. Download the Mozilla NSS open source library, and build it for your respective platform (see the Mozilla website for details).
3. Determine whether you want to run in NSS FIPS or crypto mode.

4. If you'll be running in FIPS mode, create an NSS FIPS configuration file. If you'll be running in crypto mode, create an NSS crypto configuration file.
5. If you'll be running in FIPS mode, create an NSS database, and store the appropriate certificates in it (both client and server). If you'll be running in crypto mode, create/use normal JKS keystores/truststores.
6. Set up LD_LIBRARY_PATH (or equivalent, depending on your platform) to point to the location of the NSS binary distribution. This can be done in either a shell, command line, shell/batch script or configuration file.
7. Write any API programs that are FIPS enabled (if applicable).
8. For client and/or server, set `-DenableFIPS=nssfips` in a command line, shell/batch script or configuration file if you'll be running in FIPS mode. Set `-DenableFIPS=nsscrypto` if you'll be running in crypto mode (again, for client and/or server).
9. Restart the Universal Messaging server(s), and configure SSL/HTTPS interfaces with FIPS as needed. Restart the Universal Messaging client(s).

5 Performance, Scalability and Resilience

- Overview of Performance, Scalability and Resilience 110
- Clustered Server Concepts 111
- Multicast: An Overview 140
- Shared Memory (SHM) 142
- Scalability 142

Overview of Performance, Scalability and Resilience

Performance, Scalability and Resilience are design themes that are followed in all areas of Universal Messaging's design and architecture. Specific implementation features have been introduced into server and client components to ensure that these themes remain constantly adhered to.

Performance

Universal Messaging is capable of meeting the toughest of Low Latency and High Throughput demands. Universal Messaging's server design, sophisticated threading models and heavily optimized IO subsystem ensures peak performance. Multicast and shared memory communication modes allow Universal Messaging to consistently achieve low microsecond latencies. Universal Messaging is constantly being benchmarked by clients and other 3rd parties and continues to come out on top.

The benchmarking section provides detailed information on performance in a variety of scenarios. Information on performance tuning is also available on the website, helping clients achieve optimal performance in their deployments.

Scalability

Scalability in terms of messaging middleware typically means *connection scalability*, which is the ability to support large numbers of concurrent connections; this is something Universal Messaging does out of the box. However in defining truly global enterprise applications a single system often needs to scale across more than one processing core, often in more than one geographic location.

Universal Messaging servers can be configured in a variety of ways to suit scalability (and resilience) requirements. Multiple Universal Messaging servers can exist in a single federated name space. This means that although specific resources can be put onto specific Universal Messaging realm servers any number of resources can be managed and access centrally from a single entry point into Universal Messaging's federated namespace. In addition to high availability and resilience features Universal Messaging Clusters also offer a convenient way to replicate data and resources among a number of realm servers thus supporting higher rates of concurrent connections.

Resilience

Business contingency and disaster recovery planning demand maximum availability from messaging Middleware systems. Universal Messaging provides a number of server and client features to ensure data always remains readily accessible and that outages are transparent to clients.

Universal Messaging Clusters replicate all resources between realms. Channels, Topics, Queues and the data held within each one is always accessible from any realm server in the cluster.

Universal Messaging clients are given a list of Universal Messaging realms in the cluster and automatically move from one to another if a problem or outage occurs.

Clustered Server Concepts

Clusters: An Overview

A Universal Messaging Cluster is a collection of Universal Messaging Realms (servers) that contain common messaging resources such as channels/topics or queues. Each clustered resource exists in every Realm within the cluster. Whenever the state of a clustered resource changes, the state change is updated on all realms in the cluster.

Clustering also offers a convenient way to replicate content between servers and ultimately offers a way to split large numbers of clients over different servers in different physical locations.

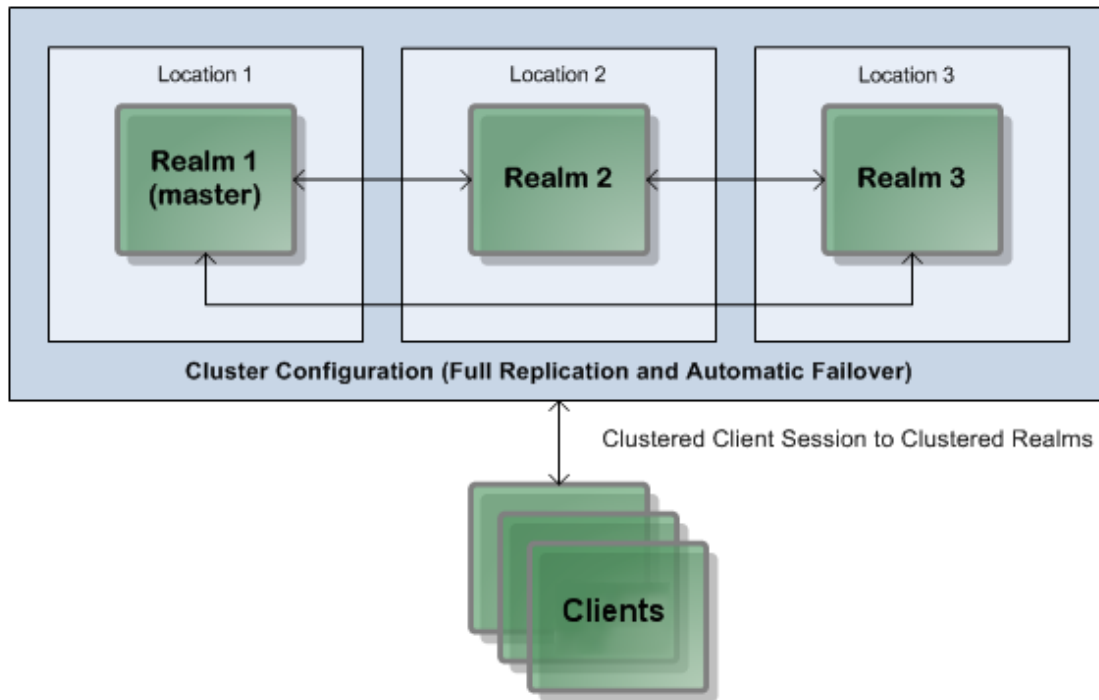
Note: Clustering and Realm Federation are mutually exclusive. If a realm is a member of a cluster, you cannot use the realm for federation. Similarly, if a realm is part of a federation, the realm cannot be used for clustering.

Universal Messaging provides built-in support for clustering in the form of Universal Messaging Clusters and Universal Messaging Clusters with Sites. Universal Messaging clients can also use the same clustering functionality to communicate with individual, non-clustered Universal Messaging realms in Shared Storage (see "[Shared Storage Configurations](#)" on page 135) server configurations.

Universal Messaging clusters provide support for business contingency and disaster recovery natively or in conjunction with existing BCP enterprise solutions.

From a client perspective a cluster offers *resilience* and *high availability*. Universal Messaging clients automatically move from realm to realm in a cluster as required or when specific realms within the cluster become unavailable to the client for any reason. The state of all client operations is maintained so a client moving will resume whatever operation they were previously carrying out.

The following diagram represents a typical three-realm cluster distributed across three physical locations:



Tip: Since the underlying purpose of a cluster is to provide resilience and high availability, we advise against running all the servers in a cluster on a single physical or virtual machine in a production environment.

Clustered Resources

A Universal Messaging realm server is a container for a number of messaging resources that can be clustered:

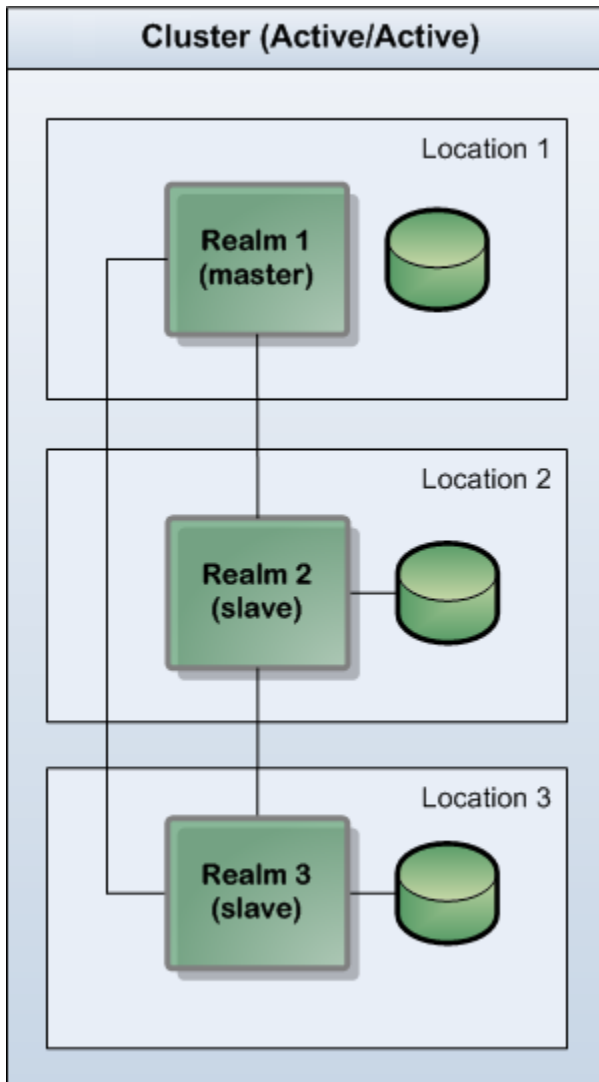
- Universal Messaging Channels
- JMS Topics
- Universal Messaging Queues
- JMS Queues
- DataGroups
- Access Control Lists
- Resource Attributes including Type, Capacity, TTL
- Client Transactions on Universal Messaging Resources

Within the context of a cluster, a single instance of a channel, topic or queue can exist on every node within the cluster. When this is the case all attributes associated with the resource are also propagated amongst every realm within the cluster. The resource in question can be written to or read from any realm within the cluster.

The basic premise for a Universal Messaging cluster is that it provides a transparent entry point to a collection of realms that share the same resources and are, in effect, a mirror image of each other.

A Universal Messaging cluster achieves this by the implementation of some basic concepts described below:

- [" Masters & Slaves" on page 118](#)
- ["Quorum" on page 119](#)
- ["Master Election" on page 120](#)
- ["Message Passing" on page 125](#)
- [" Outages & Recovery" on page 126](#)
- [" Clustered Resources" on page 126](#)

Configuration Option 1: Universal Messaging Clusters

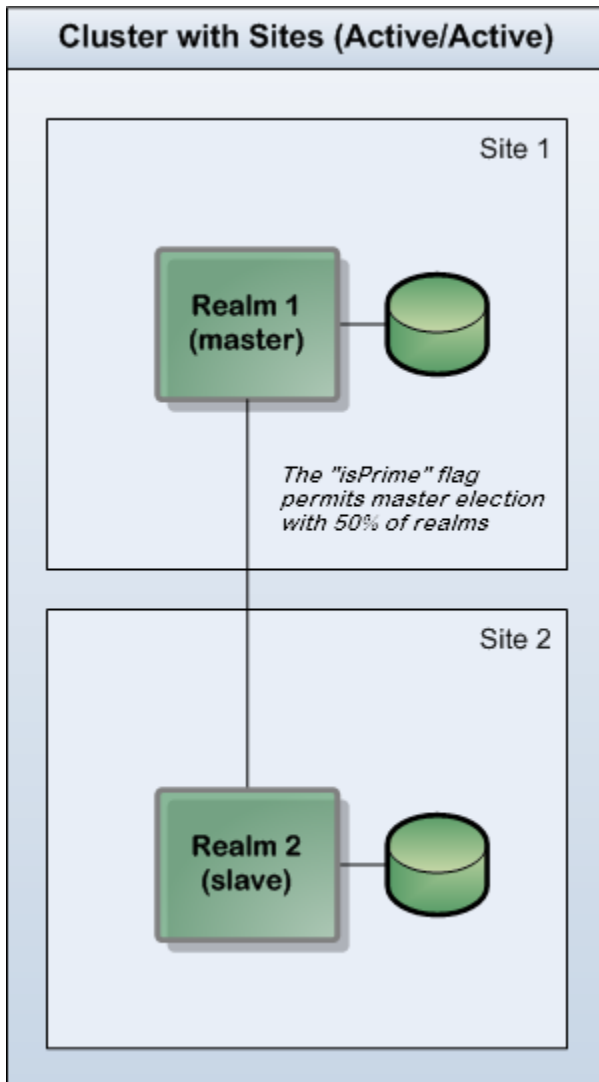
This approach offers the following features:

- Active/Active
- Transparent Client Failover
- Transparent Realm Failover
- Provides Load Balancing and Connection Scalability

Universal Messaging clusters are our recommended solution for high availability and redundancy. State is replicated across all active realms.

With 51% of realms required to form a functioning cluster (see the section ["Quorum" on page 119](#) for an explanation of this percentage figure), this is an ideal configuration for fully automatic failover across a minimum of three realms.

Configuration Option 2: Universal Messaging Clusters with Sites



This approach offers the following features:

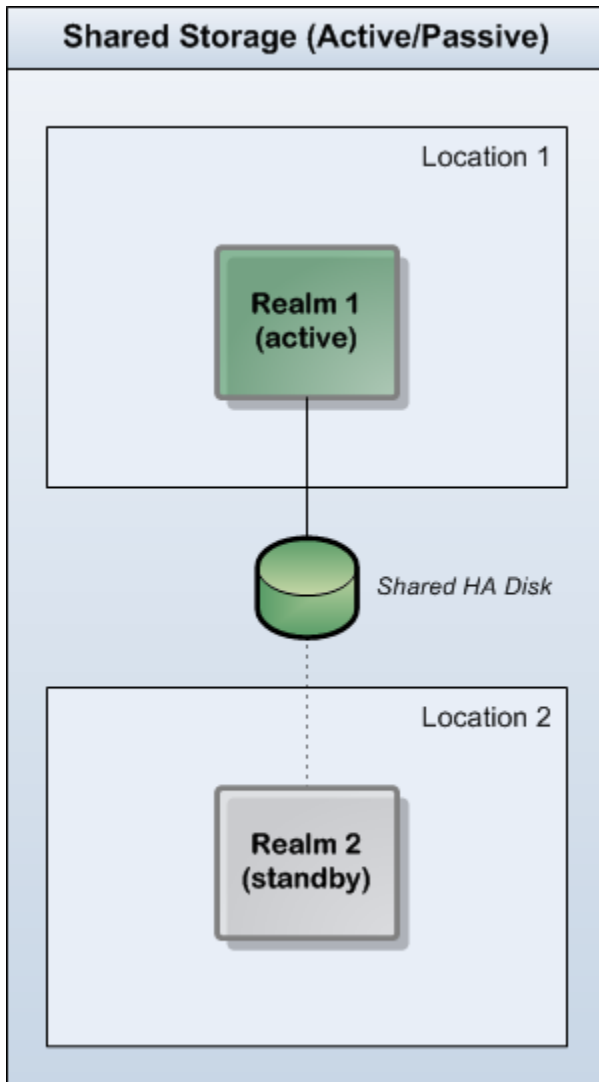
- Active/Active
- Transparent Client Failover
- Semi-Transparent Realm Failover
- Provides Load Balancing and Connection Scalability

Universal Messaging Clusters with Sites provide most of the benefits of Universal Messaging Clusters but with less hardware and occasional manual intervention.

This configuration is designed for two sites, such as Production and Disaster Recovery sites, containing an equal number of realms (for example, one realm on each site or two realms on each site). In such a configuration, if the communication between the sites is lost, neither site can achieve the quorum of 51% of reachable realms required

for a functioning cluster. This situation can be resolved by defining one of the sites to be the so-called *prime site*. If the prime site contains exactly 50% of reachable realms in the cluster, the prime site is allowed to form a functioning cluster. Failover is automatic should the "non-prime" site fail, and requires manual intervention only if the prime site fails. For details of clusters with sites, see the section ["Clusters with Sites" on page 127](#).

Configuration Option 3: Shared Storage Configurations



This approach offers the following features:

- Active/Passive
- Transparent Client Failover
- Manual Realm Failover
- No Load Balancing Features

As an alternative to native Universal Messaging Clusters, Shared Storage (see ["Shared Storage Configurations" on page 135](#)) configurations can be deployed to provide disaster recovery options.

This approach does not make use of Universal Messaging's built-in cluster features, but instead allows storage to be shared between multiple realms - of which only one is active at any one time.

In general, we recommend the use of Universal Messaging Clusters or Universal Messaging Clusters with Sites.

Client Connection

A Universal Messaging client, whether using the Universal Messaging API or JMS, accesses Universal Messaging realms and their resources through a custom URL called an RNAME. When accessing resources in a cluster, clients use a comma separated array of RNAMEs. This comma separated array can be given to the client dynamically when the client connects to any member of a cluster. If a connection is terminated unexpectedly, the Universal Messaging client automatically uses the next RNAME in its array to carry on.

For example, if we have a cluster consisting of 3 realms, your `nSession` object can be constructed using the 3 RNAME URLs associated with each of the realms in the cluster.

Once connected to a realm in a cluster, you can then obtain references to `nChannel` and `nQueue` objects (or in JMS, create a `Session` followed by a topic or queue).

Each event/message within Universal Messaging is uniquely identified by an event ID regardless of whether it is stored on a channel, topic or queue. A clustered channel, topic or queue guarantees that every event published to it via any realm within the cluster will be propagated to every other realm in the cluster and will be identified with the same unique event ID. This enables clients to seamlessly move from realm to realm after disconnection and ensure that they begin from the last event consumed based on this unique event ID.

For scenarios where failover is handled at the network level, Universal Messaging sessions can be moved to alternate servers transparently without the use of multiple RNAMEs.

Client Failover Using Multiple RNAMEs

Using an array of RNAME URLs allows client applications to seamlessly fail over to different cluster nodes without the use of any third party failover software.

For example, in a three realm clustered scenario, a client's RNAME string may contain the following RNAME URLs:

```
nsp://host1:9000,nsp://host2:9000,nsp://host3:9000
```

When we first connect, the `rnames[0]` (in this example, `nsp://host1:9000`) will be used by the session, and the client application will connect to this realm. However, should we disconnect from this realm, for example if `host1` crashes, the client API will

automatically reconnect the client application to the cluster member found at `rnames[1]`, which is `nsp://host2:9000`.


The Universal Messaging client API for Java allows you to configure a session so that it is always connected to the master node in a cluster (the client will follow the master). This behavior can be configured through the method `setFollowTheMaster` of the `nSessionAttributes` class (Universal Messaging client API), or in the JMS `ConnectionFactory`. The default configuration is that the client will not follow the master.

You can also configure the client to follow the master by setting the system property `FollowTheMaster` on the client side when you start your client application. To do this, use `"-DFollowTheMaster=true"` on the command line. Alternatively, you can invoke `java.lang.System.setProperty("FollowTheMaster", "true")` in your client application code.

Masters and Slaves

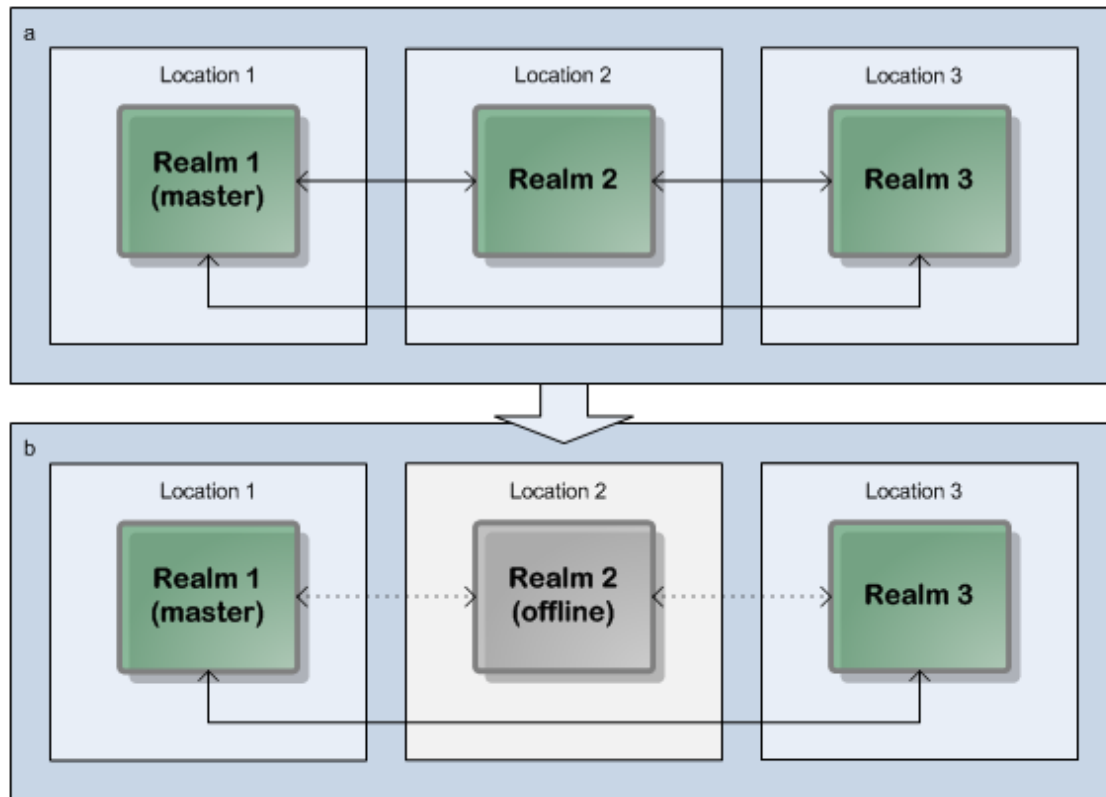
As explained in the Clustering Overview (see ["Clusters: An Overview" on page 111](#)), a *cluster* is a collection of Universal Messaging Realm Servers (realms).

Each cluster has one realm which is elected as master, and all other realms are deemed slaves. The master is the authoritative source of state for all resources within the cluster.



Three-realm cluster over three locations: one master and two slaves.

Should a realm or location become unavailable for any reason, the cluster's remaining realms should be able to carry on servicing clients:



Three-realm cluster over three locations: cluster continuation with one missing slave.

Note: Dotted lines represent interrupted communication owing to server or network outages.

For publish/subscribe resources (see "[Messaging Paradigms supported](#)" on page 17), each published event will be allocated a unique event ID by the master, which is then propagated to each slave.

Quorum

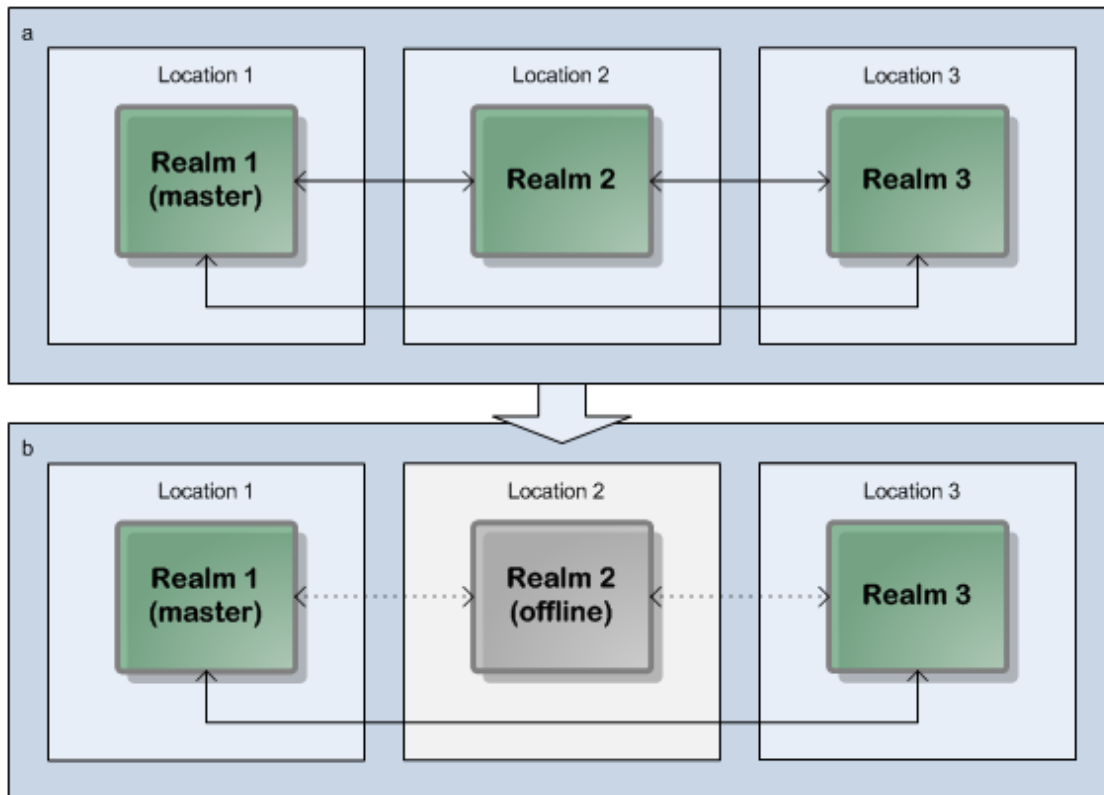
Quorum is the term used to describe the state of a fully formed cluster with an elected master. In order to achieve quorum, certain conditions need to be met. Most importantly, *51% or more* of the cluster nodes must be *online and intercommunicating* in order for quorum to be achieved.

There is an exception to the 51% rule if you use Clusters with Sites. This allows quorum to be achieved in certain circumstances with just 50% of the cluster nodes online and intercommunicating. See the section "[Clusters with Sites](#)" on page 127 for details.

Example: Quorum in a Three-Realm Cluster

In this example, we examine a three-realm cluster, distributed across three physical locations (such as a primary location and two disaster recovery locations). The 51%

quorum requirement means there must always be a minimum of two realms active for the cluster to be online and operational:



Three-realm cluster over three locations: a 67% quorum is maintained if one location/realm fails.

Note: Dotted lines represent interrupted communication owing to server or network outages.

Split-Brain Prevention

Quorum, in conjunction with our deployment guidelines, prevents *split brain* (the existence of multiple masters) scenarios from occurring. By requiring at least 51% of realms to be online and intercommunicating before an election of a new master realm can take place, it is impossible for two sets of online but not intercommunicating realms to both elect a new master.

To simplify the reliable achievement of quorum, we recommend a cluster be created with an *odd number of member realms*, preferably in at least three separate locations. In the above three-realm/three-location example, should any one location become unavailable, sufficient realms remain available to achieve quorum and, if necessary, elect a new master (see ["Election of a new Master" on page 120](#)).

Election of a new Master

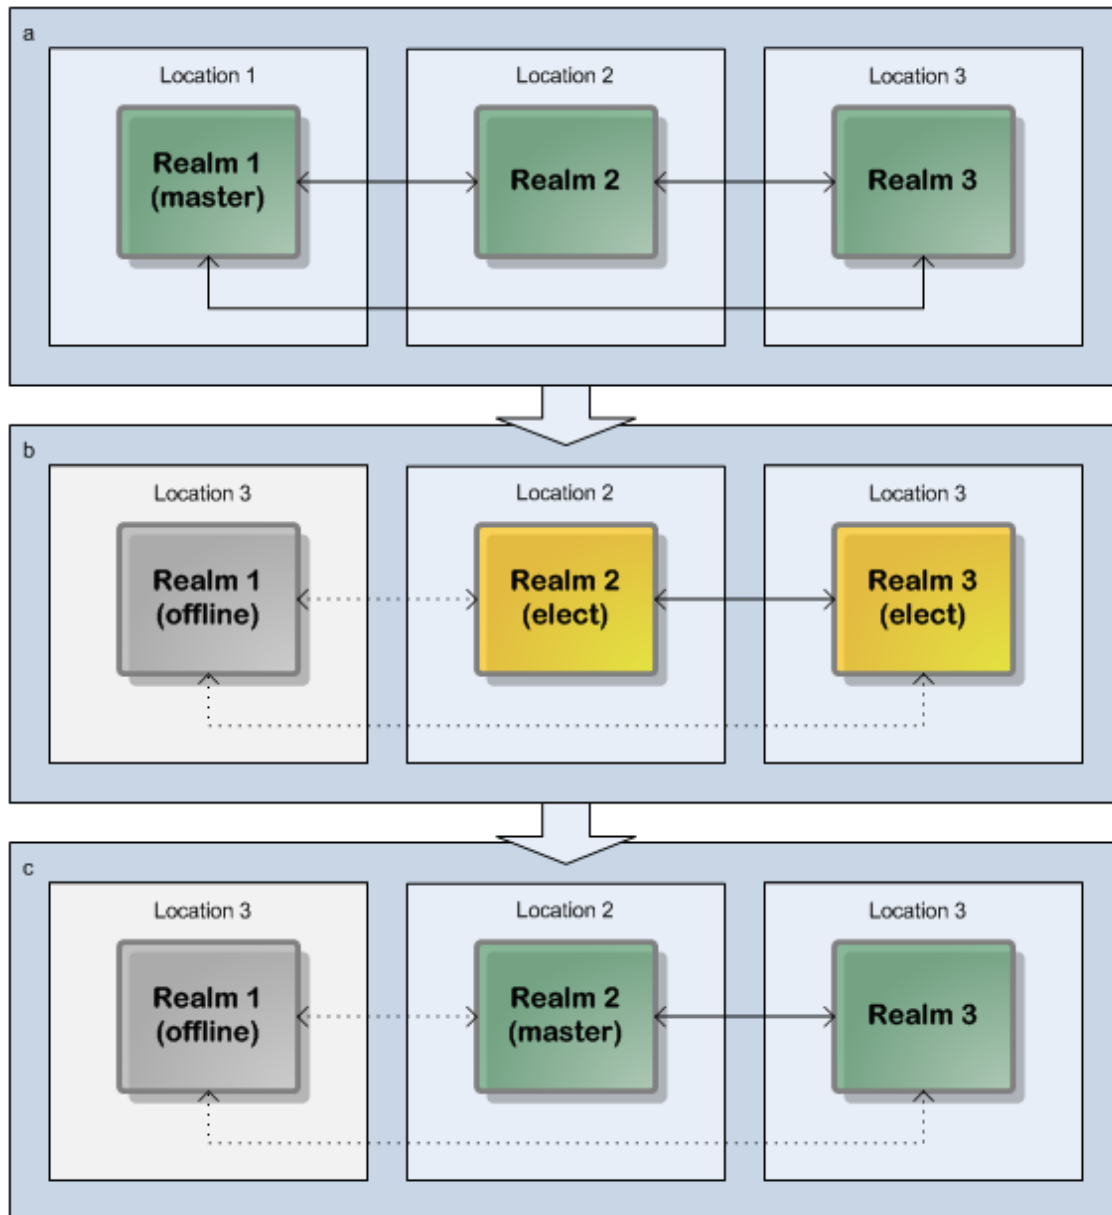
A master realm may unexpectedly exit or go offline owing to power or network failure. In this event, if the remaining cluster nodes achieve 51% or greater quorum

(see ["Quorum" on page 119](#), they will elect a new master realm between them and continue to function as a cluster.

The process of the master election involves all remaining realms in the cluster. Each remaining realm submits a vote across the cluster that results in the new master once all votes are received and the number of votes is greater than or equal to 51% of the total cluster members.

Example: Master Election in a Three-Realm Cluster

In this example, we examine a three-realm cluster, distributed across three physical locations (such as a primary location and two disaster recovery locations). The master realm has failed, but the remaining two realms achieve a quorum of 67% (which satisfies the 51% quorum minimum requirement), so will elect a new master and continue operating as a cluster:




Three-realm cluster over three locations: quorum permits election of a new Master and cluster continuation.

Note: Dotted lines represent interrupted communication owing to server or network outages.

Examples: Insufficient Quorum for Master Election

In this example, we again examine a three-realm cluster, distributed across three physical locations. In this case, both the master realm and one slave realm have failed, so the remaining realm represents only 33% of the cluster members (which does not satisfy the 51% quorum minimum requirement). As a result, it cannot elect a new master, but

will instead disconnect its clients and attempt to re-establish communications with the other realms with the aim of reforming the cluster:




Three-realm cluster over three locations: insufficient quorum prevents election of a new Master or cluster continuation.

Note: Dotted lines represent interrupted communication owing to server or network outages.

A second example highlights both a realm's perspective of quorum, and prevention of split-brain (multiple masters) configurations. In this example, one realm server has failed, while two realms are still active. Also, in this particular example, we assume network connectivity between all realms has failed. As far as each active realm is concerned, therefore, it is the only functioning realm, representing only 33% of the

cluster members. As you might expect, this is insufficient for the 51% required quorum, and is thus also insufficient for the continued operation of the cluster. Both "active" realms will disconnect their clients and attempt to re-establish communications with the other realms, with the aim of achieving quorum and reforming the cluster:



Three-realm cluster over three locations: lack of network interconnectivity prevents cluster continuation.

Note: Dotted lines represent interrupted communication owing to server or network outages.

Note that in the above example, although the two active realms were unable to communicate with each other, it is possible that they were able to communicate with clients. Here, the importance of the 51% quorum rule can be clearly seen: without it,

both realms would have elected themselves master, which could have led to logically unresolvable conflicts once the cluster was reformed. It is therefore essential that the cluster was disabled until such time as a 51% quorum could be achieved.

Clearly, for the situation where we have a total cluster failure, i.e. all realms or locations are offline, the cluster is also deemed inaccessible.

Message Passing

Message passing between cluster realms enables state to be maintained across the member realms. The complexity of the message passing differs somewhat depending on the scenario.

Message Passing in Topics

It is possible to publish to topics on either master or slave realm nodes.

When we publish to the master and subscribe from both master and slave nodes, the master will simply pass the event onto each slave for delivery with the correct event ID, and each slave will maintain the same event ID as set by the master.

When publishing to a topic on a slave node, the slave has to contact the master for the correct event ID assignment before the event is then propagated to each slave.

Message Passing in Queues

When using queues, the message passing is much more complex, since each read is destructive (i.e. it is immediately removed from the queue after it is delivered successfully).

Consider a situation where we have a cluster of 5 realms, and each realm has a consumer connected to a queue, *somequeue*. Assume we publish 2 events directly to the master realm's *somequeue* object.

If the first event happens to be consumed by a consumer on the master realm, each slave realm will be notified of the consumption of the event from *somequeue* and thus remove the event from its own local copy of *somequeue*.

If the next event is consumed by a consumer on some slave realm, then the slave realm will notify the master of the event being consumed. The master will update its local *somequeue*, and then propagate this change to all other slave realms (to update their own local copies of *somequeue*).

Both the Universal Messaging API and the JMS standard define transactional semantics for queue consumers which add to the complexity of the message passing. For example, a consumer may effectively roll back any number of events it has consumed but not acknowledged. When an event is rolled back, it must then be re-added to the queue for re-delivery to the next available queue consumer (which may exist on any of the slave realms). Each event that is rolled back requires each slave realm to maintain a cache of the events delivered to transactional consumers in order for the event to be effectively restored should it be required. The state of this cache must also be maintained identically by all cluster members. Once an event is acknowledged by the consumer (or

the session is committed), these events are no longer available to any consumer, and no longer exist in any of the cluster member's queues.

By now it should be clear that certain scenarios will require more message passing between cluster members than others. Nevertheless, there are a huge number of benefits associated with using clusters in terms of scalability.

Outages and Recovery

Should any cluster member realm exit unexpectedly or become disconnected from the remaining cluster realms, it needs to fully recover the current cluster state as soon as it restarts or attempts to rejoin the cluster.

When a cluster member rejoins the cluster, it automatically moves into the *recovery state* until all local stores are recovered and its state is fully validated against the current master realm.

In order to achieve this, each clustered resource must recover the state from the master. This involves a complex evaluation of its own local stores against the master realm's stores to ensure that they contain the correct events, and that any events that no longer exist in any queues or topics are removed from its local stores. With queues for example, events are physically stored in sequence, but may be consumed non-sequentially (for example using message selectors that would consume and remove, say, every fifth event). Such an example would result in a fairly sparse and fragmented store, and adds to the complexity of recovering the correct state. Universal Messaging clusters will, however, automatically perform this state recovery upon restart of any cluster member.

Creating Clustered Resources

Channels, topics and queues can be created *cluster wide*, which ensures state is maintained across the cluster as described in the Clustering Overview (see "[Clusters: An Overview](#)" on page 111). Once a channel, topic or queue is created as cluster wide, any operations upon that resource are also propagated to all cluster members.

There are a number of ways to create cluster resources once you have created your cluster. The Enterprise Manager application is a tool that provides access to all resources on any realm within a cluster or on standalone realms. This graphical tool is written using the Universal Messaging Client and nAdmin APIs and allows resources to be created, managed and monitored from one central point.

Because the Enterprise Manager tool is written using our own APIs, all operations you can perform using the tool are also available programmatically using the Universal Messaging APIs, allowing you to write customized applications for specific areas of interest within a realm or a cluster.

Once a realm is part of a cluster, you can centrally manage its resources and configuration. For example, realm access control lists (ACLs) can be updated on any member realm and the change will be propagated to all other member realms. Clustered channels and queue ACLs can also be changed on any cluster member and the change is then propagated to the other cluster members. Configuration changes (in the Enterprise

Manager tool) can also be made on one realm and propagated to all other realms in the cluster.

This provides a powerful way of administering your entire Universal Messaging environment and its resources.

Inter-Cluster Connections

Inter-cluster connections can be formed between clusters in order to allow joins to be created between stores on these separate clusters. Inter-cluster connections are bi-directional, allowing joins to be formed between clusters in either direction once the inter-cluster connection has been set up.

In this way, these connections can facilitate inter-cluster routing of messages. Inter-cluster connections do not, however, provide full namespace federation across remote clusters. They are designed to support inter-cluster message propagation on explicitly joined stores, rather than mounting clusters in the namespace of remote clusters, as in realm federation.

Note: Inter-cluster connections and realm federation are mutually exclusive, and they *cannot be used together* in any combination.


Inter-Cluster connections can be added either using the Enterprise Manager or programmatically.

Clusters with Sites

Sites - an exception to the Universal Messaging Cluster Quorum Rule (see ["Quorum"](#) on page 119).

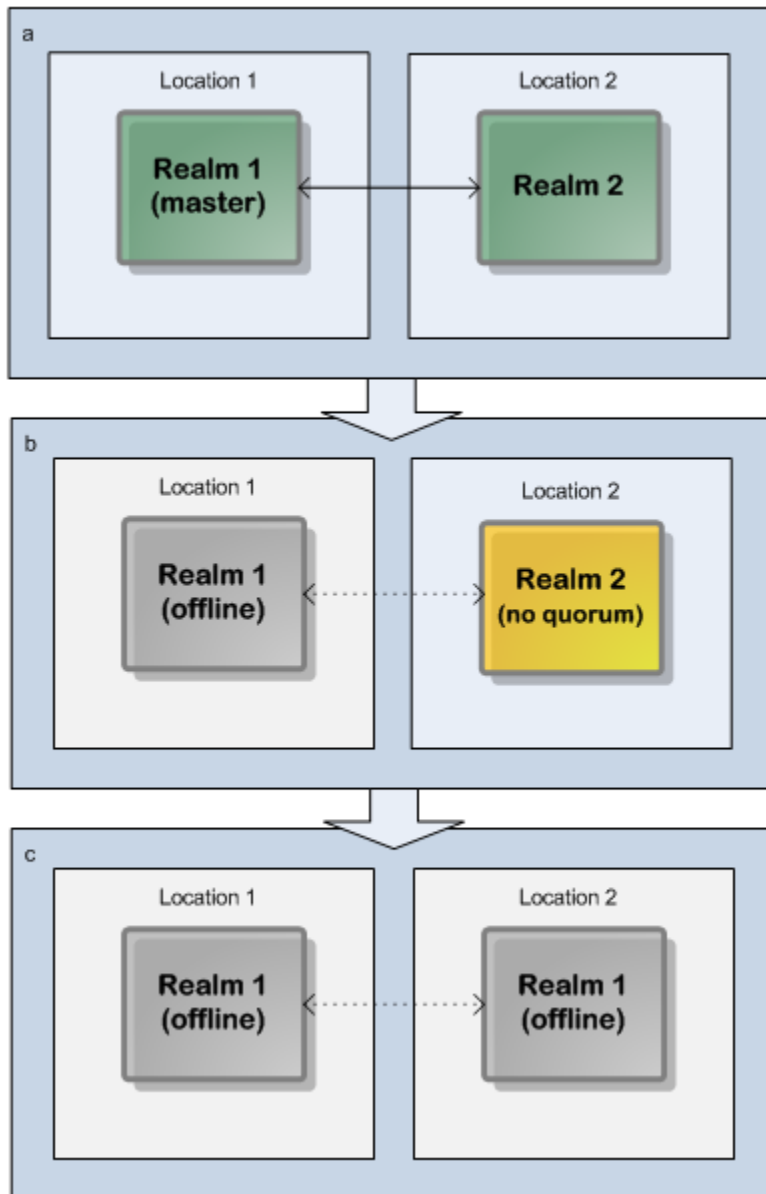
General overview of Clusters with Sites

Although our recommended approach to deploying a cluster is a minimum of three locations and an odd number of nodes across the cluster, not all organizations have three physical locations with the required hardware. In terms of BCP (Business Continuity Planning), or DR (Disaster Recovery), organizations may follow a standard approach with just two locations: a primary and backup:



Two-realm cluster over two sites, using Universal Messaging Clusters with Sites.

With only two physical sites available, the quorum rule of 51% or more of cluster nodes being available is not reliably achievable, either with an odd or even number of realms split across these sites. For example, if you deploy a two-realm cluster, and locate one realm in each available location, then as soon as either location is lost, the entire cluster cannot function because of the 51% quorum rule:



Two-realm cluster over two locations: a 51% quorum is unachievable if one location/realm fails.

Note: Dotted lines represent interrupted communication owing to server or network outages.

Similarly, if you deployed a three-node cluster with one realm in Location 1 and two in Location 2, and then lost access to Location 1, the cluster would still be available; if, however, you lost Location 2, the cluster would not be available since only 33% of the cluster's realms would be available.

This problem has been addressed through the introduction of *Universal Messaging Clusters with Sites*. The basic concept of Sites is that if only two physical locations are available, and Universal Messaging Clustering is used to provide High Availability and Disaster Recovery, it should be possible for one of those sites to function with less

than 51% of the cluster available. This is achieved by allowing an additional vote to be allocated to either of the sites in order to achieve cluster quorum.

If, for example, you have two sites and each site contains just one realm, making a total of 2 realms in the cluster, the additional vote for one of the sites raises the availability from "one out of two" to "two out of three", i.e. 67%, so the required quorum of 51% is achieved.

As a further example, if you have two sites and each site contains two realms, making a total of 4 realms in the cluster, the additional vote for one of the sites raises the availability from "two out of four" to "three out of five", i.e. 60%, so the required quorum of 51% is achieved.

The site that gets the additional vote is called the *prime* site. There can only be one prime site at any given time in a cluster, and any other sites in the cluster are *non-prime* sites.

In a cluster with a production site and a disaster recovery site, you can make either site the prime site, but we recommend you to make the production site the prime site due to the following considerations (assuming a setup with equal numbers of realms on each site):

Prime site	Effect
production site	<p>If the disaster recovery site fails, the production site continues to provide 50% of the available realms in the cluster. Since the production site is the prime site, it gets an additional vote, so the cluster can continue to run.</p> <p>If the production site fails, the disaster recovery site cannot take over automatically, since it cannot achieve the quorum of 51%. This situation requires manual intervention to set the disaster recovery site to be the prime site, so that the cluster can be restarted on the disaster recovery machine.</p>
disaster recovery site	<p>The idea behind setting the disaster recovery site to be the prime site is that if the production site fails, the disaster recovery site can achieve quorum and immediately take over from the production site.</p> <p>This may appear at first to be a good setup, but has a big disadvantage: If the production site is not the prime site, a failure on the disaster recovery site would cause the production site to halt, since the production site by itself cannot achieve a quorum of 51%. This is clearly not what a disaster recovery setup is intended for - a failure in the recovery machine shouldn't halt the production machine.</p>

How sites determine if a cluster can form

The general rule regarding the effect of sites when forming a cluster is as follows: If **exactly** 50% of **all** realms in the cluster are online, and if **all** prime site realms are contactable, a cluster can be formed. In all other cases, a cluster can only form if at least 51% of the cluster's realms are contactable. All other quorum and voting rules are identical, with or without sites.

Defining Sites in the API and in the Enterprise Manager

Within the Universal Messaging Admin API, and specifically within a cluster node, you can define individual *Site* objects and allocate each realm within the cluster to one of these physical sites. Each defined site contains a list of its members, and a flag to indicate whether the site as a whole can cast an additional vote. This flag is known as the *isPrime* flag. When the *isPrime* flag is activated for a site, the site becomes the *prime* site.

Setting the *isPrime* flag can also be achieved using the Enterprise Manager tool.

Examples: Achieving Quorum using Universal Messaging Clusters with Sites

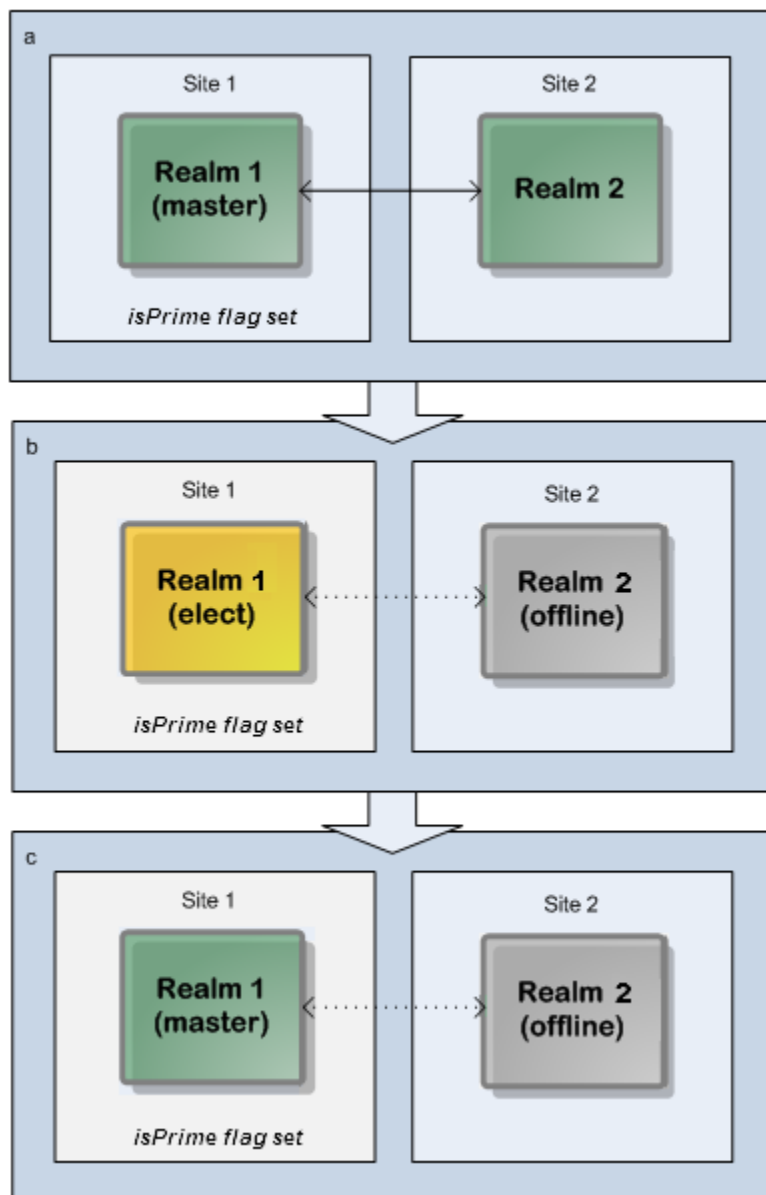
Consider an example scenario where we have a cluster across two physical locations: the default production site and a disaster recovery site, and each site has one realm. Without using sites, this configuration wouldn't be able to satisfy the 51% quorum rule in the event of the loss of one location/realm. The same technique can be used for sites with as many realms as required.

In a disaster recovery situation, where the production site is lost, the disaster recovery site will achieve quorum with only one of the two nodes available because the *isPrime* flag provides an additional vote for the disaster recovery site.

In these situations it is always advisable to discover the cause of the outage so any changes to configuration are made with the relevant facts at hand.

Example with production site as prime site

Here is the situation if the production site is the prime site (the recommended setup) and the disaster recovery node fails:




Two-realm cluster over two sites: Production site as prime site makes a 51% quorum achievable if non-prime site fails.

Note: Dotted lines in the diagram represent interrupted communication owing to server or network outages.

In this scenario, the cluster's master realm is initially on the production node and stays on the production node after the election of the new master.

Example with disaster recovery site as prime site

Here is the situation if the disaster recovery site is the prime node (as stated earlier, this is not the recommended setup) and the production node fails:



Two-realm cluster over two sites: Disaster recovery site as prime site makes a 51% quorum achievable if non-prime site fails.

Note: Dotted lines in the diagram represent interrupted communication owing to server or network outages.

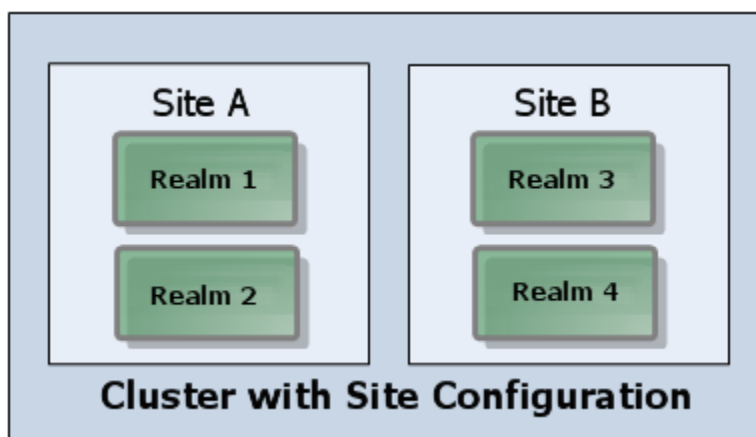
In this scenario, the cluster's master realm is initially on the production site but the new master realm is on the disaster recovery site.

Disaster Recovery using Clusters with Sites

Disaster recovery situations can be managed by using sites in Universal Messaging. In these situations, provided that 50% of the nodes are online, it is possible for the cluster to remain operational.

When using the prime functionality of sites, an administrator should always segregate a cluster into exactly two sites. Each site within the cluster should have exactly the same number of nodes present.

In the sample scenario below, the instructions for disaster recovery using sites assume a 4-node cluster split across two sites: Site A and Site B. Site A is initially the prime site.



Disaster Recovery when the Non-Prime site is lost

When the non-prime site is lost, Universal Messaging will automatically ensure that the cluster is formed with the prime site. The administrator of the system does not need to undertake any action and the cluster will continue to function as normal, providing that at least 50% of the nodes are online to establish a quorum.

When the realms within the non-prime site are brought back online, they will rejoin the cluster and continue to function as normal.

Disaster Recovery when the Prime site is lost

When the prime site is lost, Universal Messaging will not automatically reform the cluster. This is to ensure that a split-brain scenario, where both sites continue to operate independently of each other with conflicting state, does not occur.

Instead the administrator of the system must switch the prime site of the Universal Messaging cluster manually to allow the cluster to reform.

The steps to do this are:

1. Shut down all members of the former prime site (Site A).
2. Change the value of the prime site to the non-prime site (Site B).

3. Observe that a quorum has been achieved using members of new prime site (Site B) only.

When the realms within the former prime site (Site A) are ready to be brought back online, there are additional steps to ensure that they rejoin the existing cluster correctly:

1. Bring each member of the former prime site (Site A) up individually.
2. Ensure that each member joins the new cluster as a slave and acknowledges the change in the prime site flag before bringing up further members of the cluster.
3. Once all members of the cluster are online, ensure that a quorum is achieved and each node agrees on the master and prime site.

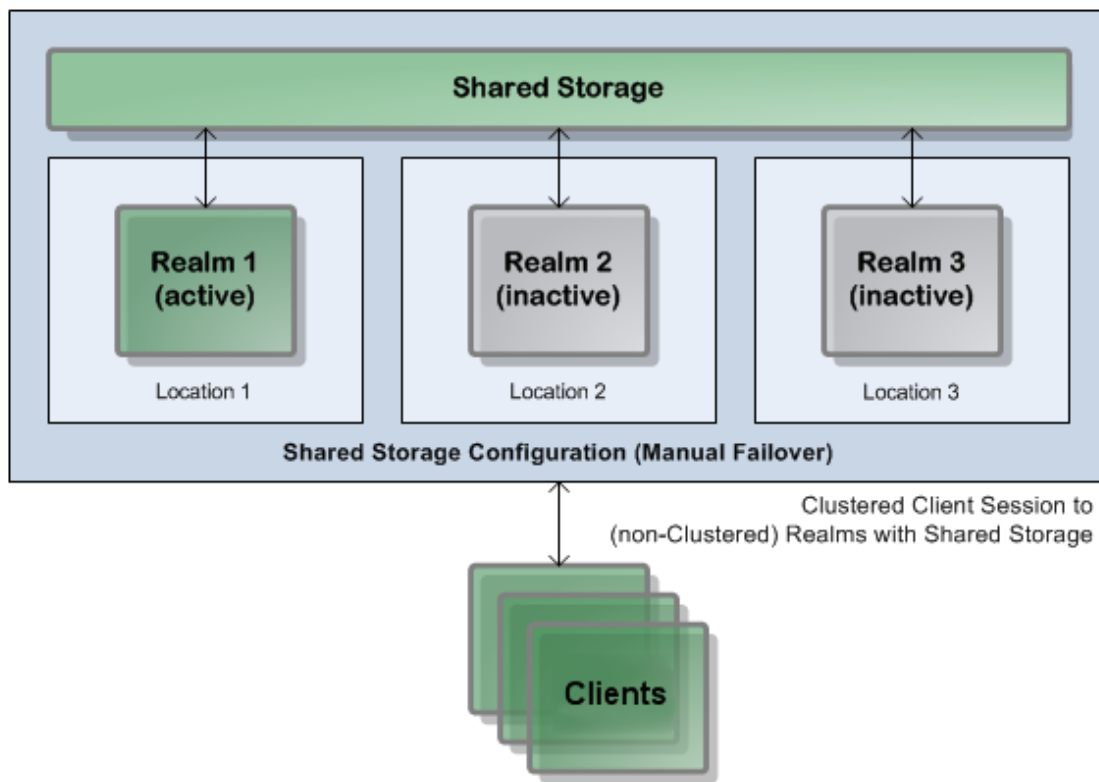
It is important that these instructions are carried out exactly to ensure that only one cluster remains active at any one point in time. Situations where two prime sites exist at the same time may cause 2 separate clusters to form. In these situations it is difficult for these nodes to resynchronize and agree state.

For information on how to perform these steps, see the following topics in the section *Cluster Administration* of the documentation of the Enterprise Manager:

- "Creating and Managing Clusters with Sites"
- " Viewing Cluster Information"

Shared Storage Configurations

Universal Messaging Realms may be set up to share a single virtual or physical disk, in *Shared Storage* configurations. While not technically a cluster, Shared Storage configurations do provide a basic mechanism for rapid failover between realms:



Single active realm with two inactive backup realms in a Shared Storage configuration.

Note that simple Shared Storage configurations do not provide any form of load balancing or scalability, and rely on administrators to perform manual or scripted failover processes. We therefore strongly recommend the use of native Universal Messaging clusters (see ["Clusters: An Overview" on page 111](#)) over Shared Storage configurations.

Setting Up a High Availability (HA) Failover Cluster

Universal Messaging servers can be clustered together to form part of a single logical High Availability (HA) server.

Server Configuration

As an example, let us look at the steps involved in creating a simple 2-node cluster:

- *Realm1* running on *host1.mycompany.com*
- *Realm2* running on *host2.mycompany.com*

Firstly, use the Enterprise Manager tool to create a cluster (see ["Clusters: An Overview" on page 111](#)) with *Realm1* and *Realm2*.

Next, create cluster channels and cluster queues, which ensures these objects exist in both realm servers.

Client Configuration

The next step is to set up your clients so that they are configured to swap between Realm1 and Realm2 in case of failures.

When you initialise a client session with a Universal Messaging server, you provide an array of RNAME URLs (see "[Communication Protocols and RNAMEs](#)" on page 27) as the argument to the `nSessionAttributes` object. This ensures that if you lose the connection to a particular Universal Messaging realm, the session will be automatically reconnected to the next realm in the RNAME array.

Using the configuration above where cluster channels exists on each realm, disconnected clients will automatically continue publishing/subscribing to the channel or queue on the newly connected realm.

For example, to use the two realms described above for failover you could use the following as your RNAME value using a comma separated list of individual RNAMEs:

```
RNAME=nhp://host1.mycompany.com:80,nsp://host2.mycompany.com:9000
```

In this example, note the optional use of different protocols and ports in the specified RNAMEs.

Failover/HA Scenarios

If all subscribers and publishers are configured in this way, then failover is provided in each of the following scenarios:

Scenario I: Subscriber loses connection to a Realm

If a subscriber is consuming data from the sales channel on Realm1 and loses its connection it will automatically attempt to connect to its additional RNAMEs (in this case `nsp://host2.mycompany:9000`) and resume consuming from where it left off.

Scenario II: Publisher loses connection to a Realm

If a publisher loses a connection to its Realm, it will automatically reconnect to the alternative realm and continue publishing there.

Scenario III: Publisher and Subscriber are connected to different Realms

As the above channels on Realm1 and Realm2 are cluster channels, events published to a channels named, say, `/sales` on either Realm will be passed to the `/sales` channel on the other realm. As long as subscribers are consuming from the `/sales` channel on one of the realms they will receive all events. Thus full guaranteed delivery is provided even if the publisher is publishing to Realm1 and subscribers are consuming from Realm2.

For more information on HA configuration options please contact the support team who will be happy to outline the pros and cons of the various HA configurations available.

Separating Client and Cluster Communication

When configuring Universal Messaging for clustering, it is not essential but it is often recommended to have a dedicated interface for client communications and a dedicated interface for cluster communications. There are various reasons for this, as described in this section.


Occasionally it is necessary to stop a client connection while some operation is being carried out on the server or cluster. For example, you may want to delete and recreate a channel, so you need to stop your applications from publishing while this is done. However, you do not want to stop the cluster communication during this time, so it is good practice to separate your cluster communication from client communication. Replication of data, state and other cluster negotiation traffic is very high priority so we need to ensure that it is not delayed in any avoidable way. If the cluster cannot form, then no client traffic will be able to flow, so it is vital that the client traffic cannot impact the cluster formation. If you have lots of clients constantly trying to access the server and getting disconnected because the cluster is not formed, they will be competing to connect at the same time as nodes within the cluster are trying to connect to one another.

By adding separate interfaces onto your servers you can separate client communication from cluster communication and also disable client communication if necessary.

There is no server-defined limit on the number of interfaces that you can create, so it might be a good idea to have different interfaces for different applications. This will allow you to easily disable connections from particular applications.

Adding a Dedicated Interface for Cluster Communication

The first thing to do is ensure that you have more than one interface per server. You can add an interface by selecting the server under the **Realms** node in Enterprise Manager, then opening the **Comms** panel. Click the **Add Interface** button to bring up the dialog for adding an interface.



In the dialog, enter the interface details as required. Also ensure that the **Auto Start** checkbox is enabled, so that if you stop and start the server, your interface is automatically enabled.

Now you will have two interfaces. In this example we have two NHP interfaces, one on port 11000 and one on port 11010. We are going to make the interface running on port 11010 the interface used for cluster communication. The default settings will allow the interface to be used for cluster communication (and the nodes will automatically be notified about this new interface). We recommend you to disable **Advertise Interface**, which is available as a checkbox under the panel **Comms > Basic**. This will stop the interface information being sent to clients.

Note: Do *not* disable the checkbox **Allow Client Connections** under **Comms > Basic**; this is important because Administration API connections, e.g. Enterprise Manager, still count as client connections and we always want to be able to connect an Enterprise Manager. You can optionally disable **Allow for InterRealm** on the *other* interface but it is not necessary.

So now we have our interfaces set up correctly and the interface for cluster communication will not be advertised for clients to use. It is now the responsibility of the system administrator to ensure that the cluster interface is not used by clients for normal connections; all applications should use the other interface.

Temporarily Disconnecting the Client Interface

Let us assume that we want to temporarily disable client connections. All we have to do is stop the interface that we are using for our client connections. This is the best way to stop client traffic, but an alternative is to disable the **Allow Client Connections** checkbox temporarily.

Note: your Enterprise Manager may be connected to the client interface, so if you stop this interface or disable client communication, you will get disconnected. It is important that you still allow client connections on the cluster interface simply so that you can connect an Enterprise Manager to that interface while your other applications remain disconnected.

Once you are ready to allow clients back in, you can re-enable the client interface.

Further Considerations

Network partitioning

Rather than relying on clients not knowing about the cluster interface, it is often a good idea to bind that interface to a different network or VLAN. This is done by specifying an adapter to bind to rather than using 0.0.0.0. This way you restrict connections on this interface to users within that network. It is in fact best practice to bind to a specific network interface rather than 0.0.0.0, especially if running containers.

Stricter Access Control

In the example above we simply ensure that clients do not know about the cluster interface. To enforce this you could use network partitioning as explained above, or you could restrict access using ACLs or VIA entries. It is not possible to set interface-specific ACL entries, so to restrict connections you would need to temporarily stop them connecting using the server level ACL. To do this you would add a specific ACL entry for your administration-related connection and then disable access for everyone else.

Another alternative is to use VIA entries, which are set on individual interfaces. A VIA entry lets you restrict what interfaces a user and host are allowed to connect to. You would need to add an entry that covers all clients into the VIA list for the client interface and leave the VIA list empty for the cluster interface.

For more information on VIA entries, see *Interface VIA Rules* in the Administration Guide.

One more alternative is to use three interfaces, an administration interface, a cluster interface and a client interface. You could disable client communication on the cluster interface to limit traffic to cluster communication only. You could then use SSL on the administration interface and only provide administration users with the certificate required to connect. This way you can disable the client interface to stop client communication and no client will be able to connect to the cluster interface at all.

Multicast: An Overview

Universal Messaging's ability to provide 'ultra-low latency' messaging has been further developed with the introduction of multicast options in addition to unicast to distribute messages to client applications.

With *unicast*, the server must physically write the message once for each destination client:




Figure 1: Universal Messaging in Unicast Mode

With *multicast*, a message is written to the network once where it is then routed to all connections in that multicast group:




Figure 2: Universal Messaging in Multicast Mode

Multicast has clear performance improvements in terms of scalability. Where the performance of unicast gradually degrades as the number of destinations increases, multicast performance does not as the server still only has to write a message once per destination.

Universal Messaging supports multicast for cluster communication and can be enabled on individual DataGroups. Universal Messaging will then transparently operate in a multicast delivery mode for clients on networks capable of supporting the multicast protocol else continue to use unicast.

When you create a DataGroup you can pass a flag to enable multicast delivery. The Universal Messaging realm will automatically begin delivering events published to that group via multicast. When a client DataStream is added to a multicast enabled Data Group, it will transparently receive the information it needs to begin consuming via multicast. Multicast delivery may not be possible for all clients so initially the client will be sent events via both unicast and multicast. The user will only be delivered the data once however.

The Universal Messaging server will quickly detect whether or not the multicast packets are reaching the client as it does or does not receive acknowledgements. If the client does not support multicast, the server will simply continue to send events to that client via unicast. If the multicast packets do reach the client then after a period of time when both unicast and multicast are in sync, the Universal Messaging Server will stop sending unicast events to that DataStream.

Multicast addresses packets using the User Datagram Protocol (UDP) as opposed to TCP which is used in unicast. UDP is a connectionless protocol and does therefore not guarantee delivery or packet ordering in way that TCP does. However Universal Messaging still provides these guarantees by implementing the required checks at the application layer. These checks happen completely transparently to the user and any packet retransmissions can be monitored using the nAdminAPI or the Enterprise Manager.

Shared Memory (SHM)

Shared Memory (SHM) is our lowest latency communications driver, designed for inter-process communication (IPC). SHM can be used for client and inter-realm (cluster and join) communication. It is operating system independent as it does not depend on any OS specific libraries, making it simple to set up and use.

As the name suggests, shared memory allocates blocks of memory that other processes can access - allowing the same physical computer to make connections without network overhead. This has many advantages, one of which is that when the realm and the data source (publisher) are located on the same physical computer, there is no network latency added between them. This results in less latency for delivery to clients.

Advantages

- Lowest latency
- No network stack involved
- Efficient use of resources with no Network IO required

Disadvantages

- Same physical machine only
- Only currently supported by certain JVMs such as Oracle JDK1.6.0_32, JDK 7, JDK 8 and Azul Zing. On HP-UX systems, shared memory drivers are currently not supported.

Scalability

Realm Benchmarks

The benchmarks detailed in this section of the website are designed to give indications of the performance levels achievable with Universal Messaging.

These benchmarks and their accompanying results are presented as a guide only. They provide indications of the levels of performance for the given context in which they have been run. Performance in contexts or environments different to one we present are likely to give different results.

The performance of these benchmarks is limited primarily by the available network infrastructure and machine which the Universal Messaging Realm Server resides on. The results were produced by running the benchmark using commercially available machines which may typically be used to host services in an enterprise environment.

If you would like access to the benchmarking tools to run these tests yourselves, please email our support team.

Test Scenarios

Descriptions of several tests are detailed below, with each test running with a different configuration to provide indications of performance in varying types of context. These tests include:

- High Update Rates with Multi-cast delivery
- High Update Rates
- Medium Update Rates
- Low Update Rates


The low and medium rate tests are designed to model traffic which may be typical of non-time critical applications which may instead focus the number of concurrent connections a server can handle.

The high rate tests are designed to model latency critical applications which use small, frequent messages. This kind of traffic is more typical of trading systems which some of our customers deploy.

High Update Rate (Using Multi-Cast Delivery)

Multi-cast delivery is a feature that allows for ultra low latency delivery of streaming data to clients. The characteristics of the test are the same as the high update rate test, leveraging the benefits of multi-cast delivery to cope with the most demanding performance requirements.


DataGroups	500
Clients	5,000
Increment Rate	100 every 60 seconds
Subscriptions per Client	50
Messages per second per Client	50
Peak Message Rate	250,000



High Update Rate

The high update rate test is designed to model applications which require low latency delivery times for data which is streamed in small intervals to a small set of clients. These characteristics are often found in trading systems.


DataGroups	500
Clients	5,000
Increment Rate	100 every 60 seconds
Subscriptions per Client	50
Messages per second per Client	50
Peak Message Rate	250,000



Medium Update Rate

The medium update rate test is designed to model applications which require low latency delivery times to data which is streamed to clients at a moderate rate.


DataGroups	500
Clients	30,000
Increment Rate	100 every 60 seconds
Subscriptions per Client	10
Messages per second per Client	10
Peak Message Rate	300,000



Low Update Rate

The low update rate test is designed to model services which update infrequently with the focus on scaling delivery to large numbers of clients in a reasonable time frame.

DataGroups	500
Clients	100,000
Increment Rate	1,000 every 60 seconds
Subscriptions per Client	1
Messages per second per Client	1
Peak Message Rate	100,000



JavaScript High Update Rate

The JavaScript High Update Rate test uses headless simulation clients which communicate with the server using the WebSocket protocol. This provides an indication of the ability of the Universal Messaging Realm Server to scale to serve large numbers of web clients.

DataGroups	500
Clients	5,000
Increment Rate	100 every 60 seconds
Subscriptions per Client	50
Messages per second per Client	50
Peak Message Rate	250,000

Failover

Universal Messaging clients, whether server to client or server to server, automatically handle disconnect and reconnection to any number of alternative Universal Messaging Realm servers.

If a Universal Messaging server becomes unreachable the Universal Messaging session will automatically try to connect to the next defined Universal Messaging server in

the list of RNAME values provided when the nSession object was created. Universal Messaging clients can get this list of RNAMEs from the server during the session handshake or programmatically at any time. The ability to pass RNAME values from server to client dynamically makes it very easy to manage client failover centrally.

Single Universal Messaging Realm

In a single server scenario upon receipt of an abnormal disconnection a Universal Messaging client will automatically try to reconnect. The back-off period and interfaces handling reconnection are all fully exposed in a public API. They can be used to trigger off any number of specific events.

Multiple Universal Messaging Realms

In a multiple server scenario Universal Messaging clustering (see "[Clusters: An Overview](#)" on page 111) allows for multiple Realms to act as live replicas of each other. This has huge benefits both in terms of load balancing and high availability.

Universal Messaging is also fully compatible with high availability and business contingency products implemented within the underlying operating system. This means that Universal Messaging is compatible with existing HA/BCP policies as well as providing its own in built fail over functionality.

Connections Scalability

Single server

A single Universal Messaging server has no hard limit set on the number of client connections that can be made to the server. Universal Messaging servers are implemented using Java's NIO (non-blocking IO) for all client communications including SSL. This means that there are no additional threads created and dedicated to each new client connection.

Multiple Servers

Universal Messaging messaging can be achieved across multiple server instances either by the use of a federated name space (see "[Federation Of Servers](#)" on page 41) or by using a Universal Messaging Cluster (see "[Clusters: An Overview](#)" on page 111).

Universal Messaging servers are aware of each other within a name space. Universal Messaging channels can appear in multiple places in a single name space allowing Universal Messaging to be infinitely scalable. In addition, Universal Messaging's licensing model places no constraints on the number of message servers run. Instead it is based on the number of channels used, thus allowing your client numbers to scale massively without additional cost.

6 MQTT: An Overview

MQTT (Message Queuing Telemetry Transport), is a publish/subscribe, simple and lightweight messaging protocol, designed for constrained devices and low-bandwidth by IBM / Eurotech in 1999. The simplicity and low overhead of the protocol make it ideal for the emerging "machine-to-machine" (M2M) or "Internet of Things" (IoT) world of connected devices, and for mobile applications where bandwidth and battery power are at a premium. The protocol is openly published with a royalty-free license, and a variety of client libraries have been developed especially on popular embedded hardware platforms such as arduino/netduino, mbed and Nanode.

In addition to Universal Messaging's own protocol, NSP interfaces are capable of also accepting MQTT connections over TCP sockets, while NSPS interfaces can accept MQTT connections over SSL/TLS for client implementations that support it.

Connecting

In order to connect to a Universal Messaging server using MQTT, your application needs to use a tcp://host:port URL (NSP Interfaces) or ssl://host:port URL (NSPS Interfaces). MQTT connections are treated in the same way as any other connections by the Universal Messaging realm. If the username is present, the Universal Messaging subject is username@hostname, otherwise the subject is anonymous@hostname.

protocol	user	host	connection	language
mqtt	anonymous	localhost	127.0.0.1:61401	MQTT 4
nsp	krital	localhost	127.0.0.1:61201	Java 1.7.0_60

Figure 1: Connection List with an MQTT connection

This way you can define realm and channel / queue ACLs as you would for any Universal Messaging connection. For example using the IBM WMQTT sample application without a username/password to connect to tcp://localhost:1883 will result in a Universal Messaging subject of anonymous@localhost.

Publishing

MQTT applications can publish events to channels. If the specified channels do not already exist in the Universal Messaging realm, they will be automatically created by the server as MIXED type with a JMS engine.

These channels are regular Universal Messaging channels but all events published from MQTT will be persisted to disk. While it is possible to create other channel types using the Administration API / Enterprise Manager, the mixed type with JMS engine is the recommended combination. Events published via MQTT only contain a byte[] payload and are tagged MQTT. They are fully interoperable with any Universal Messaging subscriber on any client platform supported and can be snooped using the Universal Messaging Enterprise Manager:




Figure 2: Snooping an MQTT message

All messages published from MQTT are flagged as JMS BytesMessage objects.

Note: Publishing to queues via MQTT is not supported.

Subscribing

MQTT applications can subscribe to channels. If the specified channels do not already exist in the Universal Messaging realm, they will be automatically created by the server as MIXED type with a JMS engine.

These channels are regular Universal Messaging channels with all messages being persistent, regardless of whether they are published by MQTT or Universal Messaging applications.

Note: Subscribing to queues via MQTT is not supported.

Quality of Service

Universal Messaging supports QOS levels 0 and 1 as defined by the MQTT standard. This is driven by the MQTT client connection and describes the effort the server and client will make to ensure that a message is received, as follows:

1. QOS 0 (*At most once delivery*): The Universal Messaging realm will deliver the message once with no confirmation
2. QOS 1 (*At least once delivery*): The Universal Messaging realm will deliver the message at least once, with confirmation required.

Note: Universal Messaging does not support QOS level 2 (*Exactly once delivery*). Connections requesting QoS level 2 will be downgraded to QoS level 1 at connection time, as allowed by the MQTT specification.

Will

Universal Messaging fully supports connections with Will settings, which indicate messages that need to be published automatically if the MQTT application disconnects unexpectedly.

7

AMQP

■ AMQP Overview	154
■ AMQP Plugin Configuration	156
■ AMQP Security Configuration	158
■ AMQP Message Transformation	159

AMQP Overview

The Advanced Message Queuing Protocol (AMQP) is an open internet protocol for business messaging. The specification of the AMQP protocol is available at <http://docs.oasis-open.org/amqp/core/v1.0/os/amqp-core-overview-v1.0-os.html>. It defines a binary wire-level protocol that allows for the reliable exchange of business messages between two parties.

In addition to Universal Messaging's own protocol, NSP interfaces are capable of also accepting AMQP connections over TCP sockets, while NSPS interfaces can accept AMQP connections over SSL/TLS for client implementations that support it.

Connecting

In order to connect to a Universal Messaging server using AMQP, your application needs to use an `amqp://host:port` URL (NSP Interfaces) or `amqps://host:port` URL (NSPS Interfaces). AMQP connections are treated in the same way as any other connections by the Universal Messaging realm. If the username is present, the Universal Messaging subject is `username@hostname`, otherwise the subject is `<AnonymousUser>@hostname` where `<AnonymousUser>` is configured under **Configuration > Protocol AMPQ Config > AnonymousUser**.

We recommend you to establish an AMQP JMS connection using the QPID Proton, QPID Legacy or the Swift MQ client libraries and using a SASL plain over TLS connection.

Figure 1: Connection List with an AMQP connection


protocol	user	host	connection	language
AMQP	anonymous	127.0.0.1	127.0.0.1:64839	AMQP 1.0.0
nsp	dada	127.0.0.1	127.0.0.1:64224	Java 1.8.0_51

This way you can define realm and channel / queue ACLs as you would for any Universal Messaging connection. For example, using the one of the JMS Sample applications with AMQP without a username/password to connect to `amqp://localhost:10000` will result in a Universal Messaging subject of `<AnonymousUser>@localhost`.

Publishing

AMQP applications can publish events to channels and queues, which should exist in advance.

Events published via AMQP are fully interoperable with any Universal Messaging subscriber on any client platform supported and can be snooped using the Universal Messaging Enterprise Manager:

Figure 2: Snooping an AMQP message

Subscribing

AMQP applications can subscribe to pre-existing channels and queues.

Temporary Topics and Queues

The Universal Messaging server supports the creation of temporary topics and queues using an AMQP JMS session's `createTemporaryTopic()` and `createTemporaryQueue()` methods. Currently only the "delete-on-close" lifetime policy is supported; a Node dynamically created with this lifetime policy will be deleted at the point that the Link which caused its creation ceases to exist.

Transactions over AMQP

Client applications can perform transactional work over AMQP. Universal Messaging implements the local transaction operations defined in the AMQP specification. For example, if an application communicates to the realm server using a JMS AMQP client library (e.g. Apache Qpid JMS client) it can take advantage of the local transaction functionalities defined in the JMS specification.

Note: We do not currently support the Transactional Acquisition operation defined in the AMQP specification. However, this sets no limitations on using JMS transactions over AMQP.


JMS sample applications over AMQP

The JMS sample applications provided with the distribution of the Universal Messaging server can be configured to communicate over the AMQP protocol.

For more information, see the page "Using the AMQP Protocol" in the JMS code examples in the Developer Guide.

AMQP Plugin Configuration

You can configure AMQP behavior using the configuration parameters that are shown under the **Config** tab in the Enterprise Manager.



The configuration parameters are described in the following table:

Parameter	Description
AllowUserTransformation	When set to "true" the server will try to load a user class for event transformation. This full class name should be provided as an AMQP_TRANSFORM JVM parameter to the server, e.g. <pre>-D AMQP_TRANSFORM=foo.bar.MyTransformer</pre>

Parameter	Description
	and this class should extend the abstract class: <pre>com.pcbsys.nirvana.server.plugins.protocols .amqp.transformation.aTransformation</pre>
AnonymousUser	The user name which the server will use for anonymous users.
BufferSize	The size of the buffer in bytes that will be used to read/write on the AMQP connection.
DefaultNodeMode	The default type of node (topic or queue) if the server is not able to detect it.
Enable	If "true", the server will accept incoming AMQP connections.
MaxFrameSize	Maximum size of an AMQP frame in bytes.
QueuePrefix	The address prefix for specifying queue nodes as required by some clients.
SASL_Anonymous	Enable Anonymous SASL authentication.
SASL_CRAM-MD5	Enable CRAM-MD5 SASL authentication.
SASL_DIGEST-MD5	Enable DIGEST-MD5 SASL authentication.
SASL_Plain	Enable Plain SASL authentication.
SubscriberCredit	Sets initial credit for AMQP subscribers.
Timeout	Sets the network timeout in milliseconds.
TopicPrefix	The address prefix for specifying topic nodes as required by some clients.
TransformToUse	Select which transformation to use on AMQP messages. For more information on transformers see the section " AMQP Message Transformation " on page 159.

Note: The `idle time` of the AMQP connection for the Universal Messaging server can be configured by modifying the **Connection Config > KeepAlive** property.

Note: We recommend that clients use SASL plain over TLS connections.

AMQP Security Configuration

The Universal Messaging server supports the following security configurations when the AMQP protocol is used:

1. Plain AMQP
2. AMQP over SASL
3. AMQP over alternative TLS

Currently, AMQP over negotiated TLS is not supported.

AMQP over SASL

The Universal Messaging server supports the following SASL mechanisms:

1. Anonymous
2. CRAM-MD5
3. DIGEST-MD5
4. Plain

You can enable and disable these mechanisms by using the AMQP configuration options in Enterprise Manager as explained in the section ["AMQP Plugin Configuration" on page 156](#).

We recommend using the SASL Plain mechanism over a TLS Connection.

By default, the SASL anonymous authentication is used and the client can connect without providing a username or password. After the connection has been established, the ability to perform various operations is derived from the realm's ACL lists. If any other SASL type is used, then Basic Authentication also needs to be configured on the realm.

For more information on configuring basic authentication, see the page ["Basic Authentication > Server-side Authentication"](#) in the Java section of the Universal Messaging Developer Guide.

AMQP over alternative TLS

The AMQP connection over alternative TLS can be established when the server has a running NSPS interface. In that case, the client should set the following system variables:

```
set CAKEystore=<TRUST KEYSTORE PATH>
set CAKEystorePASSWD=<TRUST KEYSTORE PASSWORD>
set CKEystore=<CLIENT KEYSTORE PATH>
```

```
set CKEYSTOREPASSWD=<CLIENT KEYSTORE PASSWORD>
```

Then use the `amqp://<hostname>:<port>` URL to establish the connection.

AMQP Message Transformation

The Universal Messaging server can be configured to perform transformations on any messages received and sent over the AMQP protocol. The transformations are controlled from the `TransformToUse` setting of the AMQP plugin configuration. The following settings are available:

Setting	Description
0 - No transformation	When this option is selected, no transformation is applied to the message, and it is passed as a byte array between the clients in the communication. This option is best used when only AMQP clients exchange messages.
1 - Basic Transformation	<p>With this option selected when a message is exchanged over the AMQP protocol, the following header entries are converted:</p> <ol style="list-style-type: none"> 1. priority 2. delivery count / redelivery count 3. durable / delivery mode <p>This option is again best used when only AMQP clients exchange messages, but it also allows for server side filtering, based on the event's header and properties.</p>
2 - Complete Transformation	<p>With this option selected when a message is exchanged over the AMQP protocol, the following header entries are converted:</p> <ol style="list-style-type: none"> 1. priority 2. delivery count / redelivery count 3. durable / delivery mode 4. Time To Live (TTL) 5. First Acquirer / JMS_AMQP_FirstAcquirer <p>The AMQP message's footer is converted to an event dictionary object and put under the "footer" key in the UM event's dictionary and vice versa - if an UM event being sent to an AMQP client has an "footer" dictionary object in</p>

Setting	Description
	<p>the event dictionary, the content of that dictionary will be put as application properties of the AMQP message.</p> <p>If the message being transformed is a JMS message, then the JMS headers will be properly converted.</p>
3 - User Configurable	<p>This option allows you to provide your own message converters. In order to use this option, do the following:</p> <ol style="list-style-type: none">1. Enable the "AllowUserTransformation" option.2. Provide a full class name as an <code>AMQP_TRANSFORM</code> JVM parameter to the server, e.g. <pre>-D AMQP_TRANSFORM=foo.bar.MyTransformer</pre>and this class should extend the abstract class: <pre>com.pcbsys.nirvana.server.plugins.protocols .amqp.transformation.aTransformation</pre>

8 Commonly Used Features

■ Overview	162
■ Sessions	162
■ Channel Attributes	162
■ Storage Properties of Channels and Queues	167
■ Channel Publish Keys	169
■ Queue Attributes	171
■ Events	173
■ Named Objects, Shared Named Objects and Serial Named Objects	173
■ Event Filtering	178
■ Event Filtering using Selectors	179
■ Event Filtering using Subjects and Principals	182
■ Event Dictionaries	185
■ Durable Consumers	185
■ Durable Event Browsing and Purging	185
■ Native Communication Protocols	186
■ Comet Communication Protocols	189
■ Google Protocol Buffers	191
■ Using the Shared Memory Protocol	192
■ Rollback	193
■ Out-of-Memory Protection	193
■ Pause Publishing	195

Overview

This section summarizes commonly used features of Universal Messaging. The features are available using a variety of methods, such as in the Enterprise Manager or in the Server or Client APIs.

Sessions

A session in Universal Messaging represents a logical connection to a Universal Messaging Realm. It consists of a set of session attributes, such as the protocol and authentication mechanism to be used, the host and port the message server is running on and a reconnect handler object.

Most of the session parameters are defined in a string that is called RNAME and resembles a URL. All the sample applications provided use an RNAME Java system property to obtain the necessary session attributes. The following section discusses this in further detail. The RNAME takes the following format.

The RNAME entry can contain an unlimited number of comma-separated values each one representing an interface on a Universal Messaging Realm.

The current version of the Universal Messaging Realm and the Universal Messaging client API supports 4 TCP wire protocols. These are the Universal Messaging Socket Protocol (nsp), the Universal Messaging HTTP Protocol (nhp), the Universal Messaging SSL Protocol (nsps) and the Universal Messaging HTTPS protocol (nhps). These wire protocols are available wherever a connection is required, i.e. client to Realm and Realm to Realm.

See the section *Communication Protocols and RNAMEs* in the Universal Messaging Concepts guide for more information on RNAMEs and communication protocols.

Channel Attributes

Universal Messaging channels provide a set of attributes that define the behaviour of the events published and stored by the Universal Messaging Realm Server. Each event published onto a channel has a unique ID within the channel called Event ID. Using this event ID, it is possible for subscribers to re-subscribe to events on a channel from any given point. The availability of the events on a channel is defined by the chosen attributes of the channel upon creation. Channels can be created either using the Universal Messaging Enterprise Manager or programmatically using any of the Universal Messaging Enterprise APIs.

There are a number of important channel attributes which are discussed below.

Channel Type

Each channel has an associated channel type. Universal Messaging channels can be of the types described in the following table. The difference lies in the type of physical storage used for each type and the performance overhead associated with each type.

Channel Type	Description
Persistent	<p>Persistent channels have their messages stored in the Universal Messaging Realm's persistent channel disk-based store. The persistent channel store is a high performance file-based store that uses one or more files for each channel on that Realm, thus facilitating migrating whole channels to different Realms.</p> <p>For information about setting up a multi-file disk store for persistent channels or persistent queues, see the description of the storage parameter <code>Events per Spindle</code> in the section "Storage Properties of Channels and Queues" on page 167.</p> <p>All messages published to a persistent channel will be stored to disk, hence it is guaranteed that they will be kept and delivered to subscribers until it is purged or removed as a result of a "Time to Live" (TTL) or capacity policy.</p> <p>Messages purged from a persistent channel are marked as deleted; however the store size will not be reduced immediately. If the store is backed by a single file, the store size will not be reduced until maintenance is performed on the channel using the Universal Messaging Enterprise Manager or an Administration API call. Stores backed by multiple files will be cleaned up through the removal of individual files once the events stored on a particular file have all been marked as deleted. This augments the high performance of the Universal Messaging Realm.</p>
Mixed	<p>Mixed channels allow you to specify whether the event is stored persistently or in memory, as well as the Time To Live (TTL) of the individual event. On construction of a Mixed channel, the TTL and Capacity can be set; if you supply a TTL for an event, this is used instead of the channel TTL.</p> <p>Events stored persistently are either stored on a single or multiple files in the same manner as a persistent channel.</p>
Reliable	<p>Reliable channels have their messages stored in the Universal Messaging Realm's own memory space. The first fact that is implied is that the maximum number of bytes that all messages across all reliable channels within a Universal Messaging Realm is limited by the maximum heap size available to the Java Virtual</p>

Channel Type	Description
	Machine hosting that Realm. The second fact implied is that if the Universal Messaging Realm is restarted for any reason, all messages stored on reliable channels will be removed from the channel as a matter of policy. However, as Universal Messaging guarantees never to reuse event IDs within a channel, new messages published in those channels will get assigned event IDs incremented from the event ID of the last message prior to the previous instance stopping.
Simple	<p>Simple channels have their messages stored in the Universal Messaging Realm's own memory space, thus supplying a high-speed channel type. The difference between Simple and Reliable is that the event IDs are reset to 0 in a Simple channel whenever the Universal Messaging Server is restarted, whereas Reliable channels increment and store the Event ID on disk rather than in memory so that in the event of failure, event IDs are not reset to 0.</p> <p>Simple channels are not supported across realms in a cluster.</p>
Transient	<p>A transient channel is like a simple channel in that no event characteristics are stored persistently. In addition to this, data is only ever written to a transient channel when 1 or more consumers are connected to the channel and are able to consume this data. Unlike the simple channel which stores event data in memory, transient channels do not store anything, not even in memory. Transient channels can be thought of as a relay between 1 or more publishers and 1 or more subscribers.</p> <p>Transient channels are not supported across realms in a cluster.</p>
Off-heap	Off-heap channels, similar to reliable channels, store the events in memory, but this channel type uses off-heap memory. This allows the normal JVM heap memory to be left free for short lived events, while longer living events can be stored off the JVM heap. This reduces the work the garbage collector needs to do to manage these events since they are out of range of the garbage collector.
Paged	Paged channels allows users access to high speed, off-heap memory mapped files as the basis for the event storage. As for the other channel types, each channel or queue has its own unique file mapping. While similar to the off-heap channel type, where the events are stored in the off heap memory, the paged channel store is also persistent, thereby allowing recovery of events that are written into this store.

Channel TTL

The TTL ("time to live") for a channel defines how long (in milliseconds) each event published to the channel will remain available for subscribers to consume. Specifying a TTL of 0 will mean that events will remain on the channel indefinitely. If you specify a TTL of 10000, then after each event has been on the channel for 10000 milliseconds, it will be automatically removed by the server.

Channel Capacity

The capacity of a channel defines the maximum number of events that may remain on a channel once published. Specifying a capacity of 0 will mean that there is no limit to the number of events on a channel.

Example: If you specify a capacity of 10000, then if there are 10000 events on a channel, and another event is published to the channel, the oldest event on the channel will be automatically removed by the server to make room for the new event.

Note: You can use the storage property `Honour Capacity` to prevent further events being published to the channel if the channel capacity has been reached. See the section ["Storage Properties of Channels and Queues" on page 167](#) for details.

Dead Event Store

When events are removed automatically, either by the capacity policy of the channel or the age (TTL) policy of the individual events, and they have not been consumed, it may be a requirement for those events to be processed separately. If so, channels or queues can be created with a dead event store so any events that are purged automatically before being consumed will be moved into the dead event store. Dead event stores themselves can be a channel or a queue and can be created with any attributes you wish.

Use JMS Engine / Use Merge Engine

By default, Universal Messaging retains all events on any given channel until the events reach their individual time-to-live limit (see the description of Channel TTL above).

This default behavior can be modified by selecting either of the following optional engines:

Engine	Description
JMS Engine	<p>The JMS Engine deals with JMS topics within Universal Messaging.</p> <p>The JMS engine causes events to remain on the channel as long as there is "interest" in receiving the events. In this sense, "interest" means that at least one of the following conditions is met:</p> <ul style="list-style-type: none"> ■ there are one or more active non-durable subscribers on the channel who have not yet consumed the events;

Engine	Description
	<ul style="list-style-type: none"> ■ there are one or more durable subscribers on the channel who have not yet consumed the events. <p>When there is no more "interest" for an event, the event will be deleted automatically.</p> <p>For further information on the JMS engine, see the section <i>Fanout Engine</i> in the Developer Guide.</p>
Merge Engine	<p>The merge engine is similar to the default engine, in that events are retained until they reach their TTL limit.</p> <p>The Merge Engine is used for Registered Events, and allows delivery of just the portion of an event's data that has changed, rather than of the entire event.</p> <p>In this scenario, the first time an event is published, it is published with all of its data fields. Thereafter, if another event occurs in which only a small number of the data fields have changed compared to the original published event, Universal Messaging allows you to send an event containing just the modified data fields. This process is known as <i>delta delivery</i>, and the events that contain just the modified fields are called <i>event deltas</i>.</p> <p>This pattern can be continued over a set of published events, so that each published event just contains the delta from the previous published event in the set. By publishing just delta events rather than full events, the amount of data being transferred on the channel can be considerably reduced.</p> <p>For further details on the merge engine and event deltas, see the appropriate language-specific section <i>The Merge Engine and Event Deltas</i> in the Developer Guide.</p>

The merge engine and the JMS engine are mutually exclusive: you can choose one but not both of these engines. If you choose neither of these engines, the default engine is used.

Channel Keys

Channels can also be created with a set of channel keys which define how channel events can be managed based on the content of the events. For more information, please see the section "[Channel Publish Keys](#)" on page 169.

Note on differences between Administration API and Enterprise Manager for Channel Creation

If you use the Administration API to create a channel rather than using the Enterprise Manager, some channel attributes must be set explicitly that in the Enterprise Manager are set implicitly according to context.

For example, the API offers a "cluster wide" flag, indicating that the new channel will be created on all cluster realm nodes. In the Enterprise Manager, this flag is not available directly, but will be applied automatically to the channel if you create the channel in the context of the cluster node rather than in the context of a realm node.

Storage Properties of Channels and Queues

Each channel or queue has a number of storage properties associated with it. You can set values for these properties when you create the channel or queue.

A summary of these properties can be seen below:

Property	Description
Auto Maintenance	Controls whether persistent store is maintained automatically (i.e. events reaching their TTL, or events which have been purged are cleared from the channel storage file.
Honour Capacity	<p>Controls whether the channel / queue capacity setting will prevent publishing of any more data once the channel / queue is full.</p> <p>If set to true, the client will get an exception on further publishes if a channel / queue is already full. A transactional publish will receive an exception on the commit call, a non-transactional publish will receive an asynchronous exception through the <code>nAsyncExceptionHandler</code>.</p> <p>If set to false, the oldest event in the channel / queue will be purged to make room for the newest event.</p> <p>Note: If this property is set to true, the capacity can nevertheless be exceeded by a transaction if the capacity had not yet been reached at the beginning of the transaction.</p>
Enable Caching:	Controls the caching algorithm within the server, if you set caching to false, all events will be read from the file store. If true then if server has room in memory, they will be stored in memory and reused.
Cache on Reload	When a server restarts it will scan all file based stores and check for corruption. During this test the default behaviour is to disable caching to conserve memory, however, in some instances it would be better if the server had actually cached the events in memory for fast replay.

Property	Description
Enable Read Buffering	Controls the read buffering logic for the store on the server.
Enable Multicast	This activates multicast processing, thereby allowing multicast clients to receive events over multicast connections.
Read Buffer Size	If Read Buffering is enabled, then this function sets the size in bytes of the buffer to use.
Sync Each Write	Controls whether each write to the store will also call sync on the file system to ensure all data is written to the Disk
Sync Batch Size	Controls how often in terms of number of events to sync on the file system to ensure all data is written to the Disk
Sync Batch Time	Controls how often in terms of time elapsed to sync on the file system to ensure all data is written to the Disk
Fanout Archive Target	Controls whether all events fanned out are written to an archive
Priority	<p>Sets a default priority for all events in the channel. The priority is a numeric value in the range 0 (lowest priority) to 9 (highest priority).</p> <p>You can override the default priority for an event by setting its value explicitly.</p> <p>For more information about event priorities, see the section <i>Priority Messaging</i> in the corresponding language-specific section of the Developer Guide.</p>
Stamp Dictionary	<p>Configures whether events on this channel/queue should be stamped by the server. Possible values are:</p> <ul style="list-style-type: none"> ■ <code>nChannelAttributes.DICTIONARY_STAMPING_ENABLED</code> - event dictionary stamping is explicitly enabled

Property	Description
	<ul style="list-style-type: none"> ■ <code>nChannelAttributes.DICTIONARY_STAMPING_DISABLED</code> - event dictionary stamping is explicitly disabled ■ <code>nChannelAttributes.DICTIONARY_STAMPING_DEFAULT</code> - the server-wide configuration takes effect (default value in case the parameter is not set)
Events Per Spindle	Defines the maximum number of events allowed per file in a multi-file disk store for persistent channels or persistent queues. If all existing files are full and a new event is published, a new file will be created to append further events for this store. If this value is set to 0 then the store will persist all events in a single file.

Channel Publish Keys

Channels can be created with a set of Channel Publish Key objects, as well as the default attributes that define behaviour of a channel and the events on a channel.

Channel Keys allow a channel or queue to automatically purge old events when new events of the same "type" are received. Two events are of the same "type" if the value in their dictionary (*nEventProperties*) for the key defined as the **Key Name** are identical. The channel will store the specified number of most recent events whose values match the **Key Name**.

Channel Publish Keys enable the implementation of Last Value Caches. In a last value cache, only the most recent value for a given type of event is kept on the channel. In high-update situations, where only the most recent values are of interest, Channel Publish Keys can greatly improve efficiency in this way. By altering the depth associated with the channel publish key, a recent values cache (where a set number of the most recent events of the same type are stored) can also easily be implemented.

Using Channel Keys to Automatically Purge Redundant Data

For example, if you have a channel called `BondDefinitions` which should only contain the most recent event published for each `Bond`, you can enforce this automatically by using a channel key. This functionality vastly simplifies data publication, since the publisher will not have to check the value of data currently on the channel.

In the above example you would create a `BondDefinition` channel that has a Channel Key called `BONDNAME` with a depth of 1, as shown in the following Enterprise Manager screen shot:

Add channel to realm node1

Channel Attributes

Channel Name: BondDefintions

Channel Type: Persistent

Channel TTL:

Channel Capacity:

Parent Realm: node1

Use JMS Engine: ☐

Channel Keys

Select Key To Edit:

Key Properties

Key Name: BONDNAME

Depth: 1

Save Delete

OK Cancel

The channel key defines the key in the `nEventProperties` which identifies events as being of the same type if their value for this key match. In order to add a Channel Key, type the name of the key into the **Key Name** box on the dialog and click **Save**. If you want the key to have a depth of greater than 1 then click the up arrow adjacent to the **Depth** field or enter the number manually.

If this is configured, as soon as an event is published to the `BondDefinitions` channel with a Dictionary entry called `BONDNAME`, the server checks to see if there is another event with the same value for that key. For example, if an event is published with a dictionary containing a key of `BONDNAME` and value of `bondnameA` and there is already an event with `BONDNAME=bondnameA`, then the old event will be removed, and the new one will take its place as the latest definition for `bondnameA`.

Another example would be if you wanted to keep the latest definition and the 2 before it, then you would create the channel key with depth 3 (implying that a maximum of 3 events with the same value for key name `BONDNAME` can exist on the channel).

If you wanted to keep an archive of all `bondname` values that were published to the channel, you could add a join from the `BondDefinitions` channel to, for example, a `BondDefinitionsArchive` channel. On this channel the absence of a Channel Key called `BONDNAME` will mean that it will store all events that have been published to the `BondDefinitions` channel.

Queue Attributes

Universal Messaging queues provide 3 main attributes. Depending on the options chosen, these define the behaviour of the events published and stored by the Universal Messaging Realm Server. The availability of the events on a queue is defined by the chosen attributes of the queue upon creation.

Each of these attributes are described in the following sections.

Queue TTL

The TTL for a queue defines how long (in milliseconds) each event published to the queue will remain available to consumers. Specifying a TTL of 0 will mean that events will remain on the queue indefinitely. If you specify a TTL of 10000, then after each event has been on the queue for 10000 milliseconds, it will be automatically removed by the server.

Queue Capacity

The capacity of a queue defines the maximum number of events may remain on a queue once published. Specifying a capacity of 0 will mean that there is no limit to the number of events on a queue. If you specify a capacity of 10000, then if there are 10000 events on a queue, and another event is published to the queue, the 1st event will be automatically removed by the server.

Queue Type

Each queue has an associated queue type. A Universal Messaging queue can be one of the types shown in the following table. The difference between the types lies in the type of physical storage used for each type and the performance overhead associated with each type.

Queue Type	Description
Simple	Simple queues have their messages stored in the Universal Messaging Realm's own memory space supplying a high-speed queue type. The difference between a Simple and Reliable is the fact that the event ids are reset to 0 in a Simple queue whenever the Universal Messaging Server is restarted. Simple queues are not supported across realms in a cluster.
Reliable	Reliable queues have their messages stored in the Universal Messaging Realm's own memory space. The first fact that is implied is that the maximum number of bytes that all messages across all reliable queues within a Universal Messaging Realm is limited by the maximum heap size available to the Java Virtual

Queue Type	Description
	<p>Machine hosting that Realm. The second fact implied is that if the Universal Messaging Realm is restarted for any reason, all messages stored on reliable queues will be removed from the queue as a matter of policy. However, as Universal Messaging guarantees not to ever reuse event ids within a queue, new messages published in those queues will get assigned event ids incremented from the event id of the last message prior to the previous instance stopping.</p>
Persistent	<p>Persistent queues have their messages stored in the Universal Messaging Realm's persistent queue disk-based store. The persistent queue store is a high performance file-based store that uses one or more files for each queue on that Realm, thus facilitating migrating whole queues to different Realms.</p> <p>For information about setting up a multi-file disk store for persistent channels or persistent queues, see the description of the storage parameter <code>Events per Spindle</code> in the section "Storage Properties of Channels and Queues" on page 167.</p> <p>All messages published to a persistent queue will be stored to disk, hence it is guaranteed that they will be kept and delivered to subscribers until it is purged or removed as a result of a "Time to Live" (TTL) or capacity policy.</p> <p>Messages purged from a persistent queue are marked as deleted, however the store size will not be reduced immediately. If the store is backed by a single file, the store size will not be reduced until maintenance is performed on the queue using the Universal Messaging Enterprise Manager or an Administration API call. Stores backed by multiple files will be cleaned up through the removal of individual files once the events stored on a particular file have all been marked as deleted. This augments the high performance of the Universal Messaging Realm.</p>
Mixed	<p>Mixed queues allow users to specify whether the event is stored persistently or in memory, as well as the "Time To Live" (TTL) of the individual event. On construction of a Mixed queue, the TTL and Capacity can be set, and if the user supplies a TTL for an event this is used instead of the queue. Events stored persistently are either stored on a single file or on multiple files in the same manner as a persistent queue.</p>
Transient	<p>A transient queue is like a simple queue in that no event characteristics are stored persistently. In addition to this, data is only ever written to a transient queue when 1 or more consumers</p>

Queue Type	Description
	<p>are connected to the queue and are able to consume this data. Unlike the simple queue which stores event data in memory, transient queues do not store anything, not even in memory. Transient queues can be thought of as a relay between 1 or more publishers and 1 or more subscribers.</p> <p>Transient queues are not supported across realms in a cluster.</p>
Off-heap	<p>Off-heap queues, similar to reliable queues, store the events in memory, but this queue type uses off-heap memory. This allows the normal JVM heap memory to be left free for short lived events, while longer living events can be stored off the JVM heap. This reduces the work the garbage collector needs to do to manage these events since they are out of range of the garbage collector.</p>
Paged	<p>Paged queues allows users access to high speed, off-heap memory mapped files as the basis for the event storage. As for the other queue types, each queue has its own unique file mapping. While similar to the off-heap queue type, where the events are stored in the off heap memory, the paged queue store is also persistent, thereby allowing recovery of events that are written into this store.</p>

Events

A Universal Messaging Event (`nConsumeEvent`) is the object that is published to a Universal Messaging channel or queue. It is stored by the server and then passed to consumers as and when required.

Events can contain simple byte array data, or more complex data structures such as an Event Dictionary (see ["Event Dictionaries" on page 185](#)).

Named Objects, Shared Named Objects and Serial Named Objects

Universal Messaging provides the ability for the server to maintain state for the last event that was consumed by a consumer on a channel. By providing a unique name, you can create a so-called *named object* on a channel. Then, even if your application is stopped and restarted, you will only consume available events from the last event ID that the server stored as successfully consumed by that named object.

When using normal named objects (as opposed to *shared named objects* and *serial named objects*, see below for description), there is only one store of events and that is the

channel itself. Each event remains on the channel until all named objects have consumed (and acknowledged) that event, and then the event is removed from the channel.

Named objects can be persistent, i.e. the last event ID is written to disk, so that if the Universal Messaging Realm Server is restarted, the last event ID consumed is retrievable for each named object on a channel.

Background to using Named Objects

A Universal Messaging channel (in JMS terminology this would be a "topic") can have several different "engines", one of which is the JMS engine. The JMS engine means that the channel will only store events whilst they are required by a consumer. So if you have a slow synchronous consumer, the channel will store the events until they are pulled off. If you have a named object, the channel will store the events until they are acknowledged. If you use the normal fanout engine, the events will remain on the channel regardless of any named objects. Only capacity or TTL or purges will delete events.

One reason for using named objects is to facilitate recovery if a consumer is terminated. For example, if the consumer crashes or you turn off the computer and start up another one, the server will remember where that consumer finished off because it stored the event ID on the named object.

Comparison of Named Objects and Shared Named Objects

Normally, Universal Messaging maintains a store of events for each named object (known in JMS as a *durable subscription*) until a consumer has received the events. When the consumer informs Universal Messaging that it has received a certain event from the store, Universal Messaging deletes this event from the store and all other events in the store that have a lower event ID (i.e. were created at an earlier time). There can be no more than 1 consumer of a given named object at any time; if a consumer tries to access a named object that is already being accessed by another consumer, an error will be returned to the second consumer.

There are some cases in which it may be desirable for two or more consumers to access a named object at the same time. Universal Messaging offers a variation of a named object known as a *shared named object*.

A shared named object allows multiple consumers to access the same named object. Universal Messaging delivers events to the consumers in round-robin fashion. So, for example, if there are 10 events in the named object and two consumers who share the named object, then Universal Messaging will deliver events 1, 3, 5, 7, 9 to the first consumer and events 2, 4, 6, 8, 10 to the second consumer.

A typical use case is load balancing, where multiple servers are available to process a stream of events.

As for normal named objects, you can specify that the shared named object is "persistent" and/or "cluster wide". If the shared named object is persistent, Universal Messaging will retain it if there is a system restart. If the shared named object is cluster wide and it exists on a cluster wide channel, the named object will also be replicated across the cluster.

Comparison of Shared Named Objects and Serial Named Objects

Universal Messaging also offers another variation of a named object, known as a *serial named object*, which can be accessed by multiple consumers.

A serial named object allows multiple consumers to access the named object but only in a serial manner and not concurrently. This means that any single consumer will not be sent events while another has pending unacknowledged events. This allows load balancing between multiple consumers whilst also ensuring that events are processed in the order that they are published across all consumers.

Each consumer will receive a number of events up to a maximum of either its window size or the total remaining events on the durable.

When you have a single consumer, it will maintain active mode until a second consumer is added, at which point the server will cycle between them. This may result in more events than the window size being received by the first consumer.

Window Size of Shared and Serial Named Objects

The Universal Messaging client API offers the option of specifying a window size for a shared named object. The window size can be set independently for both synchronous and asynchronous subscribers which use shared named objects. The client API contains two methods that accept the window size as a parameter: one for creating an iterator for a channel and one for adding a new subscriber with `nEventListener`.

When specifying the window size, the session is by default not auto-acknowledge. This means that when the window size is reached, the client stops receiving more events if the subscription is asynchronous. When an iterator is created and the window size is exceeded, the client is asked to commit or roll back.

If the API parameter is not specified, `nConstants.setMaxUnackedEvents(x)` will be used as the default client window size for all asynchronous subscribers, if this value has been set before subscribing with listeners. If the named object is not a shared named object, an `nIllegalArgumentException` is thrown.

If you specify a negative window size for this API parameter, the default value stored in the server configurations will be used instead. For asynchronous subscribers this is "ClientQueueWindow", which is available in the "Cluster Config" group; for synchronous subscribers, this is "IteratorWindowSize", which is available in the "Fanout Values" configuration group.

Examples

You can set the size of the client window using a method from the public `nChannel` API: `channel.addSubscriber(nEventListener nel, nNamedObject name, String selector, int windowSize)`. A different window size can be configured for every asynchronous subscriber. Here is a simple code snippet for this:

```
nSessionAttributes sessionAttr = new nSessionAttributes("nsp://localhost:11000");
nSession session = nSessionFactory.create(sessionAttr);
session.init();
nChannelAttributes channelAttr = new nChannelAttributes("channel");
nChannel channel = session.createChannel(channelAttr);
```

```

boolean isPersistent = true;
boolean isClusterWide = false;
int startEventId = 0;
nDurableAttributes attr =
    nDurableAttributes.create(nDurableAttributes.nDurableType.Shared,
        "shared-named-object");
attr.setPersistent(isPersistent);
attr.setClustered(isClusterWide);
attr.setStartEID(startEventId);
nDurable namedObject = null;
// Need to check first in case shared durable already exists!
try {
    namedObject = channels.getDurableManager().get(attr.getName());
} catch (nNameDoesNotExistException ex) {
    namedObject = channels.getDurableManager().add(attr);
}
nEventListener listener = new nEventListener() {
    @Override
    public void go(nConsumeEvent evt) {
        System.out.println("Consumed event " + evt.getEventID());
    }
};
int windowSize = 4;
channel.addSubscriber(listener, namedObject, null, windowSize);

```

You can also set the size of the client window when creating `nChannelIterator` from a channel object: `channel.createIterator(nNamedObject name, String selector, int windowSize)`. Here is an example:

```

nSessionAttributes sessionAttr = new nSessionAttributes("nsp://localhost:11000");
nSession session = nSessionFactory.create(sessionAttr);
session.init();
nChannelAttributes channelAttr = new nChannelAttributes("channel");
nChannel channel = session.createChannel(channelAttr);
boolean isPersistent = true;
boolean isClusterWide = false;
int startEventId = 0;
nDurable namedObject = null;
// create the shared durable
nDurableAttributes attr =
    nDurableAttributes.create(nDurableAttributes.nDurableType.Shared,
        "shared-named-object");
attr.setPersistent(isPersistent);
attr.setClustered(isClusterWide);
attr.setStartEID(startEventId);
// Need to check first in case shared durable is already exist!
try {
    namedObject = channels.getDurableManager().get(attr.getName());
} catch (nNameDoesNotExistException ex) {
    namedObject = channels.getDurableManager().add(attr);
}
int windowSize = 4;
nChannelIterator iterator =
    channel.createIterator(namedObject, null, windowSize);

```

`nChannelIterator` supports an option for specifying an infinite window size, namely `channel.createIterator(namedObject, null, nChannelIterator.INFINITE_WINDOW)`.

Storage Considerations

There are some storage considerations to be aware of when using shared named objects as opposed to normal (i.e. non-shared) named objects.

When you define a shared named object, Universal Messaging maintains a copy of all of the events of the named object in an internal store. The copies remain in the internal store until the original events have been accessed by the consumer. The benefit for the consumer is that events are held until the consumer is ready to receive them.

If the consumer accesses the events using event filtering, some events are never selected, and therefore the copies of these events remain in the internal store. As mentioned earlier, when an event is consumed for a normal named object, all earlier events are also removed; however, for a shared named object, only the consumed event is removed from the store. This is by design, and can, in time, lead to Universal Messaging retaining large numbers of events that are never used. If you have made the shared named object persistent, a copy of the internal store will be maintained on disk, and this disk copy will also contain copies of the events that are not used.

In order to delete messages from the internal store you can purge events from the channel that the named object is created on. That purge will also flow through to the internal store.

The set of events that are copied to the internal store for a given shared named object is limited according to the following rules:

- When the first consumer connects to the shared named object and specifies an event filter, then only events that match the filter will be copied to the internal store.
- If a second consumer connects to the same shared named object and also specifies an event filter, then from that time onwards, only events matching the second filter will be copied to the internal store. The first consumer still uses its filter to take events off the internal store.
- If the second consumer specifies no event filter, the filter of the first consumer remains in effect for copying events to the internal store.
- If the second consumer specifies a different filter than the first consumer, the first consumer will receive an asynchronous exception.

You can use the API method `getSharedNamedObjectOutstandingEvents()` on `nNamedObject` to get the number of events that are outstanding for the shared durable, i.e. the number of messages that have not yet been acknowledged by the shared durable and are therefore residing on the server. If you call this method on a non-shared durable, it will return the value -1.

In general, unless you really need to use shared named objects, we suggest you use normal named objects instead.

Event Filtering

Universal Messaging provides a server side filtering engine that allows events to be delivered to the client based on the content of values within the event dictionary.

This section introduces filtering and describes the basic syntax of the filtering engine, and provides examples to assist developers with designing the content of the events within Universal Messaging. The filtering capabilities described in this page are what is defined by the JMS standard.

Universal Messaging filtering can be applied at two levels. The first is between client and server and the second is between server and server.

All Universal Messaging filtering is handled by the Universal Messaging server and therefore significantly reduces client overhead and network bandwidth consumption.

For documentation on filtering functionality which extends beyond that available through the JMS standard please refer to the advanced filtering section (see ["Event Filtering using Selectors" on page 179](#)).

Basic Filtering

Each Universal Messaging event can contain an event dictionary as well as a byte array of data. Standard filtering, as defined by JMS, allows dictionary entries to be evaluated based on the value of the dictionary keys prior to delivering the data to the consumer.

The basic syntax of the filter strings is defined in the following notation :

EXPRESSION

where:

```

EXPRESSION ::=
<EXPRESSION> |
<EXPRESSION> <LOGICAL_OPERATOR> <EXPRESSION> |
<ARITHMETIC_EXPRESSION> |
<CONDITIONAL_EXPRESSION>
ARITHMETIC_EXPRESSION ::=
<ARITHMETIC_EXPRESSION> <ARITHMETIC_OPERATOR> <ARITHMETIC_EXPRESSION> |
<ELEMENT> <ARITHMETIC_OPERATOR> <ARITHMETIC_EXPRESSION> |
<ARITHMETIC_EXPRESSION> <ARITHMETIC_OPERATOR> <ELEMENT>
CONDITIONAL_EXPRESSION ::=
<ELEMENT> <COMPARISON_OPERATOR> <ELEMENT> |
<ELEMENT> <LOGICAL_OPERATOR> <COMPARISON_OPERATOR> <ELEMENT>
ELEMENT ::=
<DICTIONARY_KEY> |
<NUMERIC_LITERAL> |
<LOGICAL_LITERAL> |
<STRING_LITERAL> |
<FUNCTION>
LOGICAL_OPERATOR ::= NOT | AND | OR
COMPARISON_OPERATOR ::= <> | > | < | = | LIKE | BETWEEN | IN
ARITHMETIC_OPERATOR ::= + | - | / | *
DICTIONARY_KEY ::= The value of the dictionary entry with the specified key
LOGICAL_LITERAL ::= TRUE | FALSE
STRING_LITERAL ::= <STRING_LITERAL> <SEPARATOR> <STRING_LITERAL> |
                  Any string value, or if using LIKE,

```

```

        use the '_' wildcard character to denote exactly one character
        or the "%" wildcard character to denote zero or more characters.
NUMERIC_LITERAL ::= Any valid numeric value
SEPARATOR ::= ,
FUNCTION ::= <NOW> | <EVENTDATA> | DISTANCE

```

The above notation thus gives rise to the creation of any of the following valid example selector expressions :

```

size BETWEEN 10.0 AND 12.0
country IN ('uk', 'us', 'de', 'fr', 'es' ) AND size BETWEEN 14 AND 16
country LIKE 'u_' OR country LIKE '_e_'
country LIKE 'Can%'
country NOT LIKE '%nada'
size + 2 = 10 AND country NOT IN ('us', 'de', 'fr', 'es')
size / 2 = 10 OR size * 2 = 20
size - 2 = 8
size * 2 = 20
price - discount < 10.0 AND ((discount / price) * price) < 0.4

```

Additional references for event filtering may be found within the JMS message selector section of the JMS standard.

Event Filtering using Selectors

Universal Messaging supports *standard selector based filtering* and some advanced filtering concepts which will be described here .

Content Sensitive Filtering

Each Universal Messaging event can contain an event dictionary and a tag, as well as a byte array of data. Standard filtering, as defined by JMS, allows dictionary entries to be evaluated based on the value of the dictionary keys prior to delivering the data to the consumer.

Universal Messaging also supports a more advanced form of filtering based on the content of the event data (byte array) itself as well as time and location sensitive filtering. Universal Messaging also supports filtering based on arrays and dictionaries contained within event dictionaries. There is no limit to the dept of nested properties that can be filtered.

Filtering based on nested arrays and dictionaries

An event dictionary can contain primitive types as well as dictionaries. They can also contain arrays of primitive types and arrays of dictionaries. Universal Messaging supports the ability to filter based on these nested arrays and dictionaries.

If an nEventProperties object contains a key called NAMES which stores a String[] then it is possible to specify a filter that will only deliver events that match based on values within the array.

```
NAMES[] = 'myname'
```

- Returns events where any element in the NAMES array = 'myname'

```
NAMES[1] = 'myname'
```

- Returns events where the second element in the array = 'myname'

Similarly, if the array was an `nEventProperties[]` it would be possible to filter based on the values within the individual `nEventProperties` objects contained within the array.

For example if the event's `nEventProperties` contains a key called `CONTACTS` which stores an `nEventProperties[]` then the following selectors will be available.

```
CONTACTS[].name = 'aname'
```

- Return events where the `CONTACTS` array contains an `nEventProperties` which contains a key called `name` with the value 'aname'

```
CONTACTS[1].name = 'aname'
```

- Return events where the second element in the `CONTACTS` array of `nEventProperties` contains a key called `name` with the value 'aname'

```
CONTACTS[].NAMES[] = 'myname'
```

- Return events where the `CONTACTS` array contains a `NAMES` arrays with a value 'myname' somewhere in the `NAMES` array.

EventData Byte[] Filtering

Universal Messaging's filtering syntax supports a keyword called 'EVENTDATA' that corresponds to the actual byte array of data within the Universal Messaging event. There are a number of operations that can be performed on the event data using this keyword.

This enables a reduction in the amount of data you wish to send to clients, since rather than querying pre-determined dictionary values, you can now query the actual data portion of the event itself without having to provide dictionary entries. If you have a message structure and part of this structure includes the length of each value within the structure, then you can refer to each portion of data. Alternatively if you know the location of data within your byte array, these can be used for filtering quite easily.

Below is a list of the available operations that can be performed on the `EVENTDATA`.

```
EVENTDATA.LENGTH()
```

- Returns the length of the `byte[]` of the data in the event.

```
EVENTDATA.AS-BYTE(offset)
```

- Returns the byte value found within the data at the specified offset.

```
EVENTDATA.AS-SHORT(offset)
```

- Returns a short value found within the data at the specified offset. Length of the data is 2 bytes.

■

```
EVENTDATA.AS-INT(offset)
```

- Returns an int value found within the data at the specified offset. Length of the data is 4 bytes.

```
EVENTDATA.AS-LONG(offset)
```

- Returns a long value found within the data at the specified offset. Length of the data is 8 bytes.

```
EVENTDATA.AS-FLOAT(offset)
```

- Returns a float value found within the data at the specified offset. Length of the data is 4 bytes.

```
EVENTDATA.AS-DOUBLE(offset)
```

- Returns a double value found within the data at the specified offset. Length of the data is 8 bytes.

```
EVENTDATA.AS-STRING(offset, len)
```

- Returns a String value found within the data at the specified offset for the length specified.

```
EVENTDATA.TAG()
```

- Returns the String TAG of the event if it has one.

For example, we could provide a filter string of the following form if we know that at position 0, the first 2 bytes are a string that represents a value you wish to filter on:

```
EVENTDATA.AS-STRING(0, 2) = 'UK'
```

You can use the LIKE operator in the filter expression, for example:

```
EVENTDATA.AS-STRING(0, 2) LIKE 'U_'
```

Time Sensitive Filtering

Universal Messaging's filtering syntax also supports a function called 'NOW()' that is evaluated at the server as the current time in milliseconds using the standard Java time epoch. This function enables you to filter events from the server using a time sensitive approach. For example, if your data contained a dictionary key element called 'DATE_SOLD' that contained a millisecond value representing the data when an item was sold, one could provide a filter string on a subscription in the form of:

```
DATA_SOLD < (NOW() - 86400000)
```

Which would deliver events corresponding to items sold in the last 24 hours. This is a powerful addition to the filtering engine within Universal Messaging.

Location Sensitive Filtering

Universal Messaging's filtering engine supports a keyword called DISTANCE. This keyword is used to provide geographically sensitive filtering. This allows the calculation of the distance between a geographical reference point defined by you (a point on the earth's surface as defined by the latitude and longitude values) and the geographical location of the events contained in the event stream that you are filtering. The filter will find all events whose distance from the geographical reference point matches the given comparison criterion.

For example, if you were designing a system that tracked the location of a tornado, as the tornado moved position, the latitude and longitude would correspond to the geographical location on the earth's surface. As the position of a tornado changed, an event would be published containing the new latitude and longitude values as keys

within the dictionary ('latitude' and 'longitude' respectively). Using this premise, you could provide a filter in the form of:

```
DISTANCE (Latitude, Longitude, Units) Operator Value
```

where :

- Latitude : the floating point value representing the latitude of your geographical reference point
- Longitude : the floating point value representing the longitude of your geographical reference point
- Units : an optional string indicating the unit of measurement to be used in the calculation
 - K: Kilometers < Default >
 - M : Miles
 - N : Nautical Miles
- Operator: a numerical comparison operator, such as <, >, =
- Value: The required distance of the event's location from the geographical reference point.

For example :

```
DISTANCE ( 51.50, 0.16, M ) < 100
```

This will deliver events that correspond to tornadoes that are less than 100 miles away from the geographical reference point given by the latitude and longitude values (51.50, 0.16).

The DISTANCE keyword provides a valuable and powerful extension to Universal Messaging's filtering capabilities, if you require information that is sensitive to geographical locations.

Event Filtering using Subjects and Principals

Universal Messaging provides several mechanisms for directing events to specific consumers:

- *Subject-based filtering*: This enables publishers to direct events to specific consumers based on the 'subject' of their connection.
- *Principal-based filtering*: This allows one or more string attributes called principals to be associated with your nSession at initialization. Publishers can also publish events with one or more principals, and your nSession will receive events that have at least one principal that matches one of your nSession's principals.

Subject-based Filtering

Universal Messaging's subject-based filtering enables publishers to direct events to specific consumers based on the 'subject' of their connection. For example, if you initiate

a Universal Messaging nSession with the subject / session name of "mrsmith", and your nSession is then used to subscribe to a channel called '/deals', any publisher on the '/deals' channel can publish an nConsumeEvent onto the '/deals' channel directly to "mrsmith" by calling the following:

```
nConsumeEvent newdeal = new nConsumeEvent("New Deal",
    "Free Holiday In Barbados".getBytes());
newdeal.setSubscriberName("mrsmith".getBytes());
channel.publish(newdeal);
```

This is a powerful concept and one which is used heavily in request/response scenarios, as well as in scenarios for directly streaming events to individuals as required.

Limitations of subject-based filtering

Subject-based filtering has some limitations. If you consider large deployments of Universal Messaging where potentially there are multiple instances of 'services' which connect to Universal Messaging, and where the services would typically use the same UID for the nSession subject, it is difficult to distinguish between specific instances of a service, or have any kind of logical grouping of these services, as they will all appear as the same connected UID to the Universal Messaging server. One example of this would be the use of the Universal Messaging API by webMethods Integration Server (IS), where there could be multiple instances of IS, and within each instance of IS, multiple triggers that consume events from the same channel or queue. In such a case, there would be no possibility to distinguish between specific instances of IS, or specific triggers, or all triggers on an IS instance, since they are all running using the same UID.

For this purpose, Universal Messaging provides a client API feature called *Principals* to extend the subject-based filtering mechanism. This feature is described below.

Principal-based Filtering

All Universal Messaging Enterprise client APIs support the concept of an nSession object, which is the container for the physical connection to a Universal Messaging server. During initialisation, a handshake mechanism establishes a subject, which defaults to the local user name on the host machine (typically using the format "user@host"), but which can also be overridden to contain any arbitrary string as the UID portion of the subject.

Universal Messaging allows you to treat the subject not just as a single string but as a set of substrings called principals. These in turn can be used as filterable attributes for the nSession.

The Universal Messaging server maintains the subject in its own server-side connection that corresponds to the client nSession, and so any subject and principal list is available at the Universal Messaging server.

The nSession includes a method that enables multiple string principals to be passed during authentication. In addition, the subject filtering mechanism in the Universal Messaging server also checks the list of provided principals during the filtering based on the subscriber name(s) passed in any published events.

So you can do the following:

```
String[] group1 = {"trigger1", "IS1"};
```

```
nSessionAttributes sAttr1 = new nSessionAttributes(rname);
nSession subSession1 = nSessionFactory.create(sAttr1, "subscriber1");
subSession1.setAdditionalPrincipals(group1);
subSession1.init();
String[] group2 = {"trigger2", "IS1"};
nSessionAttributes sAttr2 = new nSessionAttributes(rname);
nSession subSession2 = nSessionFactory.create(sAttr1, "subscriber2");
subSession2.setAdditionalPrincipals(group2);
subSession2.init();
```

The nSessions created above now contain several principals. Similarly, a publisher can publish events in the following way:

```
nConsumeEvent event = new nConsumeEvent(group1[0].getBytes(),
    group1[0].getBytes());
event.setSubscriberName(group1[0]);
pubchan.publish(event);
nConsumeEvent event2 = new nConsumeEvent(group1[1].getBytes(),
    group1[1].getBytes());
event2.setSubscriberName(group1[1]);
pubchan.publish(event2);
```

You can see that the first event should go only to subSession1 / subscriber1, since the subscriberName equals the principal "trigger1", and the second will go to both subSession1 and subSession2, since they both have been created with a principal list that contains "IS1".

One extension to this is the ability to send an event to multiple subscriber names, which means that published events can also be delivered to multiple sessions with a matching principal in their list. For example, consider the following additional code:

```
String[] group3 = {"trigger3", "IS2"};
String[] group4 = {"trigger4", "IS2"};
String[] group5 = {"trigger1", "trigger2", "trigger3", "trigger4"};
nSessionAttributes sAttr3 = new nSessionAttributes(rname);
nSession subSession3 = nSessionFactory.create(sAttr3, "subscriber3");
subSession3.setAdditionalPrincipals(group3);
subSession3.init();
nSessionAttributes sAttr4 =
    new nSessionAttributes(rname2);
nSession subSession4 = nSessionFactory.create(sAttr4, "subscriber4");
subSession4.setAdditionalPrincipals(group4);
subSession4.init();
```

So a publisher can now publish the following events:

```
nConsumeEvent event3 = new nConsumeEvent("All", "All".getBytes());
event3.setSubscriberNames(group5);
pubchan.publish(event3);
nConsumeEvent event4 = new nConsumeEvent("IS1 and IS2",
    "IS1 and IS2".getBytes());
event4.setSubscriberNames(new String[]{group1[1], group3[1]});
pubchan.publish(event4);
```

Here, event3 will be delivered to all 4 sessions, since all 4 'triggerX' principal names have been added to the event, and event4 will also be delivered to all 4 sessions, since both 'IS1' and 'IS2' are added to the subscriber names. Therefore, all sessions have a matching principal for both events.

This very simple concept greatly enhances the ability to direct events to single connections or to a logical grouping of connections.

Event Dictionaries

Universal Messaging Event Dictionaries (nEventProperties) provide an accessible and flexible way to store any number of message properties for delivery within an event (for related information, see ["Events" on page 173](#)).

Event Dictionaries are quite similar to a hash table, supporting primitive types, arrays, and nested dictionaries.

Universal Messaging filtering allows subscribers to receive only specific subsets of a channel's events by applying the server's advanced filtering capabilities to the contents of each event's dictionary.

Event dictionaries can facilitate the automated purging of data from channels through the use of Publish Keys.

Durable Consumers

Universal Messaging provides the ability for both asynchronous and synchronous consumers to be durable. Durable consumers allow state to be kept at the server with regard to what events have been consumed by a specific consumer of data.

Universal Messaging supports durable consumers through use of Universal Messaging named objects. When a subscription is created using a named object, the server will ensure that the named object current position is maintained. As named object subscriptions are restarted, say after application restart, the server will begin delivering events from the last event which was successfully acknowledged by the client.

There are different ways in which events consumed by named consumers can be acknowledged. By specifying that 'auto acknowledge' is true when constructing either the synchronous or asynchronous consumers, then each event is acknowledged as consumed automatically. If 'auto acknowledge' is set to false, then each event consumed has to be acknowledged by calling the ack() method on the nConsumeEvent.

Durable Event Browsing and Purging

Durable Event Browsing

You can browse the un-acknowledged events on all durable types (Exclusive, Priority, SharedQueued, Shared).

To do this, you need the durable object associated with the channel you want to browse, which you can get as follows:

```
nDurable durable = myChannel.getDurableManager().get("your durable name");
nDurableViewer durableViewer = durable.createViewer();
```

The nDurableViewer object has two methods:

- `next()` - Returns the next un-acknowledged `nConsumeEvent`. If there are no further events to be consumed at this time, the method returns null.
- `close()` - Closes the created `nDurableViewer`.

Durable Event Purging

You can purge the un-acknowledged events on a durable.

To purge events, you need the durable object again:

```
nDurable durable = myChannel.getDurableManager().get("your durable name");
```

The durable object offers the following methods:

- `removeAll` - Removes all events on the durable.
- `remove(long start , long end)` - Removes an event range on a durable.
- `remove(String filter)` - Removes events matching a filter on a durable.
- `remove(long eid)` - Removes a single event with corresponding Event ID from the durable.

Currently, the purge capability is only supported for Shared Durables (i.e. on the `nSharedDurable` implementation of the `nDurable` object). If any of the above listed `remove` methods is called on a non-Shared durable it will throw an `nIllegalStateException`.

Native Communication Protocols

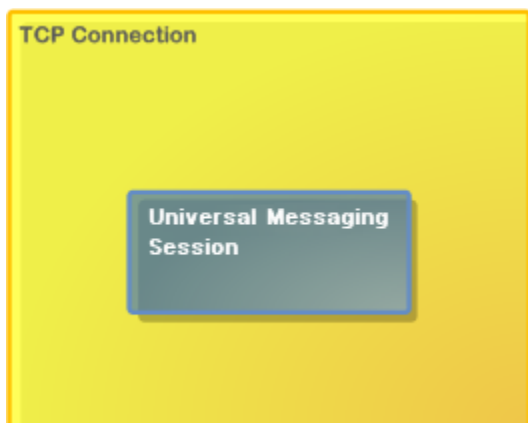
Universal Messaging supports four *Native Communication Protocols*. These TCP protocols are:

- Universal Messaging Socket Protocol (nsp)
- Universal Messaging SSL Protocol (nsps)
- Universal Messaging HTTP Protocol (nhp)
- Universal Messaging HTTPS Protocol (nhps)

These wire protocols are available for client-to-realm and realm-to-realm connections.

Universal Messaging Socket Protocol (nsp)

The Universal Messaging Socket Protocol (NSP) is a plain TCP socket protocol optimized for high throughput, low latency and minimal overhead.

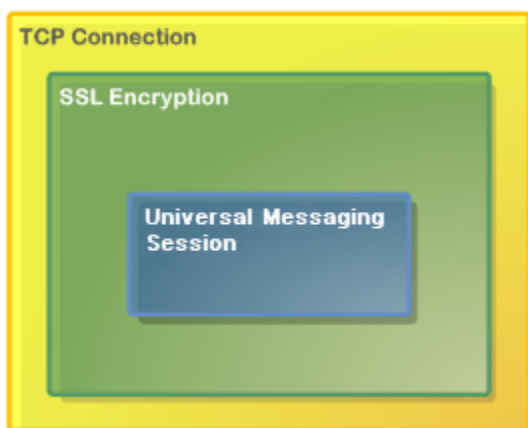


Universal Messaging Socket Protocol (nsp)

Universal Messaging SSL Protocol (nsps)

The Universal Messaging SSL (NSPS) Protocol uses SSL sockets to provide the benefits of the Universal Messaging Socket Protocol combined with encrypted communications and strong authentication.

We strongly recommend use of the NSPS protocol in production environments where data security is paramount.



Universal Messaging SSL Protocol (nsps)

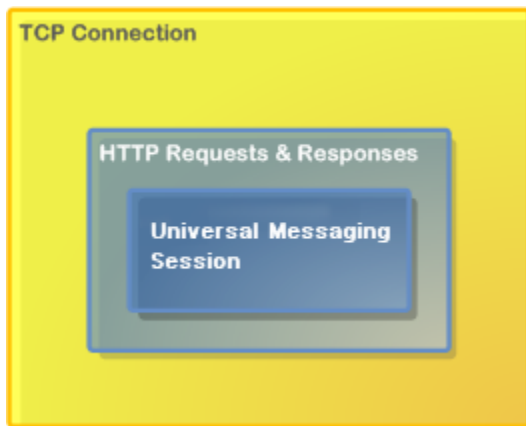
Universal Messaging HTTP Protocol (nhp)

The Universal Messaging HTTP (NHP) Protocol uses a native HTTP stack running over plain TCP sockets, to transparently provide access to Universal Messaging applications running behind single or multiple firewall layers.

This protocol was designed to simplify communication with Realms on private address range (NAT) networks, the Internet, or within another organization's DMZ.

There is no requirement for a web server, proxy, or port redirector on your firewall to take advantage of the flexibility that the Universal Messaging HTTP Protocol offers. The protocol also supports the use of HTTP proxy servers, with or without proxy user authentication.

An nhp interface will also support connections using the nsp protocol. For this reason it is suggested that you use this protocol initially when evaluating Universal Messaging.

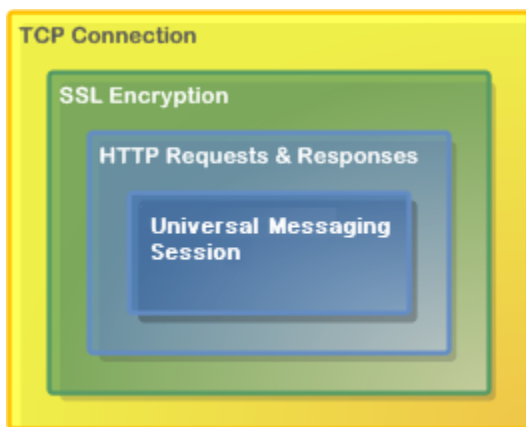


Universal Messaging HTTP Protocol (nhp)

Universal Messaging HTTPS Protocol (nhps)

The Universal Messaging HTTPS (NHPS) Protocol offers all the benefits of the Universal Messaging HTTP Protocol described above, combined with SSL-encrypted communications and strong authentication.

We strongly recommend use of the Universal Messaging HTTPS Protocol for production-level applications which communicate over the Internet or mobile networks.



Universal Messaging HTTPS Protocol (nhps)

Recommendation

We generally recommend that you initially use the *Universal Messaging HTTP Protocol (nhp)* for Universal Messaging Native clients, as this is the easiest to use and will support both nhp and nsp connections.

When deploying Internet-applications, we recommend the *Universal Messaging HTTPS Protocol (nhps)* for its firewall traversal and security features.

RNAMEs

The RNAME used by a Native Universal Messaging Client to connect to a Universal Messaging Realm server using a Native Communication Protocol is a non-web-based URL with the following structure:

```
<protocol>://<hostname>:<port>
```

where <protocol> can be one of the 4 available wire protocol identifiers:

- *nsp* (Universal Messaging Socket Protocol),
- *nspS* (Universal Messaging SSL Protocol),
- *nhp* (Universal Messaging HTTP Protocol) or
- *nhps* (Universal Messaging HTTPS Protocol)

An RNAME string consists of a comma-separated list of RNAMEs.

A Universal Messaging realm can have multiple network interfaces, each supporting any combination of Native and Comet communication protocols.

User@Realm Identification

When a Universal Messaging Native Client connects to a Universal Messaging Realm, it supplies the username of the currently-logged-on user on the client host machine. This username is used in conjunction with the hostname of the realm to create a session credential of the form `user@realm`.

For example if you are logged on to your client machine as user `fred`, and you specify an RNAME string of `nsp://realmserver.mycompany.com:9000`, then your session will be identified as `fred@realmserver.mycompany.com`.

Note, however, that if you were running the client application on the same machine as the Universal Messaging Realm and decided to use the `localhost` interface in your RNAME string, you would be identified as `fred@localhost` - which is a different credential.

The Realm and channel Access Control Lists (ACL) checks will be performed against this credential, so be careful when specifying an RNAME value.

Comet Communication Protocols

Universal Messaging supports Comet and WebSocket over two *Comet Communication Protocols*.

Streaming Comet, Long Polling or WebSocket

The Universal Messaging Comet API supports several both streaming and long polling Comet or WebSocket communications. A developer can select which method to use when starting a session with the JavaScript API.


Communication Protocols

HTTPS Protocol (https)

The Universal Messaging Comet HTTPS (SSL-encrypted HTTP) Protocol is a lightweight web-based protocol, optimized for communication over web infrastructure such as client or server-side firewalls and proxy servers.

This protocol simplifies communication between Universal Messaging Clients and Realms running behind single or multiple firewall layers or on private address range (NAT) networks. There is no requirement for an additional web server, proxy, or port redirector on your firewall to take advantage of the flexibility that the Universal Messaging HTTPS Protocol offers.

The protocol is fully SSL-encrypted and also supports the use of HTTP proxy servers, with or without proxy user authentication.




HTTPS Protocol (https)

HTTP Protocol (http)

The Universal Messaging Comet HTTP Protocol is a lightweight web-based protocol, supporting communication through proxies and firewalls at the client or server end of the network.

This protocol provides the same functionality as the Universal Messaging Comet HTTPS protocol, but without SSL encrypted communications.



HTTP Protocol (http)

Recommendation

We generally recommend the *HTTPS Protocol (https)* for Universal Messaging Comet clients, as this is both securely encrypted and easy to use.

RNAMEs

The RNAME used by a Universal Messaging Comet Client to connect to a Universal Messaging Realm server will automatically default to the same protocol/host/port as the web server from which an application is served, unless overridden by the developer when starting a session.

Note that a Universal Messaging realm can have multiple network interfaces, each supporting any combination of Native and Comet communication protocols.

Google Protocol Buffers

Overview

Google Protocol Buffers are a way of efficiently serializing structured data. They are language and platform neutral and have been designed to be easily extensible. The structure of your data is defined once, and then specific serialization and deserialization code is produced specifically to handle your data format efficiently.

Universal Messaging supports server-side filtering of Google Protocol Buffers, and this, coupled with Google Protocol Buffer's space-efficient serialization can be used to reduce the amount of data delivered to a client. If server side filtering is not required, the serialised protocol buffers could be loaded into a normal nConsume Event as the event data.

The structure of the data is defined in a .proto file, messages are constructed from a number of different types of fields and these fields can be required, optional or repeated. Protocol Buffers can also include other Protocol Buffers.

The serialization uses highly efficient encoding to make the serialized data as space efficient as possible, and the custom generated code for each data format allows for rapid serialization and deserialization.

Using Google Protocol Buffers with Universal Messaging

Google supplies libraries for Protocol Buffers in Java, C++ and Python, and third party libraries provide support for many other languages including .NET, Perl, PHP etc. Universal Messaging's client APIs provide support for the construction of Google Protocol Buffer events through which the serialized messages can be passed.

These nProtobufEvents are integrated seamlessly in Universal Messaging, allowing for server-side filtering of Google Protocol Buffer events, which can be sent on resources just like normal Universal Messaging Events. The server side filtering of messages is achieved by providing the server with a description of the data structures(constructed at the .proto compile time, using the standard protobuf compiler and the --descriptor_set_out option). The default location the server looks in for descriptor files is /plugins/ProtobufDescriptors and this can be configured through the Enterprise Manager. The server will monitor this folder for changes, and the frequency of these updates can be configured through the Enterprise Manager. The server can then extract the key value pairs from the binary Protobuf message, and filter message delivery based on user requirements.

To create an nProtobuf event, simply build your protocol buffer as normal and pass it into the nProtobuf constructor along with the message type used.

The Enterprise Manager can be used to view, edit and republish protocol buffer events, even if the EM is not running on the same machine as the server. To enable this, the server outputs a descriptor set to a configurable directory (by default the htdocs directory for the realm) and this can then be made available through a file plugin etc. The directory can be changed through the Enterprise Manager. The Enterprise Manager can then be configured to load this file using -DProtobufDescSetURL and then the contents of the protocol buffers can be parsed.

Using the Shared Memory Protocol

Universal Messaging supports a special kind of communication protocol called shm (*Shared Memory*). This communication protocol can only be used for intra host connectivity and uses physical memory to pass data rather than the network stack. Using shared memory as the communication protocol behaves just as any other nirvana communication protocol and therefore can be used within any Universal Messaging client or admin api application.

Once you have configured shared memory on your realm, it is ready to use by any Universal Messaging application you wish to run on the same host. All you need to do is set your RNAME to a the correct shared memory RNAME. For example, if you have configured shared memory to use /tmp, then your RNAME would be:

```
shm://localhost/tmp
```


To test this out, you could run any one of the example applications that are provided in the Universal Messaging download, by setting the RNAME from a Universal Messaging Java client examples prompt as described above. For example, a subscriber that subscribes to a channel called /test can be executed as follows:

```
nsubchan /test 0 1
```

Rollback

When an event is consumed from a queue transactionally, there is a chance that the transaction will be rolled back. This could be triggered by the user or by a disconnection. When this happens, the event needs to be put onto the queue for redelivery.

When a message is rolled back it is re-added to the front of the queue and is assigned a new event ID. So if an event is rolled back, that event will be added straight back at the front of the queue ready for redelivery. The API method `event.getHeader().getOriginalTransactionEID()` returns the event ID that the event had when it was first published to the queue; this value is only set on republish.

Note that for performance reasons Universal Messaging delivers a window of messages to the client. By default there will be 10 messages waiting in the client to be delivered to the user, so if the user rolls back event 0, they will receive the window of events already on the client before they consume event 0 again.

It is possible to set the client window to be only 1 event, and doing so will guarantee event ordering. The window size is set on the `nQueueReaderContext` and is only relevant for asynchronous consumers. Using a small window size affects the performance negatively, since this increases the number of windows Universal Messaging server has to deliver. Therefore, if your application does not require rolled-back events to be ordered, we suggest leaving the default windows size unaltered.

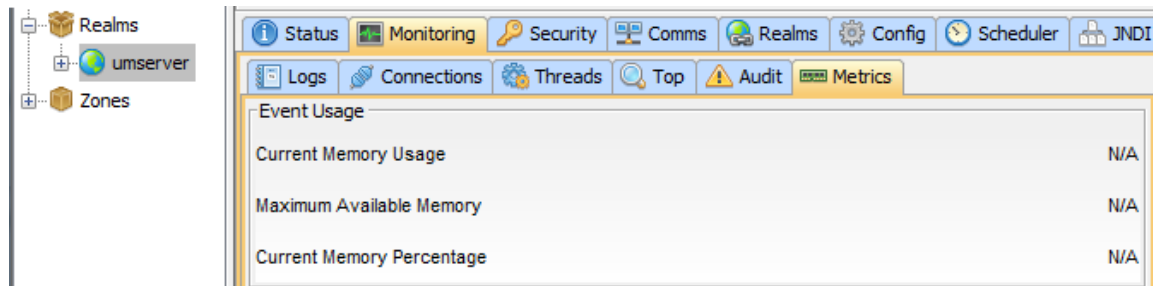
Note: Prior to version 9.9, the rolled back event was added to the back of the queue with a new event ID. The event was stamped with a redelivery count and `isRedelivered` set to true. If there were lots of events on the queue, adding the rolled back event to the back of the queue could mean it was not delivered for some time.

Out-of-Memory Protection

Universal Messaging provides methods that you can use to protect from out-of-memory situations. It is important to note that a subset of these methods may impact performance; those which do so are disabled by default.

Metric Management

Universal Messaging provides metrics pertaining to memory usage on a realm server. When you select the **Monitoring** tab in the Enterprise Manager, these metrics can be found under the **Metrics** tab, as shown below:



Metrics can be enabled/disabled on a realm server through the realm configuration. Setting the configuration option `EnableMetrics` controls the enabling/disabling of *all* metrics; each metric can also be enabled/disabled individually.

Event Usage

The event usage metric provides information on memory currently in use by on-heap events. The three statistics provided are as follows:

1. Current Memory Usage - the memory currently in use by on-heap events.
2. Maximum Available Memory - the maximum memory currently available to the JVM.
3. Current Memory Percentage - the percentage of the maximum memory available to the JVM currently occupied by on-heap events.

Event usage monitoring can be enabled/disabled through the configuration option `EnableEventMemoryMonitoring`. Enabling this metric may impact performance, and so it is disabled by default. Disabling this metric will cause "N/A" to be displayed for each of the three statistics described previously, as shown in the screenshot above.

Flow Control

Universal Messaging provides the ability to hold producing connections (publishes, transactional publishes, batch publishes, etc.) before processing their events. This is to allow consumers to reduce the number of events on connections before additional events are published, thus decreasing memory usage.

Flow control can be enabled/disabled and configured on a realm server through the realm configuration. Setting the configuration option `EnableFlowControl` controls the enabling/disabling of flow control; flow control is disabled by default. There are also three other configurable options:

- `FlowControlWaitTimeThree` - the first level of waiting, and shortest wait time; activated at 70% memory usage

- `FlowControlWaitTimeTwo` - the second level of waiting; activated at 80% memory usage
- `FlowControlWaitTimeOne` - the final level of waiting, and longest wait time; activated at 90% memory usage

Each of these options can be configured to suit the level of out-of-memory protection required. They represent the amount of time, in milliseconds, for which producer connections are throttled upon reaching the associated percentage of memory usage. In general, as memory usage gets closer to 100%, the need for throttling will be greater, so you would normally set `FlowControlWaitTimeOne` to a higher value than `FlowControlWaitTimeTwo`, which in turn will be higher than `FlowControlWaitTimeThree`.

Example: Assume a current memory usage of 90%. If there is a new publish request, the connection from which the publish originates will be throttled for `FlowControlWaitTimeOne` milliseconds, ensuring the event(s) will not be published sooner than `FlowControlWaitTimeOne` milliseconds from the time of publish.

Transactional Publishes

It is important to note that when using transactional publishes, by default the timeouts are shorter than standard publishes. This means that even if a client times out, the event may still be processed if the session has not been disconnected by the keep alive. As such, when using transactional publishes, it is best practice to call `publish` just once before committing. Every consecutive publish will be throttled, potentially leading to a timeout during the commit, and despite this timeout, the commit may still be successful.

Note: Even if flow control is enabled globally on the server, you can use Universal Messaging client API calls to bypass flow control for an individual producing connection. To do this, a boolean flag can be set on the relevant session attributes which requests that the associated connection should not be throttled. See the client API documentation on the class `nSessionAttributes` for more information.

Pause Publishing

The *pause publishing* feature allows all client publishing to the server to be paused, across channels, queues and data groups. This pause does not affect the administration API, inter-cluster communication or joins.

The feature is activated by setting the server configuration property `PauseServerPublishing` to true. Then, clients trying to publish or commit will receive `nPublishPausedException`.

See the section *Realm Configuration* in the *Developer Guide* for a description of the configuration property.

The sections below describe how publish methods are affected when publishing is paused.

JMS Publishing

Transactional

A commit will throw a `JMSEException` with a message containing "Publishing paused". The root exception will be `nPublishPausedException` which inherits `nSessionPausedException`.

The publish calls before the commit will throw an asynchronous `JMSEException` if an asynchronous exception listener is attached.

Note: The first call of a publish will not throw an exception because of an optimization, whereby the first published event is in fact not sent to the server until a commit or a second publish is in place.

Non-transactional send

The send calls throw an asynchronous `JMSEException` if an asynchronous exception listener is attached.

Note: The first call of a publish will not throw an exception because of an optimization, whereby the first published event is in fact not sent to the server until a commit or a second publish is in place.

Non-transactional SYNC send (`syncSendPersistent = true`)

The send calls will throw a synchronous `JMSEException`.

The root cause of the `JMSEException` is `nPublishPausedException`, and can be checked for example with this code:

```
} catch (JMSEException e) {
    Throwable rootCause = findRootCause(e);
    assertTrue(rootCause instanceof nPublishPausedException);
}
Throwable findRootCause(Throwable e){
    Throwable cause = e.getCause();
    return (cause != null ? findRootCause(cause) : e);
}
```

Native Java Publishing

Transactional

`nTransaction.commit` will result in `nPublishPausedException` which inherits `nSessionPausedException`.

The publish calls before the commit will throw an asynchronous `nPublishPausedException` if an asynchronous exception listener is attached.

Note: The first call of a publish will not throw an exception because of an optimization, whereby the first published event is in fact not sent to the server until a commit or a second publish is in place.

Non transactional

The publish calls will throw an asynchronous `nPublishPausedException` if an asynchronous exception listener is attached.

C++/C#

The exceptions for C++ and C# are the same as for the Java native API.

MQTT, AMQP and JavaScript

The exceptions for MQTT, AMQP and JavaScript when publishing is paused are handled as follows:

MQTT:

Currently, the MQTT protocol does not provide mechanisms to return proper errors to publishers. To handle publish exceptions (in particular, when publishing is paused), Universal Messaging provides two options: either ignore the exception, or disconnect the publisher's connection. This is configured using the `DisconnectClientsOnPublishFailure` MQTT configuration setting.

AMQP:

A standard AMQP error `amqp:internal-error` will be returned to the publisher client if publishing is paused. The error message starts with "Publishing paused", so adopters can use this substring to identify the actual cause

JavaScript:

If there is a publish or commit while publishing is paused, an error with a description starting with "Publishing paused" will be returned.