

Fineract Platform Documentation

Version 1.11.0

Table of Contents

License	1
Preface	2
Introduction	3
Platform for Digital Financial Services	3
About	3
Contribute	3
Downloads	3
Resources	3
Deployment	4
Plugins	4
HTTPS	4
Docker Compose	5
Application Server	5
Fineract Instance types	5
Kubernetes	10
AWS	10
Google Cloud	10
Apache Software Foundation Infrastructure	10
Architecture	11
History	11
Resources	11
System Overview	12
Functional Overview	13
Principles	14
Design Overview	16
Persistence	21
Idempotency	25
Validation	26
Batch execution and jobs	26
Loan account locking	35
Technology	36
Modules	37
Introducing Business Date into Fineract - Community version	37
Reliable event framework	44
Introducing Advanced payment allocation	60
Fineract Development Environment	66

Git	66
GPG	66
Docker	66
Gradle	66
IDE	66
Kubernetes	67
Tools	67
Custom Modules	69
Introduction	69
Custom Services	71
Custom Business Steps	74
Custom Loan Transaction Processors	76
Custom Batch Jobs	83
Custom Database Migration	89
Deploying Custom Modules	90
Outlook	91
Resilience	92
Introduction Resilience	92
Command	94
Jobs	97
Loan	99
Savings	100
Security	103
OAuth	103
Testing	105
Cucumber	105
Unit Testing	109
Integration Testing	109
Fineract Documentation Guide	110
File and Folder Layout	110
AsciiDoc	110
Diagrams	120
Editor	121
Antora	122
Releases	123
Configuration	123
Release Process	133
Maintenance Release Process	148

Publish Release Artifacts	149
Fineract SDKs	151
Generate Apache Fineract API Client	151
Frequently Asked Questions	153
Glossary	154
Index	155
Appendix A: Fineract Application Properties	156
Tenant Database Properties	156
Hikari Connection Pool Properties	158
SSL Properties	162
Authentication Properties	163
Tomcat Properties	164
Kafka Properties	165
Metrics Properties	168
AWS Configuration Properties	170
Resilience4j Properties	170
Appendix B: Third Party Software	175

License

Copyright 2022 Apache Foundation

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Preface

© 2015-2025

Apache Fineract

Website: fineract.apache.org

Email: dev@fineract.apache.org

Version: 1.11.0

Built on: Fri Feb 28 11:45:13 PST 2025

For document authors and changelog, view code history for the `fineract-doc` directory in github.com/apache/fineract/.

Introduction

Platform for Digital Financial Services

See fineract.apache.org

About

See fineract.apache.org/#about

Contribute

See fineract.apache.org/#contribute

Downloads

See fineract.apache.org/#downloads

Resources

See fineract.apache.org/#resources

Deployment

Plugins

Apache Fineract is extensible through plugin JARs ([FINERACT-1177](#); based on [Spring Boot's support](#)). To launch Fineract with plugin JARs in `libs/*.jar`, use:

```
java -Dloader.path=libs/ -jar fineract-provider.jar
```

The Fineract "Docker" container image's `ENTRYPOINT` uses this, see our [Dockerfile](#). You could therefore build your customized Fineract distribution container image with your own [Dockerfile](#) using e.g. `FROM apache/fineract:latest` and then drop some plugin JARs into `/app/libs/`.

The WAR distribution does not directly support such plugins, but one could "explode" the WAR and drop JARs into `WEB-INF/lib`; if you know what you are doing, and feel nostalgic of the 1990s still using WARs, instead of the recommended modern Spring Boot distribution.

Here is a list of known 3rd-party plugin projects which can be dropped into `libs/`:

- github.com/vorburger/fineract-pentaho



The reporting module became our first module experiment out of necessity. We are currently developing a strategy to split up even more internals of Fineract into proper modules. Those that have an incompatible license will be hosted in a separate Git repository (probably on Github under the Mifos organisation). We'll send out an announcement as soon as we have more to say on this topic.

HTTPS

Because Apache Fineract deals with customer sensitive personally identifiable information (PII), it very strongly encourages all developers, implementors and end-users to always only use HTTPS. This is why it does not run on HTTP even for local development and enforces use of HTTPS.

For this purpose, Fineract includes a built-in default SSL certificate. This cert is intended for development on `localhost`, only. It is not trusted by your browser (because it's self signed).

For production deployments, we recommend running Fineract behind a modern managed cloud native web proxy which includes SSL termination with automatically rotating SSL certificates, either using your favourite cloud provider's respective solution, or locally setting up the equivalent using e.g. something like NGINX combined with [Let's Encrypt](#).

Such products, when correctly configured, add the conventional `X-Forwarded-For` and `X-Forwarded-Proto` HTTP headers, which Fineract (or rather the Spring Framework really) correctly respects since

[FINERACT-914 was fixed.](#)

Alternatively, you could replace the built-in default SSL certificate with one you obtained from a Certificate Authority. We currently do not document how to do this, because we do not recommend this approach, as it's cumbersome to configure and support and less secure than a managed auto rotating solution.

The Fineract API client supports an insecure mode (`FineractClient.Builder#insecure()`), and API users such as mobile apps may expose Settings to let end-users accept the self signed certificate. This should always be used for testing only, never in production.

Docker Compose

TBD

Application Server

Tomcat

TBD

Undertow

TBD

Jetty

TBD

JBoss

TBD

Weblogic

TBD

Payara

TBD

Fineract Instance types

In cases where Fineract has to deal with high load, it can cause a performance problem for a single Fineract instance.

To overcome this problem, Fineract instances can be started in different instance types for better scalability and performance in a multi-instance environment:

Fineract instance types

- Read instance
- Write instance
- Batch instance

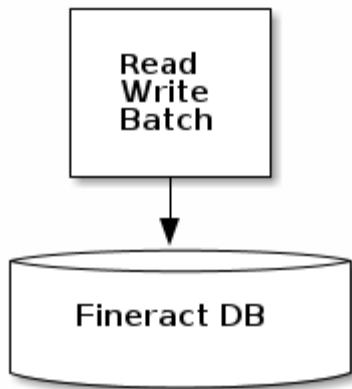
Each instance type comes with different restrictions. The specifics can be found in the table below.

Table 1. Instance types

	Read instance	Write instance	Batch instance
Using only read-only DB connection	Yes	No	No
Batch jobs are automatically scheduled or startable via API	No	No	Yes
Can receive events (business events, hook template events)	No	Yes	No
Can send events (business events, hook template events)	No	Yes	Yes
Read APIs supported	Yes	Yes	No
Write APIs supported	No	Yes	No
Batch job APIs supported	No	No	Yes
Liquibase migration initiated upon startup	No	Yes	No

Configuring instance types in single instance setup

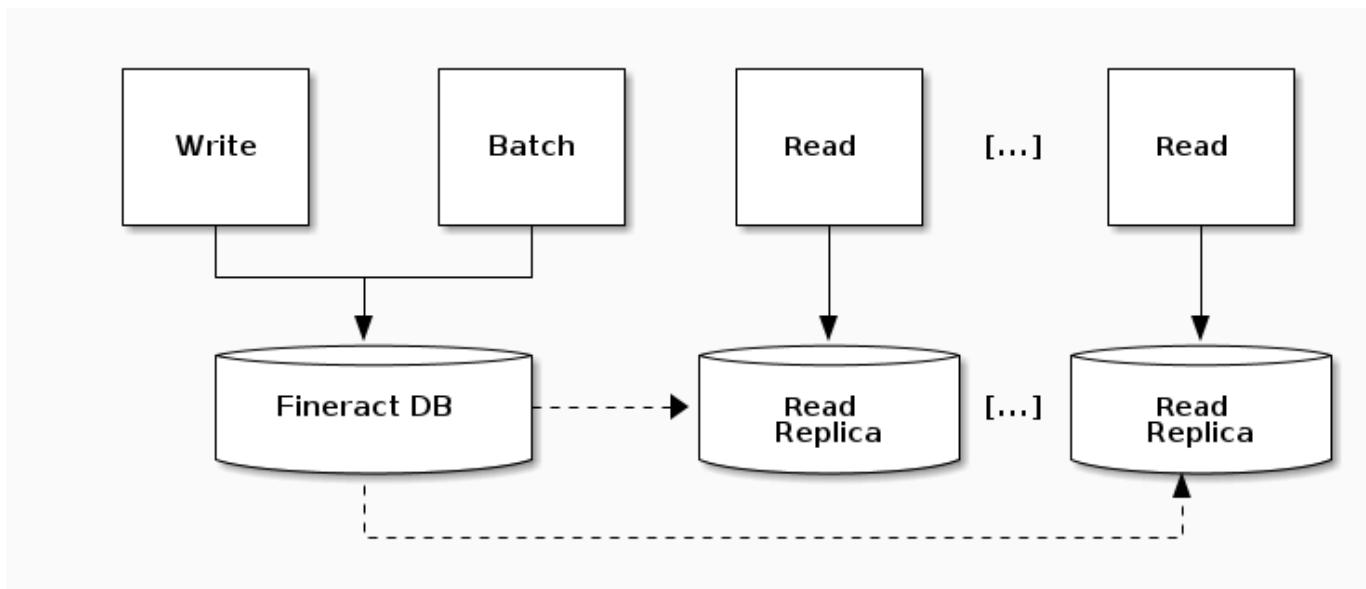
If Fineract is running as a single instance, then all of the 3 instance types should be enabled. In this case, there is no need to worry about the configuration, because this is the default behavior.



Configuring instance types in multi-instance setup

A common solution to dealing with the high load is to deploy 1 write and 1 batch instances and deploy multiple read instances with read replicas of the Fineract database.

In this case, the write instance and the database will be freed from part of the load, because read request will use the separated read instance and its read replica database.

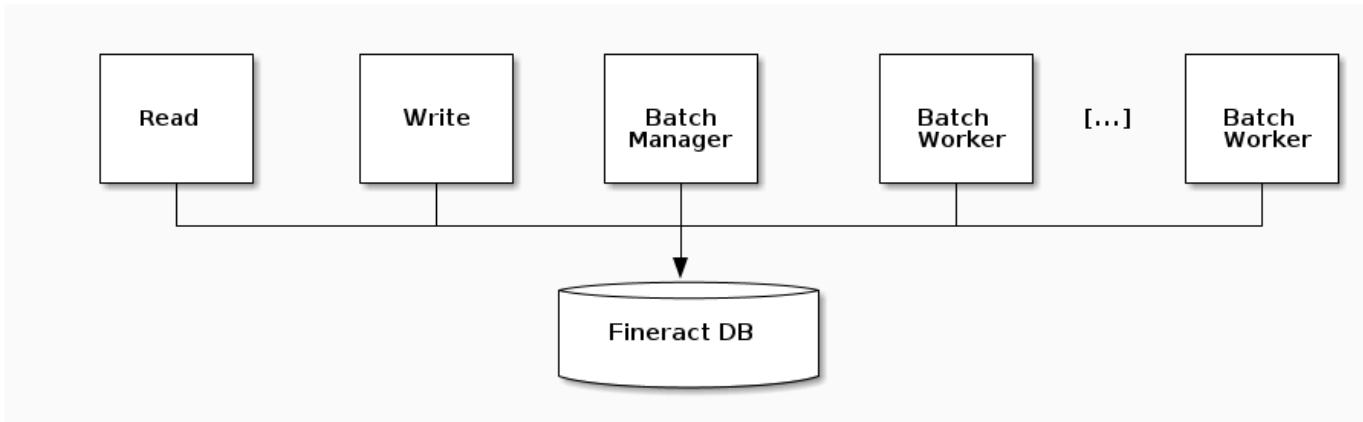


Also a common scenario when Close of Business jobs are running and Fineract has to deal with a high amount of processes.

(In a future release) Fineract (will be) is able to run this CoB jobs in batches.

In multi-instances environment these CoB jobs can run on multiple batch instances and they don't have any impact on the performance of the read and write processes.

The best practice is to deploy 1 master batch instance and multiple worker batch instances.



These solutions can be mixed with each other, based on the load of the Fineract deployment.

Configuring instance type via environment variables

The Fineract instance type is configurable via environment variables for the following 3 values:

Table 2. Environment variables

Instance type	Environment variable
Read instance	FINERACT_READ_MODE_ENABLED
Write instance	FINERACT_WRITE_MODE_ENABLED
Batch instance	FINERACT_BATCH_MODE_ENABLED

The environment variable values are booleans (true/false). The Fineract instance can be configured in any combination of these instance types, although if all 3 configurations are false, startup will fail. The default value for all 3 values is true.

The configured Fineract instance types are easily accessible via a single Spring bean, named `FineractProperties.FineractModeProperties` that has 4 methods: `isReadMode()`, `isWriteMode()`, `isBatchMode()`, `isReadOnlyMode()`

Liquibase Database Migration

Liquibase data migration is allowed only for write instances

APIs

Read APIs are allowed only for read and write instances

A Fineract instance is ONLY able to serve read API calls when it's configured as a read or write instance. In batch instance mode, it won't serve read API calls.

If it's a read or write instance, the read APIs will be served.

If it's a batch instance, the read APIs won't be served and a proper HTTP status code will be returned.

The distinction whether something is a read API can be decided based on the HTTP request method. If

it's a GET, we can assume it's a read call.

Write APIs are allowed only for write instances

A Fineract instance is ONLY able to serve write API calls when it's configured as a write instance. In read or batch instance mode, it won't serve write API calls.

If the write APIs won't be served and a proper HTTP status code will be returned.

If it's a write instance, the write APIs will be served except the ones related to batch jobs.

The distinction whether something is a write API can be decided based on the HTTP request method. If it's non-GET, we can assume it's a write call. Also, the write APIs related to batch jobs (starting/stopping jobs) will not be served either.

Batch job APIs are allowed only for batch instances

A Fineract instance is ONLY able to serve batch API calls when it's configured as a batch instance. In read or write instance mode, it won't serve batch API calls.

If the batch APIs won't be served and a proper HTTP status code will be returned.

If it's a batch instance, the batch APIs will be served.

Batch jobs

Batch job scheduling is allowed only for batch instances

Batch jobs are scheduled only if the Fineract instance running as a batch instance

Read-only instance type restrictions

If the read mode is enabled, but the write mode and batch mode are disabled, Fineract instance runs in read-only mode.

Events are disabled for read-only instances

When a Fineract instance is running in read-only mode, all event receiving/sending will be disabled.

Read-only tenant connection support

With read separation, there's a possibility to use read-only database connections for read-only instances.

If the instance is read-only, the DataSource connection used for the tenant will be read-only.

If the instance is read-only and the configuration for the read-only datasource is not set, the application startup will fail.

Batch-only instance type restrictions

If the batch mode is enabled, but the read mode and write mode are disabled, Fineract instance runs in batch-only mode.

Receiving events is disabled for batch-only instances

When a Fineract instance is running as batch, event receiving will be disabled while sending events will be still possible since the batch jobs are potentially generating business events.

Kubernetes

In a scaled Kubernetes environment where multiple Fineract instances are deployed, doing the database migrations properly is essential.

Fineract provides a way to run only the Liquibase migrations instead of starting up the whole application server so that you can easily do the migrations before actually upgrading a Fineract instance.

The `FINERACT_LIQUIBASE_ENABLED` flag controls whether Liquibase is enabled or not. For regular read/write/batch manager/batch worker instances this should be disabled.

There's a special Spring profile that should be enabled for running Liquibase only. It can be done via `SPRING_PROFILES_ACTIVE` environment variable. The profile name is `liquibase-only`. At the end of the migration process, the application will exit.

For the instance running the Liquibase migrations, the profile should be activated.

AWS

TBD

Google Cloud

The www.fineract.dev demo server runs on Google Cloud.

The [Running Fineract.dev, SRE style presentation](#) given at *ApacheCon 2020* has some related background.

Apache Software Foundation Infrastructure

We can order a server from Apache's infrastructure team and deploy a demo instance...

TBD

Architecture

This document captures the major architectural decisions in platform. The purpose of the document is to provide a guide to the overall structure of the platform; where it fits in the overall context of an MIS solution and its internals so that contributors can more effectively understand how changes that they are considering can be made, and the consequences of those changes.

The target audience for this report is both system integrators (who will use the document to gain an understanding of the structure of the platform and its design rationale) and platform contributors who will use the document to reason about future changes and who will update the document as the system evolves.

History

The Idea

Fineract was an idea born out of a wish to create and deploy technology that allows the microfinance industry to scale. The goal is to:

- Produce a gold standard management information system suitable for microfinance operations
- Acts as the basis of a platform for microfinance
- Open source, owned and driven by member organisations in the community
- Enabling potential for eco-system of providers located near to MFIs

Timeline

- 2006: Project initiated by Grameen Foundation
- Late 2011: Grameen Foundation handed over full responsibility to open source community.
- 2012: Mifos X platform started. Previous members of project come together under the name of Community for Open Source Microfinance (COSM / OpenMF)
- 2013: COSM / OpenMF officially rebranded to Mifos Initiative and receive US 501c3 status.
- 2016: Fineract 1.x began incubation at Apache

Resources

- Project URL is github.com/apache/fineract
- Issue tracker is issues.apache.org/jira/projects/FINERACT/summary
- Download from fineract.apache.org/

System Overview

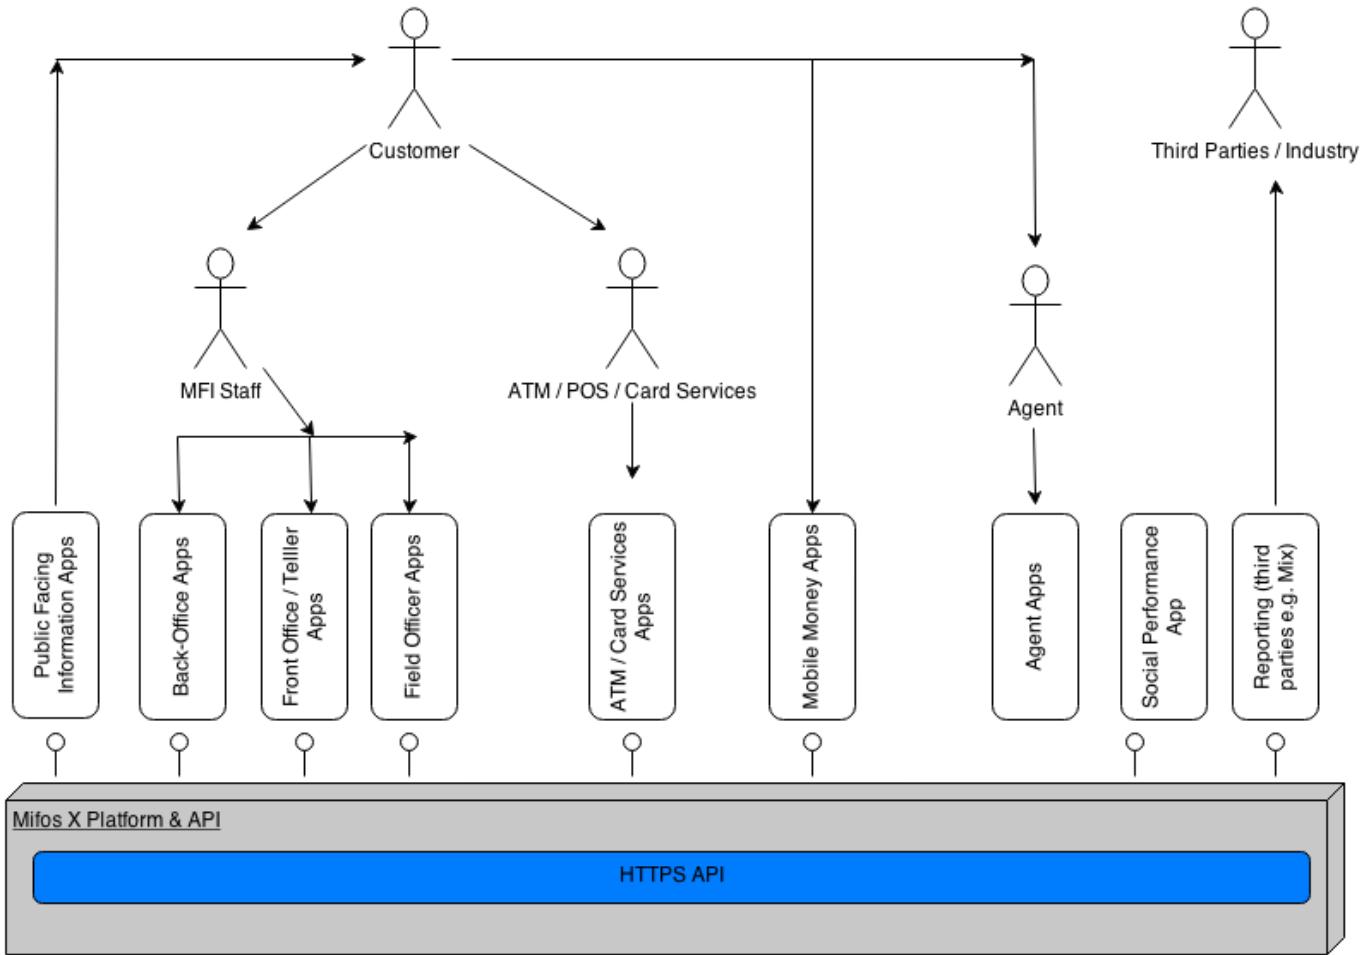


Figure 1. Platform System Overview

Financial institutions deliver their services to customers through a variety of means today.

- Customers can call direct into branches (teller model)
- Customers can organise into groups (or centers) and agree to meetup at a location and time with FI staff (traditional microfinance).
- An FI might have a public facing information portal that customers can use for variety of reasons including account management (online banking).
- An FI might be integrated into a ATM/POS/Card services network that the customer can use.
- An FI might be integrated with a mobile money operator and support mobile money services for customer (present/future microfinance).
- An FI might use third party agents to sell on products/services from other banks/FIs.

As illustrated in the above diagram, the various stakeholders leverage business apps to perform specific customer or FI related actions. The functionality contained in these business apps can be bundled up and packaged in any way. In the diagram, several of the apps may be combined into one

app or any one of the blocks representing an app could be further broken up as needed.

The platform is the core engine of the MIS. It hides a lot of the complexity that exists in the business and technical domains needed for an MIS in FIs behind a relatively simple API. It is this API that frees up app developers to innovate and produce apps that can be as general or as bespoke as FIs need them to be.

Functional Overview

As ALL capabilities of the platform are exposed through an API, The API docs are the best place to view a detailed breakdown of what the platform does. See online API Documentation.

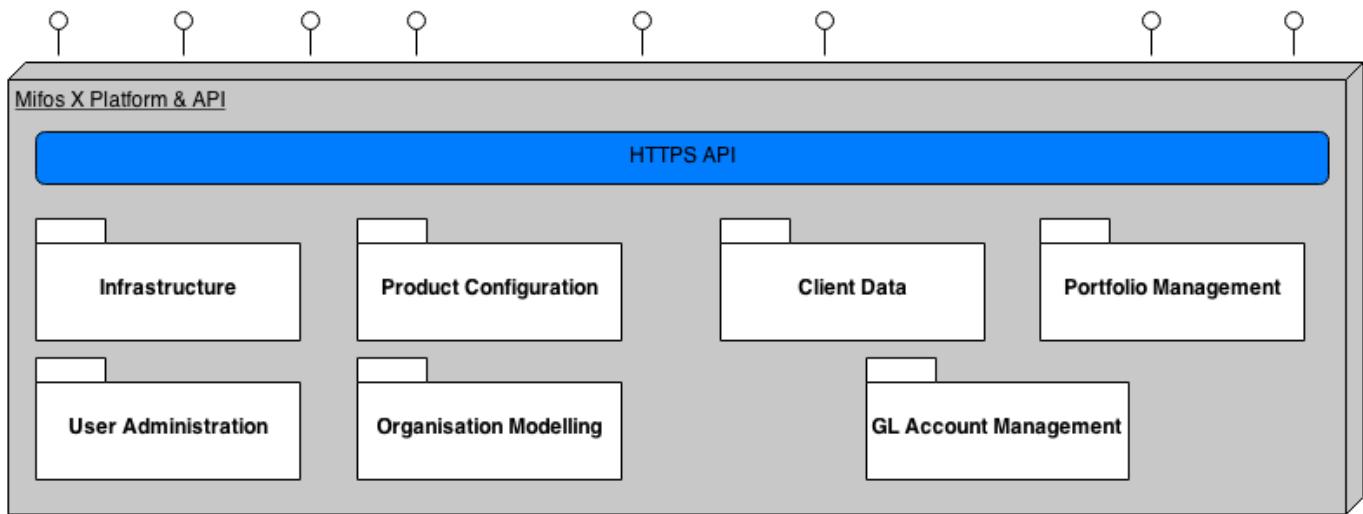


Figure 2. Platform Functional Overview

At a higher level though we see the capabilities fall into the following categories:

- Infrastructure
 - Codes
 - Extensible Data Tables
 - Reporting
- User Administration
 - Users
 - Roles
 - Permissions
- Organisation Modelling
 - Offices
 - Staff
 - Currency

- Product Configuration
 - Charges
 - Loan Products
 - Deposit Products
- Client Data
 - Know Your Client (KYC)
- Portfolio Management
 - Loan Accounts
 - Deposit Accounts
 - Client/Groups
- GL Account Management
 - Chart of Accounts
 - General Ledger

Principles

RESTful API

The platform exposes all its functionality via a **practically-RESTful API**, that communicates using JSON.

We use the term **practically-RESTful** in order to make it clear we are not trying to be fully REST compliant but still maintain important RESTful attributes like:

- Stateless: platform maintains no conversational or session-based state. The result of this is ability to scale horizontally with ease.
- Resource-oriented: API is focussed around set of resources using HTTP vocabulary and conventions e.g GET, PUT, POST, DELETE, HTTP status codes. This results in a simple and consistent API for clients.

See online API Documentation for more detail.

Multi-tenanted

The Fineract platform has been developed with support for multi-tenancy at the core of its design. This means that it is just as easy to use the platform for Software-as-a-Service (SaaS) type offerings as it is for local installations.

The platform uses an approach that isolates an FIs data per database/schema (See Separate Databases and Shared Database, Separate Schemas).

Extensible

Whilst each tenant will have a set of core tables, the platform tables can be extended in different ways for each tenant through the use of Data tables functionality.

Command Query Separation

We separate **commands** (that change data) from **queries** (that read data).

Why? There are numerous reasons for choosing this approach which at present is not an attempt at full blown CQRS. The main advantages at present are:

- State changing commands are persisted providing an audit of all state changes.
- Used to support a general approach to **maker-checker**.
- State changing commands use the Object-Oriented paradigm (and hence ORM) whilst queries can stay in the data paradigm.

Maker-Checker

Also known as **four-eyes principal**. Enables apps to support a maker-checker style workflow process. Commands that pass validation will be persisted. Maker-checker can be enabled/disabled at fine-grained level for any state changing API.

Fine grained access control

A fine grained permission is associated with each API. Administrators have fine grained control over what roles or users have access to.

Package Structure

The intention is for platform code to be packaged in a vertical slice way (as opposed to layers).

Source code starts from github.com/apache/fineract/tree/develop/fineract-provider/src/main/java/org/apache/fineract

- accounting
- useradministration
- infrastructure
- portfolio
 - charge
 - client
 - fund
 - loanaccount
- accounting

Within each vertical slice is some common packaging structure:

- api - XXXApiResource.java - REST api implementation files
- handler - XXXCommandHandler.java - specific handlers invoked
- service - contains read + write services for functional area
- domain - OO concepts for the functional area
- data - Data concepts for the area
- serialization - ability to convert from/to API JSON for functional area

Design Overview



The implementation of the platform code to process commands through handlers whilst supporting maker-checker and authorisation checks is a little bit convoluted at present and is an area pin-pointed for clean up to make it easier to on board new platform developers. In the mean time below content is used to explain its workings at present.

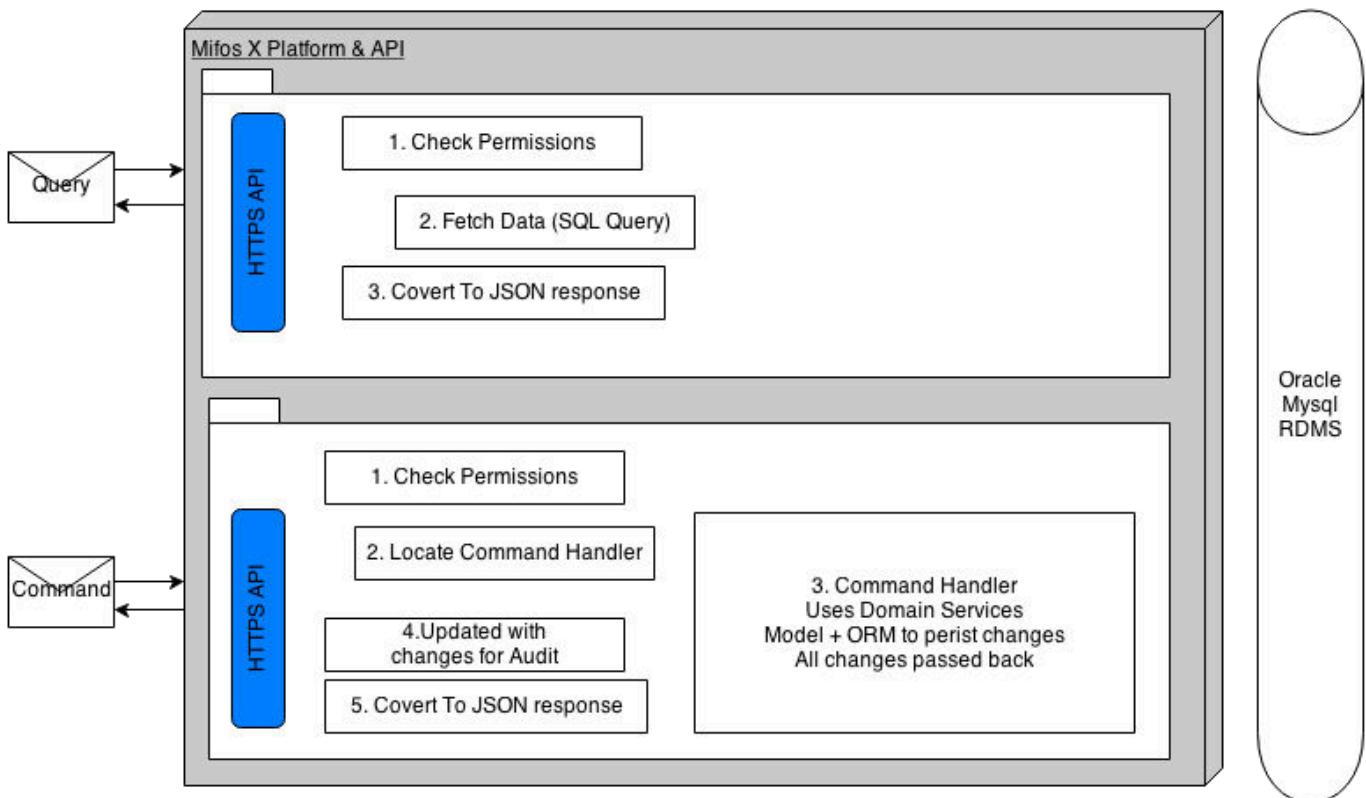


Figure 3. CQRS

Taking into account example shown above for the **users** resource.

- Query: GET /users

- HTTPS API: `retrieveAll` method on **org.apache.fineract.useradministration.api.UsersApiResource** invoked
- `UsersApiResource.retrieveAll`: Check user has permission to access this resources data.
- `UsersApiResource.retrieveAll`: Use 'read service' to fetch all users data ('read services' execute simple SQL queries against Database using JDBC)
- `UsersApiResource.retrieveAll`: Data returned to converted into JSON response
- Command: POST /users (Note: data passed in request body)
- HTTPS API: `create` method on **org.apache.fineract.useradministration.api.UsersApiResource** invoked

UsersApiResource.create

```

        return this.toApiJsonSerializer.serialize(result);
    }

    @PUT
    @Path("{userId}")
    @Operation(summary = "Update a User", description = "When updating a password you
must provide the repeatPassword parameter also.")
    @RequestBody(required = true, content = @Content(schema = @Schema(implementation =
UsersApiResourceSwagger.PutUsersUserIdRequest.class)))
    @ApiResponses({
        @ApiResponse(responseCode = "200", description = "OK", content = @Content
(schema = @Schema(implementation = UsersApiResourceSwagger.PutUsersUserIdResponse.class)
))
    })
    @Consumes({ MediaType.APPLICATION_JSON })
    @Produces({ MediaType.APPLICATION_JSON })
    public String update(@PathParam("userId") @Parameter(description = "userId") final
Long userId,
        @Parameter(hidden = true) final String apiRequestBodyAsJson) {

        final CommandWrapper commandRequest = new CommandWrapperBuilder() //
.updateUser(userId) //
.withJson(apiRequestBodyAsJson) //
.build();

        final CommandProcessingResult result = this.commandsSourceWritePlatformService
.logCommandSource(commandRequest);
    }
}

```

Create a CommandWrapper object that represents this create user command and JSON request body. Pass off responsibility for processing to PortfolioCommandSourceWritePlatformService.logCommandSource

```
validateIsUpdateAllowed();
```

```

        final String json = wrapper.getJson();
        final JsonElement parsedCommand = this.fromApiJsonHelper.parse(json);
        JsonCommand command = JsonCommand.from(json, parsedCommand, this
        .fromApiJsonHelper, wrapper.getEntityName(), wrapper.getEntityId(),
                wrapper.getSubentityId(), wrapper.getGroupId(), wrapper.getClientId(),
        wrapper.getLoanId(), wrapper.getSavingsId(),
                wrapper.getTransactionId(), wrapper.getHref(), wrapper.getProductId(),
        wrapper.getCreditBureauId(),
                wrapper.getOrganisationCreditBureauId(), wrapper.getJobName(), wrapper
        .getLoanExternalId());

        return this.processAndLogCommandService.executeCommand(wrapper, command,
isApprovedByChecker);
    }

    @Override
    public CommandProcessingResult approveEntry(final Long makerCheckerId) {
        final CommandSource commandSourceInput = validateMakerCheckerTransaction
(makerCheckerId);
        validateIsUpdateAllowed();

        final CommandWrapper wrapper = CommandWrapper.fromExistingCommand(makerCheckerId,
commandSourceInput.getActionName(),
                commandSourceInput.getEntityName(), commandSourceInput.getResourceId(),
commandSourceInput.getSubResourceId(),
                commandSourceInput.getResourceGetUrl(), commandSourceInput.getProductId(
), commandSourceInput.getOfficeId(),
                commandSourceInput.getGroupId(), commandSourceInput.getClientId(),
commandSourceInput.getLoanId(),
                commandSourceInput.getSavingsId(), commandSourceInput.getTransactionId(),
commandSourceInput.getCreditBureauId(),
                commandSourceInput.getOrganisationCreditBureauId(), commandSourceInput
.getIdempotencyKey(),
                commandSourceInput.getLoanExternalId());
        final JsonElement parsedCommand = this.fromApiJsonHelper.parse(
commandSourceInput.getCommandAsJson());
        final JsonCommand command = JsonCommand.fromExistingCommand(makerCheckerId,
commandSourceInput.getCommandAsJson(), parsedCommand,
                this.fromApiJsonHelper, commandSourceInput.getEntityName(),
commandSourceInput.getResourceId(),
                commandSourceInput.getSubResourceId(), commandSourceInput.getGroupId(),
commandSourceInput.getClientId(),
                commandSourceInput.getLoanId(), commandSourceInput.getSavingsId(),
commandSourceInput.getTransactionId(),
                commandSourceInput.getResourceGetUrl(), commandSourceInput.getProductId(
), commandSourceInput.getCreditBureauId(),
                commandSourceInput.getOrganisationCreditBureauId(), commandSourceInput
.getJobName(),

```

```

        commandSourceInput.getLoanExternalId()));

    return this.processAndLogCommandService.executeCommand(wrapper, command, true);
}

@Transactional
@Override
public Long deleteEntry(final Long makerCheckerId) {

    validateMakerCheckerTransaction(makerCheckerId);
    validateIsUpdateAllowed();

    this.commandSourceRepository.deleteById(makerCheckerId);

    return makerCheckerId;
}

private CommandSource validateMakerCheckerTransaction(final Long makerCheckerId) {
    final CommandSource commandSource = this.commandSourceRepository.findById
(makerCheckerId)
        .orElseThrow(() -> new CommandNotFoundException(makerCheckerId));
    if (!commandSource.isMarkedAsAwaitingApproval()) {
        throw new CommandNotAwaitingApprovalException(makerCheckerId);
    }
    AppUser appUser = this.context.authenticatedUser();
    String permissionCode = commandSource.getPermissionCode();
    appUser.validateHasCheckerPermissionTo(permissionCode);
    if (!configurationService.isSameMakerCheckerEnabled() && !appUser
.isCheckerSuperUser()) {
        AppUser maker = commandSource.getMaker();
        if (maker == null) {
            throw new UnsupportedCommandException(permissionCode, "Maker user is
missing.");
    }
}

```

Check user has permission for this action. if ok, a) parse the json request body, b) create a JsonCommand object to wrap the command details, c) use CommandProcessingService to handle command

```

@Retry(name = "executeCommand", fallbackMethod = "fallbackExecuteCommand")
public CommandProcessingResult executeCommand(final CommandWrapper wrapper, final
JsonCommand command,
        final boolean isApprovedByChecker) {
    // Do not store the idempotency key because of the exception handling
    setIdempotencyKeyStoreFlag(false);

    Long commandId = (Long) fineractRequestContextHolder.getAttribute
(COMMAND_SOURCE_ID, null);
    boolean isRetry = commandId != null;
}

```

```

        boolean isEnclosingTransaction = BatchRequestContextHolder.
isEnclosingTransaction();

        CommandSource commandSource = null;
        String idempotencyKey;
        if (isRetry) {
            commandSource = commandSourceService.getCommandSource(commandId);
            idempotencyKey = commandSource.getIdempotencyKey();
        } else if ((commandId = command.commandId()) != null) { // action on the command
itself
            commandSource = commandSourceService.getCommandSource(commandId);
            idempotencyKey = commandSource.getIdempotencyKey();
        } else {
            idempotencyKey = idempotencyKeyResolver.resolve(wrapper);
        }
        exceptionWhenTheRequestAlreadyProcessed(wrapper, idempotencyKey, isRetry);

        AppUser user = context.authenticatedUser(wrapper);
        if (commandSource == null) {
            if (isEnclosingTransaction) {
                commandSource = commandSourceService.getInitialCommandSource(wrapper,
command, user, idempotencyKey);
            } else {
                commandSource = commandSourceService.saveInitialNewTransaction(wrapper,
command, user, idempotencyKey);
                commandId = commandSource.getId();
            }
        }
        if (commandId != null) {
            storeCommandIdInContext(commandSource); // Store command id as a request
attribute
        }

        boolean isMakerChecker = configurationDomainService.isMakerCheckerEnabledForTask
(wrapper.taskPermissionName());
        if (isApprovedByChecker || (isMakerChecker && user.isCheckerSuperUser())) {
            commandSource.markAsChecked(user);
        }
        setIdempotencyKeyStoreFlag(true);

        final CommandProcessingResult result;
        try {
            result = commandSourceService.processCommand(findCommandHandler(wrapper),
command, commandSource, user, isApprovedByChecker,
                isMakerChecker);
        } catch (Throwable t) { // NOSONAR
            RuntimeException mappable = ErrorHandler.getMappable(t);
            ErrorInfo errorInfo = commandSourceService.generateErrorInfo(mappable);
        }
    }
}

```

```

    Integer statusCode = errorInfo.getStatusCode();
    commandSource.setResultStatusCode(statusCode);
    commandSource.setResult(errorInfo.getMessage());
    if (statusCode != SC_OK) {
        commandSource.setStatus(ERROR.getValue());
    }
    if (!isEnclosingTransaction) { // TODO: temporary solution
        commandSource = commandSourceService.saveResultNewTransaction
(commandSource);
    }
    // must not throw any exception; must persist in new transaction as the
    current transaction was already
    // marked as rollback
    publishHookErrorEvent(wrapper, command, errorInfo);
    throw mappable;
}

commandSource.setResultStatusCode(SC_OK);
commandSource.updateForAudit(result);
commandSource.setResult(toApiResultJsonSerializer.serializeResult(result));

```



if a `RollbackTransactionAsCommandIsNotApprovedByCheckerException` occurs at this point. The original transaction will of been aborted and we only log an entry for the command in the audit table setting its status as 'Pending'.

- Check that if maker-checker configuration enabled for this action. If yes and this is not a 'checker' approving the command - rollback at the end. We rollback at the end in order to test if the command will pass 'domain validation' which requires commit to database for full check.
- `findCommandHandler` - Find the correct Handler to process this command.
- Process command using handler (In transactional scope).
- CommandSource object created/updated with all details for logging to 'm_portfolio_command_source' table.
- In update scenario, we check to see if there where really any changes/updates. If so only JSON for changes is stored in audit log.

Persistence

TBD

Database support

Fineract supports multiple databases:

- MySQL compatible databases (e.g. MariaDB)

- PostgreSQL

The platform differentiates between these database types in certain cases when there's a need to use some database specific tooling. To do so, the platform examines the JDBC driver used for running the platform and tries to determine which database is being used.

The currently supported JDBC driver and corresponding mappings can be found below.

JDBC driver class name	Resolved database type
<code>org.mariadb.jdbc.Driver</code>	MySQL
<code>com.mysql.jdbc.Driver</code>	MySQL
<code>org.postgresql.Driver</code>	PostgreSQL

The actual code can be found in the `DatabaseTypeResolver` class.

Tenant database security

The tenant database schema password is stored in the `tenant_server_connections` table in the tenant database.

The password and the read only schema password are encrypted using the `fineract.tenant.master-password` property.

By default, the database property will be encrypted in the first start from a plain text.

When you want to generate a new encrypted password, you can use the `org.apache.fineract.infrastructure.core.service.database.DatabasePasswordEncryptor` class.

Database password encryption usage

```
java -cp fineract-provider.jar -Dloader.main=org.apache.fineract.infrastructure.core
.service.database.DatabasePasswordEncryptor org.springframework.boot.loader
.PropertiesLauncher <masterPassword> <plainPassword>
```

For example:

```
java -cp fineract-provider-0.0.0-48f7e315.jar -Dloader.main=org.apache.fineract
.infrastructure.core.service.database.DatabasePasswordEncryptor org.springframework.boot
.loader.PropertiesLauncher fineract-master-password fineract-tenant-password
The encrypted password: VLwG17v0P/q275ZTku+PNGWnGwW4mzzNHSNa09Pr67WT5
/NZMpBr9tGYYiYsqwL1eRew2j1703/N1EFbL1XhSA==
```

Data-access layer

The data-access layer of Fineract is implemented by using JPA (Java Persistence API) with the

EclipseLink provider.

Despite the fact that JPA is used quite extensively in the system, there are cases where the performance is a key element for an operation therefore you can easily find native SQLs as well.

The data-access layer of Fineract is compatible with different databases. Since a lot of the native queries are using specific database functions, a wrapper class - `DatabaseSpecificSQLGenerator` - has been introduced to handle these database specifics. Whenever there's a need to rely on new database level functions, make sure to extend this class and implement the specific functions provided by the database.

Fineract has been developed for 10+ years by the community and unfortunately there are places where entity relationships are configured with `EAGER` fetching strategy. This must not confuse anybody. The long-term goal is to use the `LAZY` fetching strategy for every single relationship. If you're about to introduce a new one, make sure to use `LAZY` as a fetching strategy, otherwise your PR will be rejected.

Database schema migration

As for every system, the database structure will and need to evolve over time. Fineract is no different. Originally for Fineract, Flyway was used until Fineract 1.6.x.

After 1.6.x, PostgreSQL support was added to the platform hence there was a need to make the data-access layer and the schema migration as database independent as possible. Because of that, from Fineract 1.7.0, Flyway is not used anymore but Liquibase is.

Some of the changesets in the Liquibase changelogs have database specifics into it but they only run for the relevant databases. This is controlled by Liquibase contexts.

The currently available Liquibase contexts are:

- `mysql` - only set when the database is a MySQL compatible database (e.g. MariaDB)
- `postgresql` - only set when the database is a PostgreSQL database
- configured Spring active profiles
- `tenant_store_db` - only set when the database migration runs the Tenant Store upgrade
- `tenant_db` - only set when the database migration runs the Tenant upgrade
- `initial_switch` - this is a technical context and should **NOT** be used

The switch from Flyway (1.6.x) to Liquibase (1.7.x) was planned to be as smooth as possible so there's no need for manual work hence the behavior is described as following:

- If the database is empty, Liquibase will create the database schema from scratch
- If the database contains the latest Fineract 1.6.x database structure which was previously migrated with Flyway. Liquibase will seamlessly upgrade it to the latest version. Note: the Flyway related 2 database tables are left as they are and are not deleted.
- If the database contains an earlier version of the database structure than Fineract 1.6.x. Liquibase

will NOT do anything and **will fail the application during startup**. The proper approach in this case is to first upgrade your application version to the latest Fineract 1.6.x so that the latest Flyway changes are executed and then upgrade to the newer Fineract version where Liquibase will seamlessly take over the database upgrades.

Troubleshooting

1. During upgrade from Fineract 1.5.0 to 1.6.0, Liquibase fails

After dropping the flyway migrations table (schema_version), Liquibase runs its own migrations which fails (in recreating tables which already exist) because we are aiming to re-use DB with existing data from Fineract 1.5.0.

Solution: The latest release version (1.6.0) doesn't have Liquibase at all, it still runs Flyway migrations. Only the develop branch (later to be 1.7.0) got switched to Liquibase. Do not pull the develop before upgrading your instance.

Make sure first you upgrade your instance (aka database schema with Fineract 1.6.0). Then upgrade with the current develop branch. Check if some migration scripts did not run which led to some operations failing due to slight differences in schema. Try with running the missing migrations manually.

Note: develop is considered unstable until released.

1. Upgrading database from MySQL 5.7 as advised to Maria DB 10.6, fails. If we use data from version 18.03.01 it fails to migrate the data. If we use databases running on 1.5.0 release it completes the startup but the system login fails.

Solution: A database upgrade is separate thing to take care of.

1. We are getting **ScehmaUpgradeNeededException: Make sure to upgrade to Fineract 1.6 first and then to a newer version** error while upgrading to tag 1.6.

1.6 version shouldn't include Liquibase. It will only be released after 1.6.

Make sure Liquibase is dropping **schema_version** table, as there is no Flyway it is not required. Drop Flyway and use Liquibase for both migrations and database independence. In case, if you still get errors, you can use git SHA **746c589a6e809b33d68c0596930fcaa7338d5270** and Flyway migration will be done to the latest.

```
TENANT_LATEST_FLYWAY_VERSION = 392;  
TENANT_LATEST_FLYWAY_SCRIPT_NAME =  
"V392__interest_recovery_conf_for_reschedule.sql";  
TENANT_LATEST_FLYWAY_SCRIPT_CHECKSUM = 1102395052;
```

Idempotency

Idempotency is the way to make sure your specific action is only executed once.

For example, if you have a button that is supposed to send a repayment, you don't want to repayment twice if the user clicks the button twice. Idempotency is a way to make sure that the action is only executed once.

There are two ways to use idempotency:

- HTTP Request with idempotency key header
- Batch request with batch item header

How it works

The `idempotency key` with `action name` and `entity name` is unique, and identify a specific command in the system.

If no idempotency key is assigned to the request, the system will generate one for you.

1. User send a request
2. The system checks there are already executed commands with the same `idempotency key` and `action name` and `entity name`
3. The action based on the result of the check
 - If the request is completed the system return with the already generated result
 - If not completed, return HTTP 409 response
 - If the request is not completed, we process the requests and store the results in the database

Idempotency in HTTP requests

To achieve idempotency in HTTP requests, you can use the HTTP header from `fineract.idempotency-key-header-name` configuration variables (default `Idempotency-Key`). This header is a unique identifier for the request. If you send the same request twice, the second request will be ignored and the response from the first request will be returned.

Idempotency in Batch requests

In batch requests, you can set the idempotency key for every batch item, in the batch item `header` fields. The header key is from `fineract.idempotency-key-header-name` configuration variables (default `Idempotency-Key`).

Result of the request

- When the request is already executed and completed, the system will return a `x-served-from-cache` header with the value `true` in the response and return the original request body.

- When the request is already executed but still not completed, the system will return to HTTP 409 error code
- When the request is not executed, the system runs it normally and stores the result in the date

Validation

Programmatic

Use the [DataValidatorBuilder](#), e.g. like so:

```
new DataValidatorBuilder().resource("fileUpload")
    .reset().parameter("Content-Length").value(contentLength).notBlank()
    .integerGreaterThanOrEqualTo(0)
    .reset().parameter("FormDataContentDisposition").value(fileDetails).notNull()
    .throwValidationErrors();
```

Such code is often encapsulated in [*Validator](#) classes (if more than a few lines, and/or reused from several places; avoid copy/paste), like so:

```
public class YourThingValidator {

    public void validate(YourThing thing) {
        new DataValidatorBuilder().resource("yourThing")
        ...
        .throwValidationErrors();
    }
}
```

Declarative

[FINERACT-1229](issues.apache.org/jira/browse/FINERACT-1229) is an open issue about adopting Bean Validation for *declarative* instead of *programmatic* (as above) validation. Contributions welcome!

Batch execution and jobs

Just like any financial system, Fineract also has batch jobs to achieve some processing on the data that's stored in the system.

The batch jobs in Fineract are implemented using [Spring Batch](#). In addition to the Spring Batch ecosystem, the automatic scheduling is done by the [Quartz Scheduler](#) but it's also possible to trigger batch jobs via regular APIs.

Glossary

Job	A Job is an object that encapsulates an entire batch process.
Step	A Step is an object that encapsulates an independent phase of a Job.
Chunk oriented processing	Chunk oriented processing refers to reading the data one at a time and creating 'chunks' that are written out within a transaction boundary.
Partitioning	Partitioning refers to the high-level idea of dividing your data into so called partitions and distributing the individual partitions among Workers. The splitting of data and pushing work to Workers is done by a Manager.
Remote partitioning	Remote partitioning is a specialized partitioning concept. It refers to the idea of distributing the partitions among multiple JVMs mainly by using a messaging middleware.
Manager node	The Manager node is one of the objects taking a huge part when using partitioning. The Manager node is responsible for dividing the dataset into partitions and keeping track of all the divided partitions' Worker execution. When all Workers nodes are done with their partitions, the Manager will mark the corresponding Job as completed.
Worker node	A Worker node is the other important party in the context of partitioning. The Worker node is the one executing the work needed for a single partition.

Batch jobs in Fineract

Types of jobs

The jobs in Fineract can be divided into 2 categories:

- Normal batch jobs
- Partitionable batch jobs

Most of the jobs are normal batch jobs with limited scalability because Fineract is still passing through the evolution on making most of them capable to process a high-volume of data.

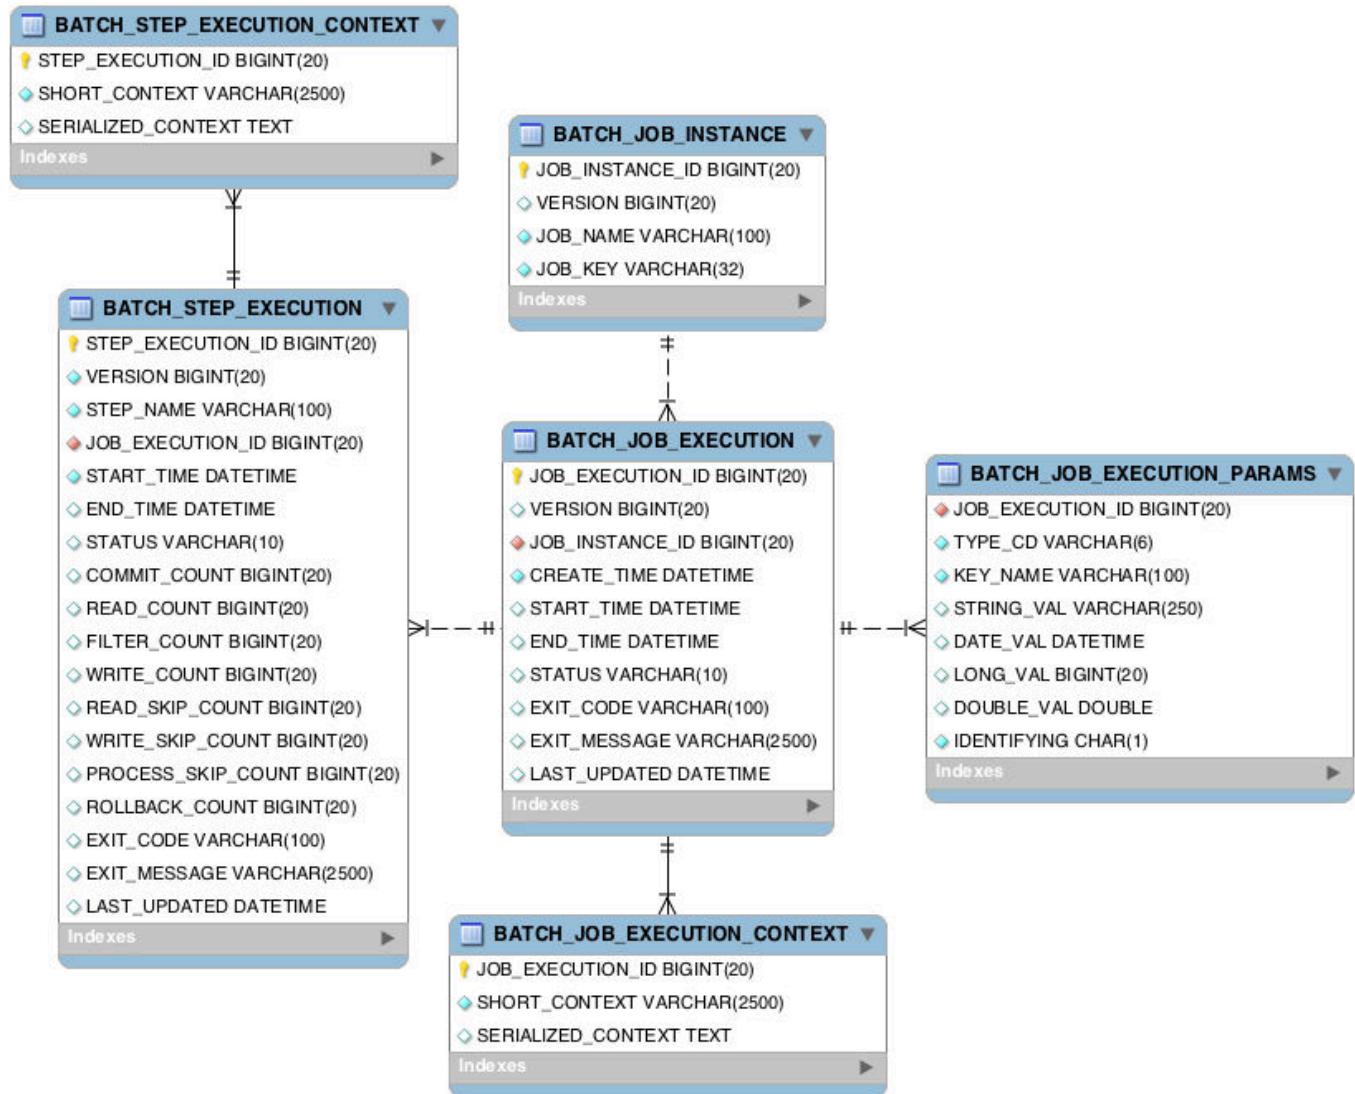
List of jobs

Job name	Active by default	Partitionable	Description
LOAN_CLOSE_OF_BUSINESS	No	Yes	TBD

Batch job execution

State management

State management for the batch jobs is done by the Spring Batch provided state management. The data model consists of the following database structure:



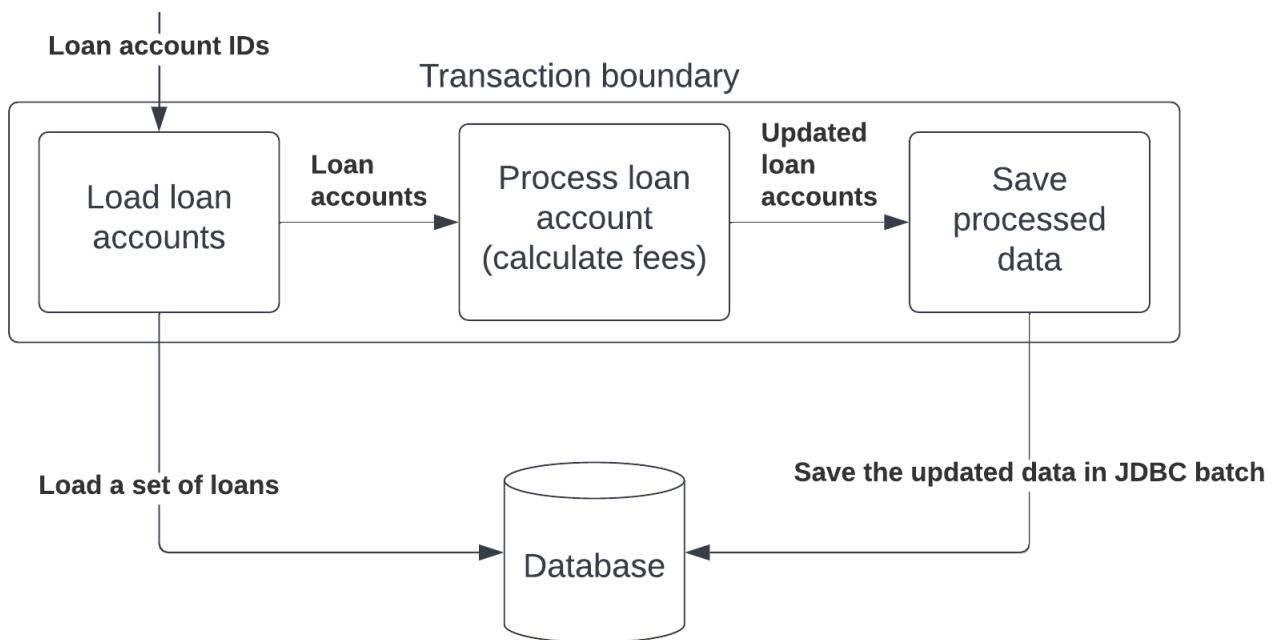
The corresponding database migration scripts are shipped with the Spring Batch core module under the `org.springframework.batch.core` package. They are only available as native scripts and are named as `schema-.sql` where `-` is the short name of the database platform. For MySQL it's called `schema-mysql.sql` and for PostgreSQL it's called `schema-postgresql.sql`.

When Fineract is started, the database dependent schema SQL script will be picked up according to the datasource configurations.

Chunk oriented processing

Chunking data has not been easier. Spring Batch does a really good job at providing this capability.

In order to save resources when starting/committing/rollbacking transactions for every single processed item, chunking shall be used. That way, it's possible to mark the transaction boundaries for a single processed chunk instead of a single item processing. The image below describes the flow with a very simplistic example.



In addition to not opening a lot of transactions, the processing could also benefit from JDBC batching. The last step - writing the result into the database - collects all the processed items and then writes it to the database; both for MySQL and PostgreSQL (the databases supported by Fineract) are capable of grouping multiple DML (INSERT/UPDATE/DELETE) statements and sending them in one round-trip, optimizing the data being sent over the network and granting the possibility to the underlying database engine to enhance the processing.

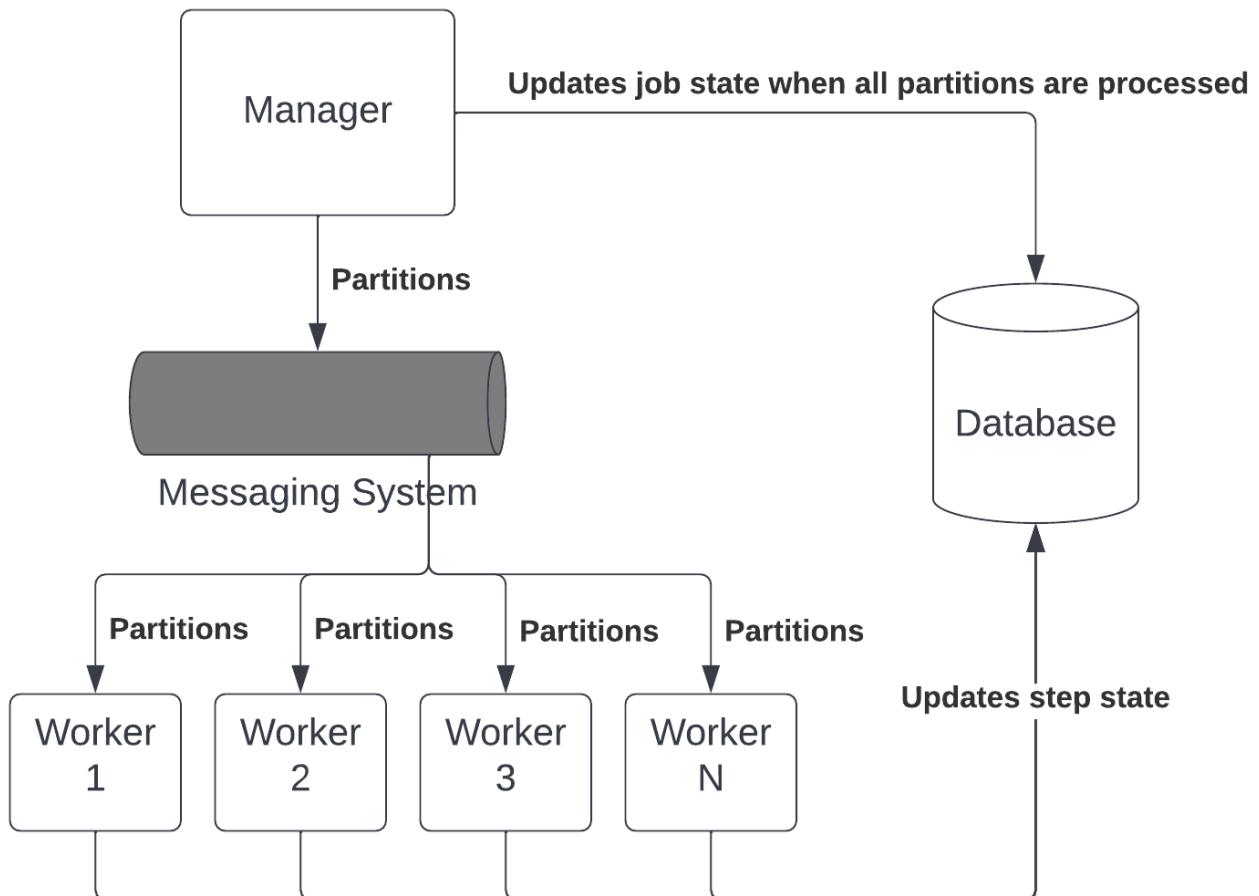
Remote partitioning

Spring Batch provides a really nice way to do remote partitioning. The 2 type of objects in this setup is a manager node - who splits and distributes the work - and a number of worker nodes - who picks up the work.

In remote partitioning, the worker instances are receiving the work via a messaging system as soon as the manager splits up the work into smaller pieces.

Remote partitioning could be done 2 ways in terms of keeping the job state up-to-date. The main difference between the two is how the manager is notified about partition completions.

One way is that they share the same database. When the worker does something to a partition - for example picks it up for processing - it updates the state of that partition in the database. In the meantime, the manager regularly polls the database until all partitions are processed. This is visualized in the below diagram.



An alternative approach to this - when the database is not intended to be shared between manager and workers - is to use a messaging system (could be the same as for distributing the work) and the workers could send back a message to the manager instance, therefore notifying it about failure/completion. Then the manager can simply keep the database state up-to-date.

Even though the alternative solution decouples the workers even better, we thought it's not necessary to add the complexity of handling reply message channel to the manager.

Also, please note that the partitioned job execution is multitenant meaning that the workers will receive which tenant it should do the processing for.

Supported message channels

For remote partitioning, the following message channels are supported by Fineract:

- Any JMS compatible message channels (ActiveMQ, Amazon MQ, etc)
- Apache Kafka

Fault-tolerance scenarios

There are multiple fault tolerance use-cases that this solution must and will support:

1. If the manager fails during partitioning
2. If the manager completes the partitioning and the partition messages are sent to the broker but while the manager is waiting for the workers to finish, the manager fails
3. If the manager runs properly and during a partition processing a worker instance fails

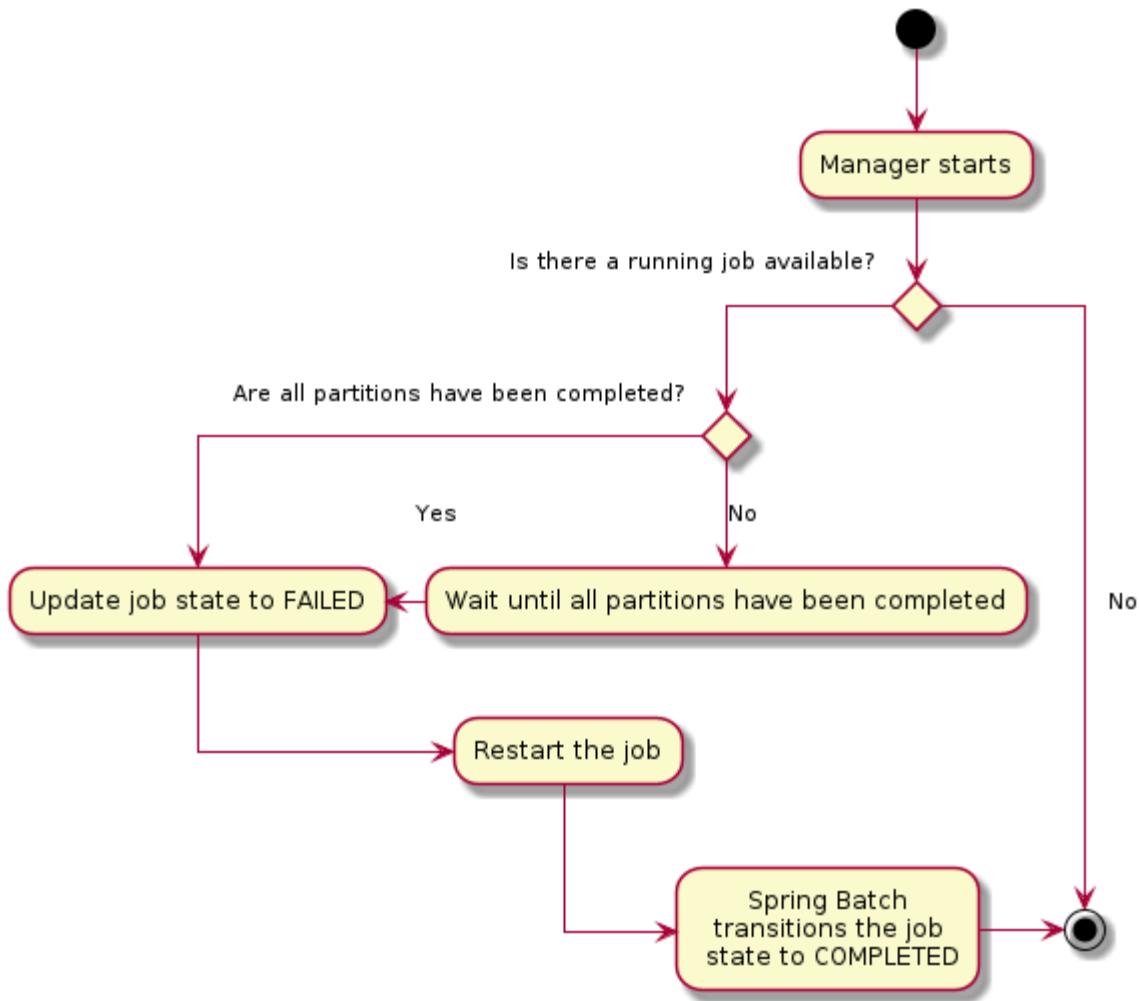
In case of scenario 1), the simple solution is to re-trigger the job via API or via the Quartz scheduler.

In case of scenario 2), there's no out-of-the-box solution by Spring Batch. Although there's a custom mechanism in place that'll resume the job upon restarting the manager. There are 2 cases in the context of this scenario:

- If all the partitions have been successfully processed by workers
- If not all the partitions have been processed by the workers

In the first case, we'll simply mark the stuck job as **FAILED** along with its partitioning step and instruct Spring Batch to restart the job. The behavior in this case will be that Spring Batch will spawn a new job execution but will notice that the partitions have all been completed so it's not going to execute them once more.

In the latter case, the same will happen as for the first one but before marking the job execution as **FAILED**, we'll wait until all partitions have been completed.



In case of scenario 3), another worker instance will take over the partition since it hasn't been finished.

Configurable batch jobs

There's another type of distinction on the batch jobs. Some of them are configurable in terms of their behavior.

The currently supported configurable batch jobs are the following:

- LOAN_CLOSE_OF_BUSINESS

The behavior of these batch jobs are configurable. There's a new terminology we're introducing called **business steps**.

Business steps

Business steps are a smaller unit of work than regular Spring Batch Steps and the two are not meant to be mixed up because there's a large difference between them.

A Spring Batch Step's main purpose is to decompose a bigger work into smaller ones and making sure that these smaller Steps are properly handled within a single database transaction.

In case of a business step, it's a smaller unit of work. Business steps live **within** a Spring Batch Step. Fundamentally, they are simple classes that are implementing an interface with a single method that contains the business logic.

Here's a very simple example:

```
public class MyCustomBusinessStep implements BusinessStep<Loan> {  
    @Override  
    public Loan process(Loan loan) {  
        // do something  
    }  
}
```

```
public class LoanCOBItemProcessor implements ItemProcessor<Loan, Loan> {  
    @Override  
    public Loan process(Loan loan) {  
        List<BusinessStep<Loan>> bSteps = getBusinessSteps();  
        Loan result = loan;  
        for (BusinessStep<Loan> bStep : bSteps) {  
            result = bStep.process(result);  
        }  
        return result;  
    }  
}
```

Business step configuration

The business steps are configurable for certain jobs. The reason for that is because we want to allow the possibility for Fineract users to configure their very own business logic for generic jobs, like the Loan Close Of Business job where we want to do a formal "closing" of the loans at the end of the day.

All countries are different with a different set of regulations. However in terms of behavior, there's no all size fits all for loan closing.

For example in the United States of America, you might need the following logic for a day closing:

1. Close fully repaid loan accounts
2. Apply penalties
3. Invoke IRS API for regulatory purposes

While in Germany it should be:

1. Close fully repaid loan accounts
2. Apply penalties

3. Do some fraud detection on the account using an external service
4. Invoke local tax authority API for regulatory purposes

These are just examples, but you get the idea.

The business steps are configurable through APIs:

Retrieving the configuration for a job:

```
GET /fineract-provider/api/v1/jobs/{jobName}/steps?tenantIdentifier={tenantId}
HTTP 200

{
  "jobName": "LOAN_CLOSE_OF_BUSINESS",
  "businessSteps": [
    {
      "stepName": "APPLY_PENALTY_FOR_OVERDUE_LOANS",
      "order": 1
    },
    {
      "stepName": "LOAN_TAGGING",
      "order": 2
    }
  ]
}
```

Updating the business step configuration for a job:

```
PUT /fineract-provider/api/v1/jobs/{jobName}/steps?tenantIdentifier={tenantId}

{
  "businessSteps": [
    {
      "stepName": "LOAN_TAGGING",
      "order": 1
    },
    {
      "stepName": "APPLY_PENALTY_FOR_OVERDUE_LOANS",
      "order": 2
    }
  ]
}
```

The business step configuration for jobs are tracked within the database in the `m_batch_business_steps` table.

Inline Jobs

Some jobs that work with business entities have a corresponding job that can trigger the job with a list of specified entities.

When the Inline job gets triggered then the corresponding existing job will run in real time with the given entities as a dataset.

List of Inline jobs

Inline Job name	Corresponding Job
LOAN_COB	LOAN_CLOSE_OF_BUSINESS

Triggering the Inline Loan COB Job:

```
POST /fineract-provider/api/v1/jobs/LOAN_COB/inline?tenantIdentifier={tenantId}

{
  "loanIds": [
    1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14
  ]
}
```

In this case the Loan COB job will work only with the given loan IDs.

Global Configuration for enabling/disabling jobs

Some jobs can be enabled/disabled with global configuration.

If the job is disabled with the global configuration then it cannot be scheduled and cannot be triggered via API.

List of jobs with global configuration

Job name	Application property	Environment variable	Default value
LOAN_CLOSE_OF_BUSIN ESS	fineract.job.loan-cob- enabled	FINERACT_JOB_LOAN_C OB_ENABLED	true

Loan account locking

Keeping a consistent state of loan accounts become quite important when we start talking about doing a business day closing each day for loans.

There are 2 concepts for loan account locking:

1. Soft-locking loan accounts

2. Hard-locking loan accounts

Soft-locking simply means that when the Loan COB has been kicked off but workers not yet processing the chunk of loan accounts (i.e. the partition is waiting in the queue to be picked up) and during this time a real-time write request (e.g. a repayment/disbursement) comes in through the API, we simply do an "inlined" version of the Loan COB for that loan account. From a practical standpoint this will mean that before doing the actual repayment/disbursement on the loan account on the API, we execute the Loan COB for that loan account, kind of like prioritizing it.

Hard-locking means that when a worker picks up the loan account in the chunk, real-time write requests on those loan accounts will be simply rejected with an **HTTP 409**.

The locking is strictly tied to the Loan COB job's execution but there could be other processes in the future which might want to introduce new type of locks for loans.

The loan account locking is solved by maintaining a database table which stores the locked accounts, it's called **m_loan_account_locks**.

When a loan account is present in the table above, it simply means there's a lock applied to it and whether it's a soft or hard lock can be determined by the **lock_owner** column.

And when a loan account is locked, loan related write API calls will be either rejected or will trigger an inline Loan COB execution. There could be a corner case here when the Loan COB fails to process some loan accounts (due to a bug, inconsistency, etc) and the loan accounts stay locked. This is an intended behavior to mark loans which are not supposed to be used until they are "fixed".

Since the fixing might involve making changes to the loan account via API (for example doing a repayment to fix the loan account's inconsistent state), we need to allow those API calls. Hence, the lock table includes a **bypass_enabled** column which disables the lock checks on the loan write APIs.

Technology

- Java: www.oracle.com/technetwork/java/javase/downloads/index.html
- JAX-RS using Jersey
- JSON using Google GSON
- Spring I/O Platform: spring.io/platform
 - Spring Framework
 - Spring Boot
 - Spring Security
 - Spring Data (JPA) backed by EclipseLink
- MySQL: www.oracle.com/us/products/mysql/overview/index.html
- PostgreSQL

Modules

We are currently working towards a fully modular codebase and will publish more here when we are ready.



Even if we are not quite there yet with full modularity you can already create your own custom modules to extend Fineract. Please see chapter [Custom Modules](#).

Introducing Business Date into Fineract - Community version

Business date as a concept does not exist as of now in Fineract. It would be business critical to add such a functionality to support various banking functionalities like “Closing of Business day”, “Having Closing of Business day relevant jobs”, “Supporting logical date management”.

Glossary

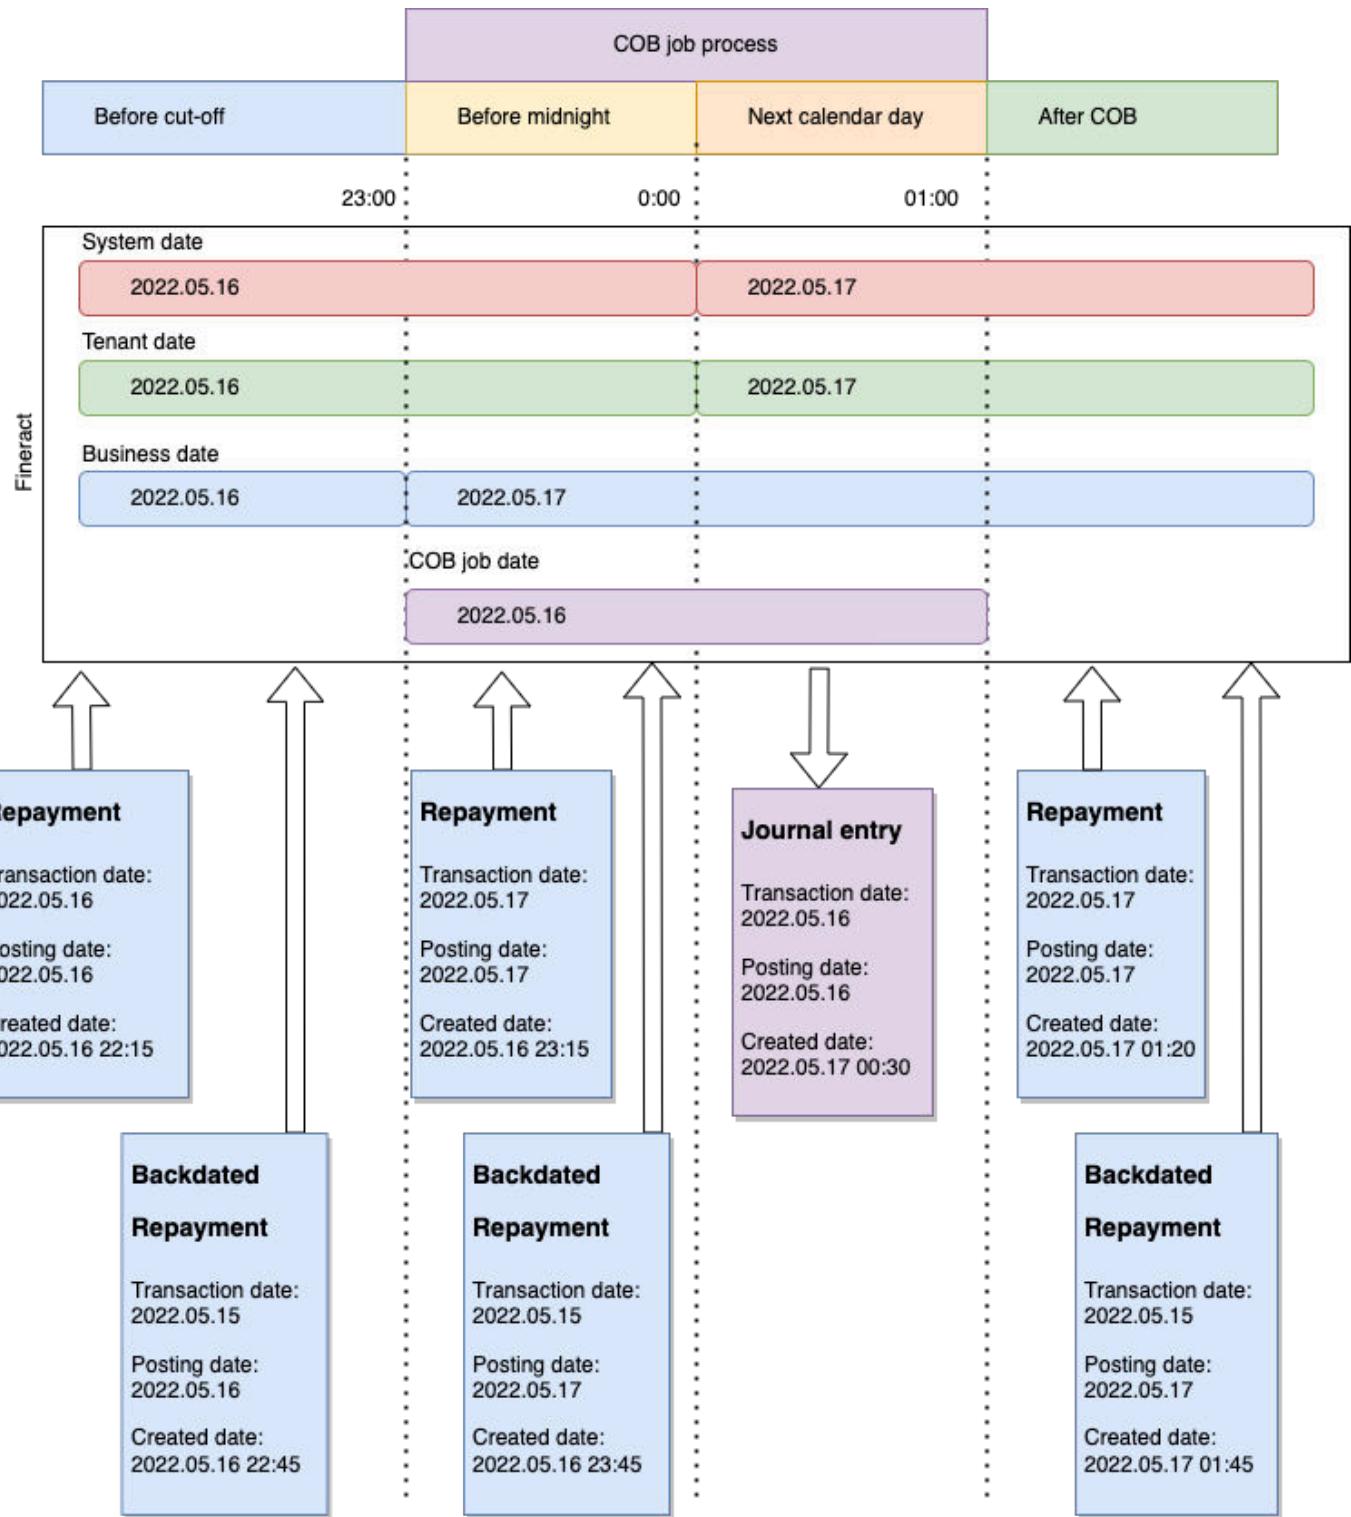
*COB	Close of Business; concept of closing a business day
*Business day	Timeframe that logically group together actions on a particular business date
*Business date	Logical date; its value is not tied to the physical calendar. Represents a business day
*Cob date	Logical date; Represents the business date for actions during COB job execution
*Created date	When the transaction was created (audit purposes). Date + time
*Last modified date	When the transaction was last modified (audit purposes). Date + time
*Submitted on date / Posting date	When the transaction was posted. Tenant date or business date (depends on whether the logical date concept was introduced or not)
*Transaction date / Value date	The date on which the transaction occurred or to be accounted for

Current behaviour

- Fineract support 3 types of dates:

- System date
 - Physical/System date of the running environment
- Tenant date
 - Timezoned version of the above system date
- User-provided date
 - Based on the provided date (as string) and the provided date format
- There is no support of logical date concept
 - Independent from the system / tenant date
- Jobs are scheduled against system date (CRON), but aligned with the tenant timezone.
- During the job execution all the data and transactions are using the actual tenant date
 - It could happen some transactions are written for 17th of May and other for 18th of May, if the job was executed around midnight
- There is no support of COB
 - No backdated transactions by jobs
 - There is no support to logically group together transactions and store them with the same transaction date which is independent of the physical calendar of the tenant
- All the transactions and business logic are tied to a physical calendar

Business date



Design

By introducing the business day concept we are not tied anymore to the physical calendar of the system or the tenant. We got the ability to define our own business day boundaries which might end 15 minutes before midnight and any incoming transactions after the cutoff will be accounted for the following business day.

It is a logical date which makes it possible to separate the business day from the physical calendar:

- Close a business day before midnight
- Close a business day at midnight
- Close a business day after midnight

Closing a Business Day could be a longer process (see COB jobs) meanwhile some processes shall still be able to create transactions for that business day (COB jobs), but others are meant to create the transactions for the next (incoming transactions): Business date concept is there to sort that out.

Business date concept is essential when:

- Having COB jobs:
 - When the COB was triggered:
 - All the jobs which processing the data must still accounted for actual business day
 - All the incoming transactions must be accounted to the next business day
- Business day is ending before / after midnight (tenant date / system date)
- Testing purposes:
 - Since the transactions and job execution is not tied anymore to a physical calendar, we can easily test a whole loan lifecycle by altering the business date
- Handling disruption of service: For any unseen reason the system goes down or there are any disruption in the workflow, the “missed days” can easily be processed one by one as nothing happened
 - There is a disruption at 2022-06-02
 - The issue is fixed by 2022-06-05
 - The COB flow can be executed for 2022-06-03 and when it is finished for 2022-06-04 and after when the time arrives for 2022-06-05

This logical date is manageable via:

- Job
- API

To maintain such separation from physical calendar we need to introduce the following new dates:

- Business date
- COB date
 - Can be calculated based on the actual business date
 - Depend on COB date strategy (see below)

Business date

The - logical - date of the actual business day, eg: 2022-05-06

- It does not support time parts
- It can be managed manually (via API call) or automatically (via scheduled job)
- All business actions during the business day shall use this date:
 - Posting / submitted on date of transactions
 - Submitted on date of actions
 - (Regular) jobs
- It will be used in every situation where the transaction date / value date is not provided by the user or the user provided date shall be validated.
 - Opening date
 - Closing date
 - Disbursal date
 - Transaction/Value date
 - Posting/Submitted date
 - Reversal date
- Will not be used for audit purposes:
 - Created on date
 - Updated on date

COB date

The - logical - date of the business day for job execution, eg: 2022-05-05

- It can be calculated based on the business date
 - COB date = business date - 1 day
 - Automatically modified alongside with the business date change
- It does not support time parts
- It is automatically managed by business date change
 - Configurable
- It is used only via COB job execution
 - When we create / modify any business data during the COB job execution, the COB date is to be used:
 - Posting date of transactions
 - Submitted on date of actions

- Transaction / value date of any actions

Some basic example

Apply for a loan

#1

Tenant date: 2022-05-23 14:22:12

Business date: 2022-05-22

Submitted on date: 2022-05-23

Outcome: **FAIL**

Message: **The date on which a loan is submitted cannot be in the future.**

Reason: Even the tenant date is 2022-05-23, but the business date was 2022-05-22 which means anything further than that date must be considered as a future date.

#2

Tenant date: 2022-05-23 14:22:12

Business date: 2022-05-22

Submitted on date: 2022-05-22

Outcome: **SUCCESS**

Loan application details:

- Submitted on date: 2022-05-22

Repayment for a loan

#1

Tenant date: 2022-05-25 11:22:12

Business date: 2022-05-24

Transaction date: 2022-05-25

Outcome: **FAIL**

Message: **The transaction date cannot be in the future.**

Reason: Even the physical date is 2022-05-25, but the business date was 2022-05-24 which means

anything further that date must be considered as a future date.

#2

Tenant date: 2022-05-25 11:22:12

Business date: 2022-05-24

Transaction date: 2022-05-23

Outcome: **SUCCESS**

Loan transaction details:

- Submitted on date: 2022-05-24
- Transaction date: 2022-05-23
- Created on date: 2022-05-25 11:22:12

Changes in Fineract

We shall modify at all the relevant places where the tenant date was used:

- With very limited exceptions all places where the tenant date is used we need to modify to use the business date.
- Replace system date with tenant date or business date (exceptions may apply)
- Add missing Value dates and Posting dates to entities
- Having a generic naming conventions for JPA fields and DB fields
- Renaming the fields accordingly
- Evaluate value date (transaction date) and posting date (submitted on date), created on date usages
- Jobs to be checked and modified accordingly
- Native queries to be checked and modified accordingly
- Reports to be checked and modified accordingly
- Every table where update is supported the AbstractAuditableCustom should be implemented
- Amend Transactions and Journal entries date handling to fit for business date concept
- For audit fields we shall introduce timezoned datetimes and store them in database accordingly
 - Storing DATETIME fields without Timezone is potential problem due to the daylight savings
 - Also, some external libs (like Quartz) are using system timezone and Fineract will use Tenant timezone for audit fields. To be able to distinct them in DB we shall use DATETIME with TIMESTAMP column types and use timezoned java time objects in the application

Reliable event framework

Fineract is capable of generating and raising events for external consumers in a reliable way. This section is going to describe all the details on that front with examples.

Framework capabilities

ACID (transactional) guarantee

The event framework must support ACID guarantees on the business operation level.

Let's see a simple use-case:

1. A client applies to a loan on the UI
2. The loan is created on the server
3. A loan creation event is raised

What happens if step 3 fails? Shall it fail the original loan creation process?

What happens if step 2 fails but step 3 still gets executed? We're raising an event for a loan that hasn't been created in reality.

Therefore, raising an event is tied to the original business transaction to ensure the data that's getting written into the database along with the respective events are saved in an all-or-nothing fashion.

Messaging integration

The system is able to send the raised events to downstream message channels. The current implementation supports the following message channels:

- ActiveMQ

Ordering guarantee

The events that are raised will be sent to the downstream message channels in the same order as they were raised.

Delivery guarantee

The framework supports the at-least-once delivery guarantee for the raised events.

Reliability and fault-tolerance

In terms of reliability and fault-tolerance, the event framework is able to handle the cases when the downstream message channel is not able to accept events. As soon as the message channel is back to operational, the events will be sent again.

Selective event producing

Whether or not an event must be sent to downstream message channels for a particular Fineract instance is configurable through the UI and API.

Standardized format

All the events sent to downstream message channels are conforming a standardized format using Avro schemas.

Extendability and customizations

The event framework is capable of being easily extended with new events for additional business operations or customizing existing events.

Ability to send events in bulk

The event framework makes it possible to sort of queue events until they are ready to be sent and send them as a single message instead of sending each event as a separate, individual one.

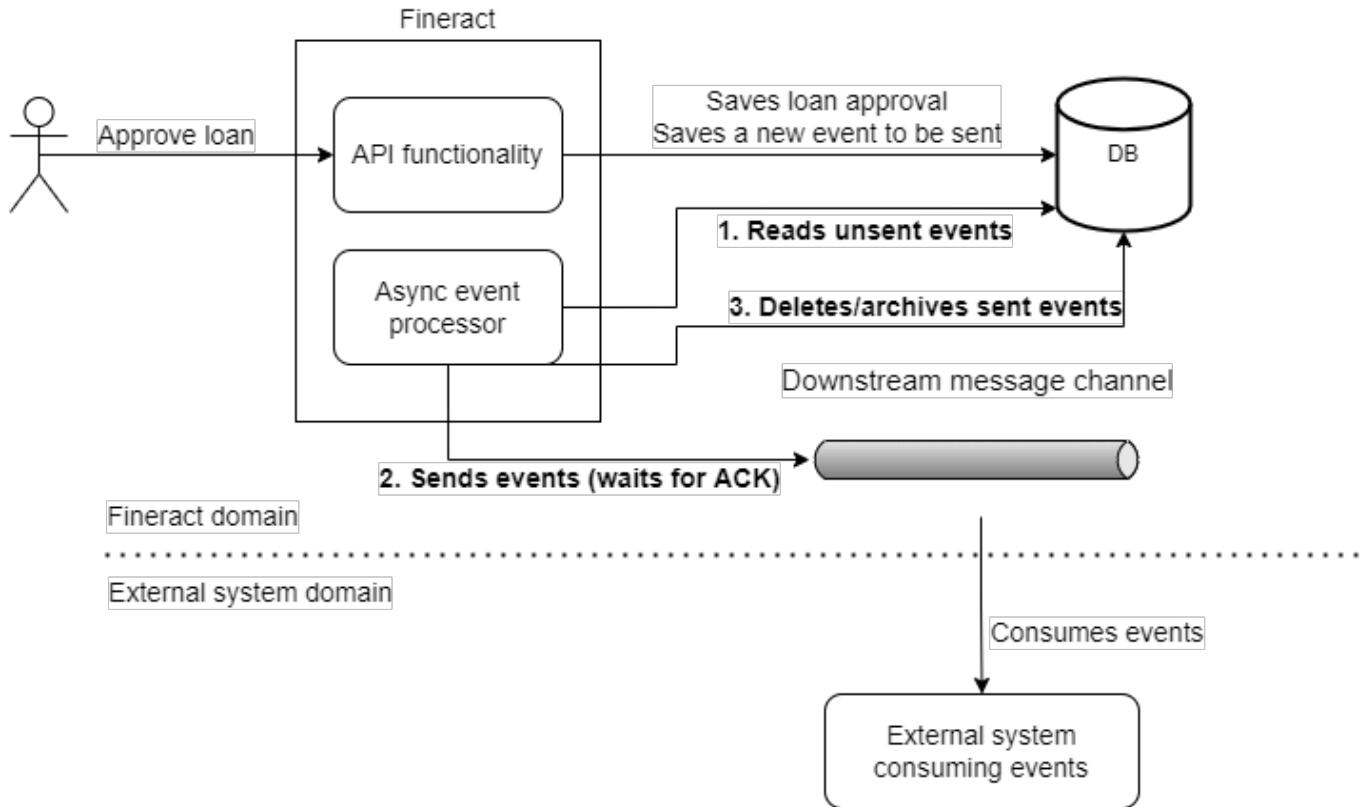
For example during the COB process, there might be events raised in separate business steps which needs to be sent out but they only need to be sent out at the end of the COB execution process instead of one-by-one.

Architecture

Intro

On a high-level, the concept looks the following. An event gets raised in a business operation. The event data gets saved to the database - to ensure ACID guarantees. An asynchronous process takes the saved events from the database and puts them onto a message channel.

The flow can be seen in the following diagram:



Foundational business events

The whole framework is built upon an existing infrastructure in Fineract; the Business Events.

As a quick recap, Business Events are Fineract events that can be raised at any place in a business operation using the [BusinessEventNotifierService](#). Callbacks can be registered when a certain type of Business Event is raised and other business operations can be done. For example when a Loan gets disbursed, there's an interested party doing the Loan Arrears Aging recalculation using the Business Event communication.

The nice thing about the Business Events is that they are tied to the original transaction which means if any of the processing on the subscriber's side fail, the entire original transaction will be rolled back. This was one of the requirements for the Reliable event framework.

Event database integration

The database plays a crucial part in the framework since to ensure transactionality, - without doing proper transaction synchronization between different message channels and the database - the framework is going to save all the raised events into the same relational database that Fineract is using.

Database structure

The database structure looks the following

Name	Type	Description	Example
<code>id</code>	number	Auto incremented ID. Not null.	1
<code>type</code>	text	The event type as a string. Not null.	<code>LoanApprovedBusinessEvent</code>
<code>schema</code>	text	The fully qualified name of the schema that was used for the data serialization, as a string. Not null.	<code>org.apache.fineract.avro.loan.v1.LoanAccountDataV1</code>
<code>data</code>	BLOB (MySQL/MariaDB), BYTEA (PostgreSQL)	The event payload as Avro binary. Not null.	
<code>created_at</code>	timestamp	UTC timestamp when the event was raised. Not null.	<code>2022-09-06 14:20:10.148627 +00:00</code>
<code>status</code>	text	Enum text representing the status of the external event. Not null, indexed.	<code>TO_BE_SENT, SENT</code>
<code>sent_at</code>	timestamp	UTC timestamp when the event was sent.	<code>2022-09-06 14:30:10.148627 +00:00</code>
<code>idempotency_key</code>	text	Randomly generated UUID upon inserting a row into the table for idempotency purposes. Not null.	<code>68aed085-8235-4722-b27d-b38674c19445</code>
<code>business_date</code>	date	The business date to when the event was generated. Not null, indexed.	<code>2022-09-05</code>

The above database table contains the unsent events which later on will be sent by an asynchronous event processor.

Upon successfully sending an event, the corresponding statuses will be updated.

Avro schemas

For serializing events, Fineract is using Apache Avro. There are 2 reasons for that:

- More compact storage since Avro is a binary format
- The Avro schemas are published with Fineract as a separate JAR so event consumers can directly map the events into POJOs

There are 3 different levels of Avro schemas used in Fineract for the Reliable event framework which are described below.

Standard event schema

The standard event schema is for the regular events. These schemas are used when saving a raised event into the database and using the Avro schema to serialize the event data into a binary format.

For example the OfficeDataV1 Avro schema looks the following:

```
{  
  "name": "OfficeDataV1",  
  "namespace": "org.apache.fineract.avro.office.v1",  
  "type": "record",  
  "fields": [  
    {  
      "default": null,  
      "name": "id",  
      "type": [  
        "null",  
        "long"  
      ]  
    },  
    {  
      "default": null,  
      "name": "name",  
      "type": [  
        "null",  
        "string"  
      ]  
    },  
    {  
      "default": null,  
      "name": "nameDecorated",  
      "type": [  
        "null",  
        "string"  
      ]  
    },  
    {  
      "default": null,  
      "name": "externalId",  
      "type": [  
        "null",  
        "string"  
      ]  
    },  
    {  
      "default": null,  
      "name": "openingDate",  
      "type": [  
        "null",  
        "string"  
      ]  
    },  
  ]  
},
```

```

{
  "default": null,
  "name": "hierarchy",
  "type": [
    "null",
    "string"
  ]
},
{
  "default": null,
  "name": "parentId",
  "type": [
    "null",
    "long"
  ]
},
{
  "default": null,
  "name": "parentName",
  "type": [
    "null",
    "string"
  ]
},
{
  "default": null,
  "name": "allowedParents",
  "type": [
    "null",
    {
      "type": "array",
      "items": "org.apache.fineract.avro.office.v1.OfficeDataV1"
    }
  ]
}
]
}

```

Event message schema

The event message schema is just a wrapper around the standard event schema with extra metadata for the event consumers.

Since Avro is strongly typed, the event content needs to be first serialized into a byte sequence and that needs to be wrapped around.

This implies that for putting a single event message onto a message queue for external consumption, data needs to be serialized 2 times; this is the 2-level serialization.

1. Serializing the event
2. Serializing the already serialized event into an event message using the message wrapper

The message schema looks the following:

MessageV1.avsc

```
{  
  "name": "MessageV1",  
  "namespace": "org.apache.fineract.avro",  
  "type": "record",  
  "fields": [  
    {  
      "name": "id",  
      "doc": "The ID of the message to be sent",  
      "type": "long"  
    },  
    {  
      "name": "source",  
      "doc": "A unique identifier of the source service",  
      "type": "string"  
    },  
    {  
      "name": "type",  
      "doc": "The type of event the payload refers to. For example  
LoanApprovedBusinessEvent",  
      "type": "string"  
    },  
    {  
      "name": "category",  
      "doc": "The category of event the payload refers to. For example LOAN",  
      "type": "string"  
    },  
    {  
      "name": "createdAt",  
      "doc": "The UTC time of when the event has been raised; in  
ISO_LOCAL_DATE_TIME format. For example 2011-12-03T10:15:30",  
      "type": "string"  
    },  
    {  
      "name": "businessDate",  
      "doc": "The business date when the event has been raised; in ISO_LOCAL_DATE  
format. For example 2011-12-03",  
      "type": "string"  
    }  
  ]  
}
```

```

},
{
  "name": "tenantId",
  "doc": "The tenantId that the event has been sent from. For example default",
  "type": "string"
},
{
  "name": "idempotencyKey",
  "doc": "The idempotency key for this particular event for consumer de-duplication",
  "type": "string"
},
{
  "name": "dataschema",
  "doc": "The fully qualified name of the schema of the event payload. For example org.apache.fineract.avro.loan.v1.LoanAccountDataV1",
  "type": "string"
},
{
  "name": "data",
  "doc": "The payload data serialized into Avro bytes",
  "type": "bytes"
}
]
}

```

Bulk event schema

The bulk event schema is used when multiple events are supposed to be sent together. This schema is used also when serializing the data for the database storing but the idea is quite simple. Have an array of other event schemas embedded into it.

Since Avro is strongly typed, the array within the bulk event schema is an array of [MessageV1](#) schemas. That way the consumers can decide which events they want to deserialize and which don't.

This elevates the regular 2-level serialization/deserialization concept up to a 3-level one:

1. Serializing the standard events
2. Serializing the standard events into a bulk event
3. Serializing the bulk event into an event message

Versioning

Avro is quite strict with changes to an existing schema and there are a number of compatibility modes available.

Fineract keeps it simple though. Version numbers - in the package names and in the schema names -

are increased with each published modification; meaning that if the `OfficeDataV1` schema needs a new field and the `OfficeDataV1` schema has been published officially with Fineract, a new `OfficeDataV2` has to be created with the new field instead of modifying the existing schema.

This pattern ensures that a certain event is always deserialized with the appropriate schema definition, otherwise the deserialization could fail.

Code generation

The Avro schemas are described as JSON documents. That's hardly usable directly with Java hence Fineract generates Java POJOs from the Avro schemas. The good thing about these POJOs is the fact that they can be serialized/deserialized in themselves without any magic since they have a `toByteBuffer` and `fromByteBuffer` method.

From POJO to ByteBuffer:

```
LoanAccountDataV1 avroDto = ...
ByteBuffer buffer = avroDto.toByteBuffer();
```

From ByteBuffer to POJO:

```
ByteBuffer buffer = ...
LoanAccountDataV1 avroDto = LoanAccountDataV1.fromByteBuffer(buffer);
```



The ByteBuffer is a stateful container and needs to be handled carefully. Therefore Fineract has a built-in ByteBuffer to byte array converter; `ByteBufferConverter`.

Downstream event consumption

When consuming events on the other side of the message channel, it's critical to know which events the system is interested in. With the multi-level serialization, it's possible to deserialize only parts of the message and decide based on that whether it makes sense for a particular system to deserialize the event payload more.

Whether events are important can be decided based on:

- the `type` attribute in the message
- the `category` attribute in the message
- the `dataschema` attribute in the message

These are the main attributes in the message wrapper one can use to decide whether an event message is useful.

If the event needs to be serialized, the next step is to find the corresponding schema definition.

That's going to be sent in the `dataschema` attribute within the message wrapper. Since the attribute contains the fully-qualified name of the respective schema, it can be easily resolved to a Class object. Based on that class, the payload data can be easily deserialized using the `fromByteBuffer` method on every generated schema POJO.

Message ordering

One of the requirements for the framework is to provide ordering guarantees. All the events have to conform a happens-before relation.

For the downstream consumers, this can be verified by the `id` attribute within the messages. Since it's going to be a strictly-monotonic numeric sequence, it can be used for ordering purposes.

Event categorization

For easier consumption, the terminology event category is introduced. This is nothing else but the bounded context an event is related to.

For example the `LoanApprovedBusinessEvent` and the `LoanWaiveInterestBusinessEvent` are both related to the `Loan` bounded contexts.

The category in which an event resides in is included in the message under the `category` attribute.

The existing event categories can be found under the [Event categories](#) section.

Asynchronous event processor

The events stored in the database will be picked up and sent by a regularly executed job.

This job is a Fineract job, scheduled to run for every minute and will pick a number of events in order. Those events will be put onto the downstream message channel in the same order as they were raised.

Purging events

The events database table is going to grow continuously. That's why Fineract has a purging functionality in place that's gonna delete old and already sent events.

It's implemented as a Fineract job and is disabled by default. It's called TBD.

Usage

Using the event framework is quite simple. First, it has to be enabled through properties or environment variable.

The respective options are the following:

- the `fineract.events.external.enabled` property

- the `FINERACT_EXTERNAL_EVENTS_ENABLED` environment variable

These configurations accept a boolean value; `true` or `false`.

The key component to interact with is the `BusinessEventNotifierService#notifyPostBusinessEvent` method.

Raising events

Raising events is really easy. An instance of a `BusinessEvent` interface is needed, that's going to be the event. There are plenty of them available already in the Fineract codebase.

And that's pretty much it. Everything else is taken care of in terms of event data persisting and later on putting it onto a message channel.

An example of event raising:

```
@Override
public CommandProcessingResult createClient(final JsonCommand command) {
    ...
    businessEventNotifierService.notifyPostBusinessEvent(new ClientCreateBusinessEvent
    (newClient));
    ...
    return ...;
}
```



The above code is copied from the `ClientWritePlatformServiceJpaRepositoryImpl` class.

Example event message content

Since the message is serialized into binary format, it's hard to represent in the documentation therefore here's a JSON representation of the data, just as an example.

```
{
    "id": 121,
    "source": "a65d759d-04f9-4ddf-ac52-34fa5d1f5a25",
    "type": "LoanApprovedBusinessEvent",
    "category": "Loan",
    "createdAt": "2022-09-05T10:15:30",
    "tenantId": "default",
    "idempotencyKey": "abda146d-68b5-48ca-b527-16d2b7c5daef",
    "dataschema": "org.apache.fineract.avro.loan.v1.LoanAccountDataV1",
    "data": "..."
}
```



The source attribute refers to an ID that's identifying the producer service. Fineract will regenerate this ID upon each application startup.

Raising bulk events

Raising bulk events is really easy as well. The 2 key methods are:

- `BusinessEventNotifierService#startExternalEventRecording`
- `BusinessEventNotifierService#stopExternalEventRecording`

First, you have to start recording your events. This recording will be applied for the current thread. And then you can raise as many events as you want with the regular `BusinessEventNotifierService#notifyPostBusinessEvent` method, but they won't get saved to the database immediately. They'll get "recorded" into an internal buffer.

When you stop recording using the method above, all the recorded events will be saved as a bulk event to the database; and serialized appropriately.

From then on, the bulk event works just like any of the event. It'll be picked up by the processor to send it to a message channel.

Event categories

TBD

Selective event producing

TBD

Customizations

The framework provides a number of customization options:

- Creating new events (that's already given by the Business Events)
- Creating new Avro schemas
- Customizing what data gets serialized for existing events

In the upcoming sections, that's what going to be discussed.

Creating new events

Creating new events is super easy. Just create an implementation of the `BusinessEvent` interface and that's it.

From then on, you can raise those events in the system, although you can't publish them to an external message channel. If you have the event framework enabled, it's going to fail with not finding the

appropriate serializer for your business event.



There are existing serializers which might be able to handle your new event. For example the `LoanBusinessEventSerializer` is capable of handling all `LoanBusinessEvent` subclasses so there's no need to create a brand new serializer.

The interface looks the following:

`BusinessEvent.java`

```
public interface BusinessEvent<T> {  
  
    T get();  
  
    String getType();  
  
    String getCategory();  
  
    Long getAggregateRootId();  
}
```

Quite simple. The `get` method should return the data you want to pass within the event instance. The `getType` method returns the name of the business event that's gonna be saved as the `type` into the database.



Creating a new business event only means that it can be used for raising an event. To make it compatible with the event framework and to be sent to a message channel, some extra work is needed which are described below.

Creating new Avro schemas and serializers

First let's talk about the event serializers because that's what's needed to make a new event compatible with the framework.

The serializer has a special interface, `BusinessEventSerializer`.

BusinessEventSerializer.java

```
public interface BusinessEventSerializer {  
  
    <T> boolean canSerialize(BusinessEvent<T> event);  
  
    Class<? extends GenericContainer> getSupportedSchema();  
  
    <T> ByteBufferSerializable toAvroDTO(BusinessEvent<T> rawEvent);  
  
}
```

An implementation of this interface shall be registered as a Spring bean, and it'll be picked up automatically by the framework.



You can look at the existing serializers for implementation ideas.

New Avro schemas can be easily created. Just create a new Avro schema file in the [fineract-avro-schemas](#) project under the respective bounded context folder, and it will be picked up automatically by the code generator.

BigDecimal support in Avro schemas

Apache Avro by default doesn't support complex types like a BigDecimal. It has to be implemented using a custom snippet like this:

```
{  
    "logicalType": "decimal",  
    "precision": 27,  
    "scale": 8,  
    "type": "bytes"  
}
```

It's a 20 precision and 8 scale BigDecimal.

Obviously it's quite challenging to copy-paste this snippet to every single BigDecimal field, so there's a customization in place for Fineract.

The type **bigdecimal** is supported natively, and you're free to use it like this:

```
{  
  "default": null,  
  "name": "principal",  
  "type": [  
    "null",  
    "bigdecimal"  
  ]  
}
```



This `bigdecimal` type will be simple replaced with the `BigDecimal` snippet showed above during the compilation process.

Custom data serialization for existing events

In case there's a need some extra bit of information within the event message that the default serializers are not providing, you can override this behavior by registering a brand-new custom serializer (as shown above).

Since there's a priority order of serializers, the only thing the custom serializer need to do is to be annotated by the `@Order` annotation or to implement the `Ordered` interface.

An example custom serializer with priority looks the following:

```

@Component
@RequiredArgsConstructor
@Order(Ordered.HIGHEST_PRECEDENCE)
public class CustomLoanBusinessEventSerializer implements BusinessEventSerializer {
    ...

    @Override
    public <T> boolean canSerialize(BusinessEvent<T> event) {
        return ...;
    }

    @Override
    public <T> byte[] serialize(BusinessEvent<T> rawEvent) throws IOException {
        ...
        ByteBuffer buffer = avroDto.toByteBuffer();
        return byteBufferConverter.convert(buffer);
    }

    @Override
    public Class<? extends GenericContainer> getSupportedSchema() {
        return ...;
    }
}

```



All the default serializers are having `Ordered.LOWEST_PRECEDENCE`.

Appendix A: Properties and environment variables

Property name	Environment variable	Default value	Description
<code>fineract.events.external.enabled</code>	<code>FINERACT_EXTERNAL_EVENTS_ENABLED</code>	<code>false</code>	Whether the external event sending is enabled or disabled.

Introducing Advanced payment allocation

Since the first repayment strategy got introduced, many followed, but there was one thing common in them:

- They were hard coding the allocation rules for each transaction type.

By introducing the "Advanced payment allocation" the idea was to have a repayment strategy which was:

- supporting dynamic configuration of the allocation rules for transaction types

- supporting configuration of more fine-grained allocation rules for future installments
- supporting reprocessing of transactions and charges in chronological order

Glossary

*Advanced payment allocation	Ability to configure allocation rules dynamically for transactions
*Payment allocation	Rule that defines which outstanding balance to be paid of first on which installment
*Re-amortization	Transaction amount to be divided into equal portions by the number of future installments and those installments to be paid by these portions.

Capabilities

- Payment allocation should be configurable for transactions:
 - Repayment
 - Goodwill credit
 - Payout refund
 - Merchant refund
 - Charge adjustments
 - etc.
- Can be configured for Loan products
 - Payment allocation rule changes on the loan product will affect only the newly created Loan accounts.
- Chronological reprocess order
 - Transactions (including disbursements) and charges are (re)processed and allocated in chronological order
- Support re-amortization between future installments
 - Transaction amount to be divided into equal portions (based on the number of future installments) and to repay each future installment by the calculated portion.
 - It's not hard coded, but usually the principal portion needs to be allocated first, but if there are still unprocessed amounts, the rest of the outstanding balances are to be allocated based on the rest of the rules
- Main allocation rules (installment level)
 - Past Due Installment(s):
 - Oldest first

- Due Installment(s):
 - Normal installment takes priority over Down-payment installment (if applicable)
- Future Installment(s):
 - Available allocation orders:
 - Next installment first
 - Last installment first
 - Re-amortization*
- Secondary allocation rules
 - Penalty
 - Fee
 - Interest
 - Principal

Configuration

Advanced repayment allocation rules can be configured for the Loan product if "Advanced payment allocation" got selected as repayment strategy.

There will be a (always required) "DEFAULT" transaction type configuration which acts as fallback ruleset, if there are no configured rules for a specific transaction type.

New repayment strategy

- Name: Advanced payment allocation
- Code: advanced-payment-allocation-strategy
- Order: 8

Allocation rules

- Past due penalty
- Past due fee
- Past due principal
- Past due interest
- Due penalty
- Due fee
- Due principal
- Due interest

- In advance penalty
- In advance fee
- In advance principal
- In advance interest

Future installment allocation rules:

- Next installment
- Last installment
- Re-amortization

Example Request

```
{
  ...
  "paymentAllocation": [
    {
      "transactionType": "DEFAULT",
      "paymentAllocationOrder": [
        {
          "paymentAllocationRule": "DUE_PAST_PENALTY",
          "order": 1
        },
        {
          "paymentAllocationRule": "DUE_PAST_FEE",
          "order": 2
        },
        {
          "paymentAllocationRule": "DUE_PAST_INTEREST",
          "order": 3
        },
        ...
        {
          "paymentAllocationRule": "IN_ADVANCE_INTEREST",
          "order": 14
        }
      ],
      "futureInstallmentAllocationRule": "NEXT_INSTALLMENT"
    }
  ],
  ...
}
```

The above request configures the "DEFAULT" allocation rules:

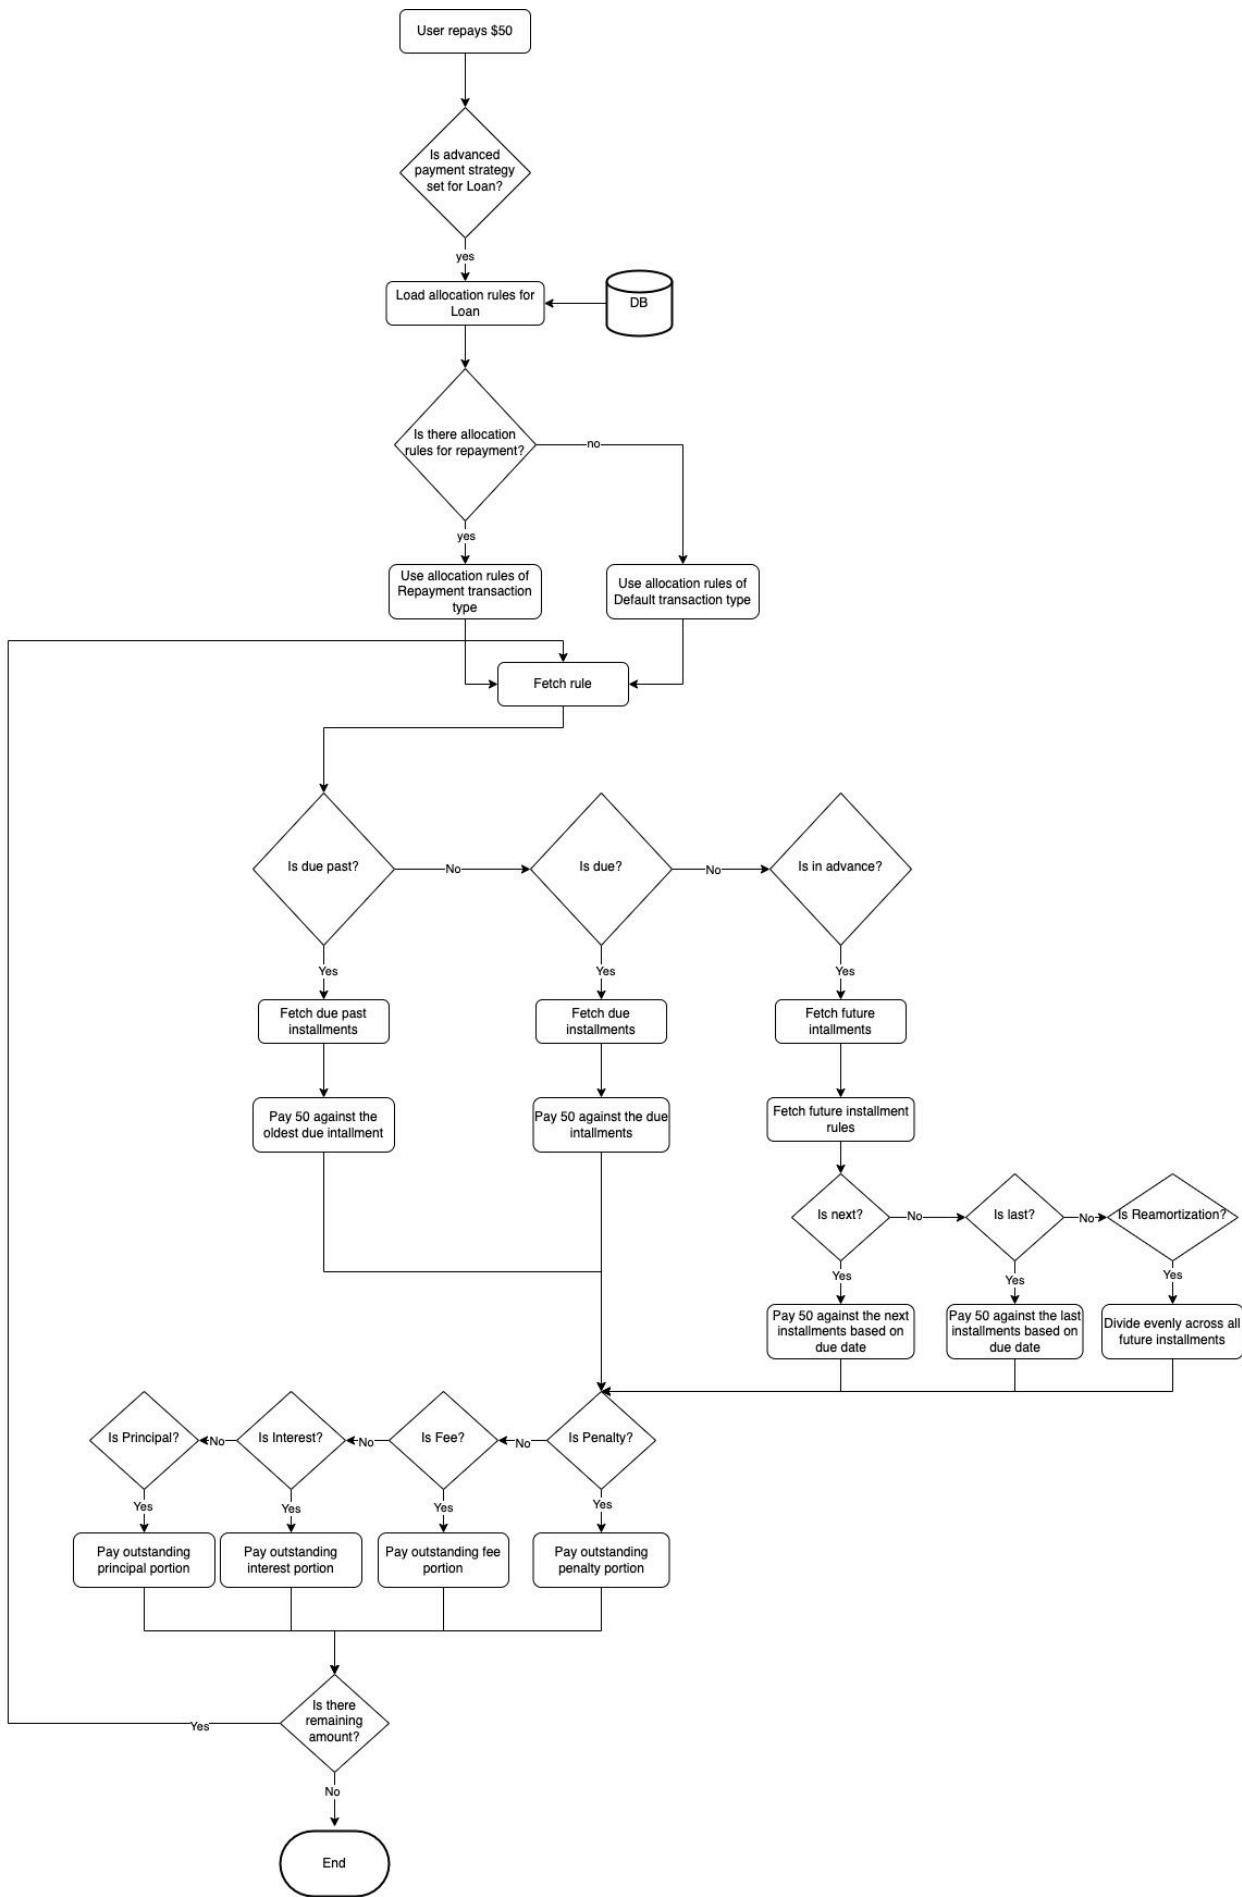
- First the already due penalties to be paid
- Second the already due fees to be paid
- Last the future interests to be paid

Also for future installments set the allocation rules as

- First future installment by due date to be paid first

High level design

Flow of advanced payment allocation processing



Fineract Development Environment

TBD

Git

TBD

GPG

TBD

Committers

Please make sure to provide your GPG fingerprint in your Apache committer profile at id.apache.org.

Docker

TBD

Docker Compose

TBD

Podman

TBD

Rancher Docker Desktop

TBD

Gradle

TBD

IDE

TBD

IntelliJ

TBD

Eclipse

TBD

VSCode

TBD

Kubernetes

TBD

Minikube

TBD

Microk8s

TBD

K3d

TBD

Helm Charts

TBD

Tools

TBD

SDKMAN

We recommend using SKDMAN to manage the following developer tools:

- JDK
- Spring Boot CLI
- Gradle (if you need a global installation)
- AsciidoctorJ

TBD

Brew

MacOS

TBD

Linux

TBD

Custom Modules



Currently, modules are a proof of concept feature in Fineract.

Introduction

Creating customizations for Fineract services is easy. The method described here will work both with our future module guidelines (aka "clean room" modules) and with the intermediary solution we will put in place to avoid major refactorings.

The folder structure for modules is based on a convention that ensures that your extensions don't clash with Fineract's internals. This is to make sure that your downstream forks of Fineract are easy to sync. In the past we had all kinds of strategies to add custom code - including editing existing sources in `fineract-provider`. This is not recommended.



At the moment the only service(s) we prepared to be overridden/replaced are `org.apache.fineract.portfolio.note.service.NoteReadPlatformService` and `org.apache.fineract.portfolio.note.service.NoteWritePlatformService`. Please reach out on the developer mailing list if you need other services.

The recommended folder structure is very simple. If you follow this recommendation you'll get some additional benefits, e. g. you don't even have to edit `settings.gradle` to include your new custom modules. Your modules will also be automatically included in a custom Fineract Docker image build that you can use for your production deployments.

Let's assume your company/org is called "ACME Inc." and you are trying to (fully/partially) replace an existing Fineract service, let's say those in `org.apache.fineract.portfolio.note`. The recommended folder structure would then look something like this:



As soon as we can publish Fineract module JARs to Maven Central you'll have more freedom to setup your projects (including to setup separate Git repos). But for now please follow these instructions:

1. Create a folder under `custom` and name it according to your company/organisation (e. g. `acme` if your company is `ACME Inc.`); this way your custom modules can't clash even with other companies' modules
2. Under your company folder create a folder for the `category` or `domain` your module is targeting; e. g. "loan", "client", "account" etc.
3. Finally, setup `library` folders for the actual modules you want to create; usually that will be to replace/extend some existing service, so there could be a `service` folder, maybe even a `core` folder, e. g. if you want to add additional DTOs etc.; we have also an example for COB business steps
4. Per `category/domain` you should have a `starter` library; means: a Spring Boot auto-configuration setup that makes including your module in Fineract easier ("hands-free"); the necessary parts for a auto-configuration library are a Spring Java configuration class (annotated with `@Configuration`) and a text file at `META-INF/spring/org.springframework.boot.autoconfigure.AutoConfiguration.imports` in your starter resource folder:

```
com.acme.fineract.portfolio.note.starter.AcmeNoteAutoConfiguration
```

Please make sure that your module libraries have proper `build.gradle` files:

```
description = 'ACME Fineract Note Service'  
  
group = 'com.acme.fineract'  
  
base {  
    archivesName = 'acme-fineract-note-service'  
}  
  
apply from: 'dependencies.gradle'
```



You don't need to edit `settings.gradle` to add your modules/libraries. If you follow above convention they'll get included automatically.

5. The dependency.gradle file could look something like this:

```
dependencies {  
    implementation(project(':fineract-core'))  
    implementation(project(':fineract-provider'))  
    compileOnly('org.springframework.boot:spring-boot-autoconfigure')  
}
```



We've included by default some basic and useful dependencies for all custom modules, like Slf4j, Lombok, the usual testing frameworks (JUnit, Cucumber, Mockito etc.)



Do not include your custom module in `fineract-provider's dependency.gradle file. This creates a circular dependency and will fail your build.

Custom Services



We are still trying to figure out which internal services make most sense to be pluggable. Please join the discussion and let us know if you have a specific requirement.

Note Service

The Note service is responsible for ... TBD



We chose the note service because it's interface is very simple and has not many cross dependencies.

Interfaces

Note Read Service Interface

```
package org.apache.fineract.portfolio.note.service;

import java.util.Collection;
import org.apache.fineract.portfolio.note.data.NoteData;

public interface NoteReadPlatformService {

    NoteData retrieveNote(Long noteId, Long resourceId, Integer noteTypeId);

    Collection<NoteData> retrieveNotesByResource(Long resourceId, Integer noteTypeId);
}
```

Note Write Service Interface

```
package org.apache.fineract.portfolio.note.service;

import java.util.Collection;
import org.apache.fineract.portfolio.note.data.NoteData;

public interface NoteReadPlatformService {

    NoteData retrieveNote(Long noteId, Long resourceId, Integer noteTypeId);

    Collection<NoteData> retrieveNotesByResource(Long resourceId, Integer noteTypeId);
}
```

Auto Start Configuration

The rules to replace the Note services are very simple. If you provide an alternative implementation of the services then the default implementations will **not** be loaded.

Note Auto Starter Configuration

```
package org.apache.fineract.portfolio.note.starter;

import org.apache.fineract.portfolio.client.domain.ClientRepositoryWrapper;
import org.apache.fineract.portfolio.group.domain.GroupRepository;
import org.apache.fineract.portfolio.loanaccount.domain.LoanRepositoryWrapper;
import org.apache.fineract.portfolio.loanaccount.domain.LoanTransactionRepository;
import org.apache.fineract.portfolio.note.domain.NoteRepository;
import
org.apache.fineract.portfolio.note.serialization.NoteCommandFromApiJsonDeserializer;
import org.apache.fineract.portfolio.note.service.NoteReadPlatformService;
import org.apache.fineract.portfolio.note.service.NoteReadPlatformServiceImpl;
import org.apache.fineract.portfolio.note.service.NoteWritePlatformService;
import
org.apache.fineract.portfolio.note.service.NoteWritePlatformServiceJpaRepositoryImpl;
import org.apache.fineract.portfolio.savings.domain.SavingsAccountRepository;
import org.springframework.boot.autoconfigure.condition.ConditionalOnMissingBean;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;
import org.springframework.jdbc.core.JdbcTemplate;

@Configuration
public class NoteAutoConfiguration {

    @Bean
    @ConditionalOnMissingBean
    public NoteReadPlatformService noteReadPlatformService(JdbcTemplate jdbcTemplate) {
        return new NoteReadPlatformServiceImpl(jdbcTemplate);
    }

    @Bean
    @ConditionalOnMissingBean
    public NoteWritePlatformService noteWritePlatformService(NoteRepository
noteRepository, ClientRepositoryWrapper clientRepository,
            GroupRepository groupRepository, LoanRepositoryWrapper loanRepository,
            LoanTransactionRepository loanTransactionRepository,
            NoteCommandFromApiJsonDeserializer fromApiJsonDeserializer,
            SavingsAccountRepository savingsAccountRepository) {
        return new NoteWritePlatformServiceJpaRepositoryImpl(noteRepository,
clientRepository, groupRepository, loanRepository,
            loanTransactionRepository, fromApiJsonDeserializer,
            savingsAccountRepository);
    }
}
```

Custom Business Steps

It is very easy to add your own business steps to Fineract's default steps:

1. Create a custom module (e. g. `custom/acme/steps`, follow the instructions on how to create a custom module)
2. Create a class that implements interface `org.apache.fineract.cob.COBBusinessStep`
3. Provide the custom database migration to add the necessary information about your business step in table `m_batch_business_steps`

Business Step Interface

```
package org.apache.fineract.cob;

import org.apache.fineract.infrastructure.core.domain.AbstractPersistableCustom;

public interface COBBusinessStep<T extends AbstractPersistableCustom<Long>> {

    T execute(T input);

    String getEnumStyledName();

    String getHumanReadableName();
}
```

Business Step Implementation

Custom Business Step Implementation Example

```
package com.acme.fineract.loan.cob;

import lombok.RequiredArgsConstructor;
import lombok.extern.slf4j.Slf4j;
import org.apache.fineract.cob.loan.LoanCOBBusinessStep;
import org.apache.fineract.portfolio.loanaccount.domain.Loan;
import org.apache.fineract.portfolio.loanaccount.domain.LoanAccountDomainService;
import org.springframework.beans.factory.InitializingBean;
import org.springframework.stereotype.Component;

@Slf4j
@Component
@RequiredArgsConstructor
public class AcmeNoopBusinessStep implements LoanCOBBusinessStep, InitializingBean {

    private static final String ENUM_STYLED_NAME = "ACME_LOAN_NOOP";

    private static final String HUMAN_READABLE_NAME = "ACME Loan Noop";

    // NOTE: just to demonstrate that dependency injection is working
    private final LoanAccountDomainService loanAccountDomainService;

    @Override
    public void afterPropertiesSet() throws Exception {
        log.warn("Acme COB Loan: '{}'", getClass().getCanonicalName());
    }

    @Override
    public Loan execute(Loan input) {
        return input;
    }

    @Override
    public String getEnumStyledName() {
        return ENUM_STYLED_NAME;
    }

    @Override
    public String getHumanReadableName() {
        return HUMAN_READABLE_NAME;
    }
}
```

As you can see this implementation is very simple and doesn't do much. There are some simple conventions though that you should follow implementing your own business steps:

1. Make sure the value returned by method `getEnumStyledName()` is unique; it's a good idea to choose a prefix that reflects the name of your organization (in this example `ACME_`)
2. You have more freedom for the value returned by `getHumanReadableName()`, but it's a good idea to keep this value as unique as possible

Business Step Database Migration

Business Step Database Migration Example

```

<databaseChangeLog xmlns="http://www.liquibase.org/xml/ns/dbchangelog"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.liquibase.org/xml/ns/dbchangelog
http://www.liquibase.org/xml/ns/dbchangelog/dbchangelog-4.1.xsd">
    <changeSet author="acme" id="1">
        <insert tableName="m_batch_business_steps">
            <column name="job_name" value="LOAN_CLOSE_OF_BUSINESS"/>
            <column name="step_name" value="ACME_LOAN_NOOP"/>
            <column name="step_order" value="5"/>
        </insert>
    </changeSet>
</databaseChangeLog>

```



See also chapter about batch jobs in this documentation.

Custom Loan Transaction Processors

Fineract has 7 built-in loan transaction processors:

1. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.CreocoreLoanRepaymentScheduleTransactionProcessor`
2. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.EarlyPaymentLoanRepaymentScheduleTransactionProcessor`
3. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.FineractStyleLoanRepaymentScheduleTransactionProcessor`
4. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.HeavensFamilyLoanRepaymentScheduleTransactionProcessor`
5. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.InterestPrincipalPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor`
6. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.PrincipalInterestPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor`
7. `org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.RBILoanRepaymentScheduleTransactionProcessor`

Default Loan Transaction Processor configuration

```

import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Conditional;
import org.springframework.context.annotation.Configuration;
import org.springframework.context.annotation.Lazy;

@Configuration
public class LoanAccountAutoStarter {

    @Bean
    @Conditional(CreocoreLoanRepaymentScheduleTransactionProcessorCondition.class)
    public CreocoreLoanRepaymentScheduleTransactionProcessor
creocoreLoanRepaymentScheduleTransactionProcessor() {
        return new CreocoreLoanRepaymentScheduleTransactionProcessor();
    }

    @Bean
    @Conditional(EarlyRepaymentLoanRepaymentScheduleTransactionProcessorCondition.class)
    public EarlyPaymentLoanRepaymentScheduleTransactionProcessor
earlyPaymentLoanRepaymentScheduleTransactionProcessor() {
        return new EarlyPaymentLoanRepaymentScheduleTransactionProcessor();
    }

    @Bean
    @Conditional(MifosStandardLoanRepaymentScheduleTransactionProcessorCondition.class)
    public FineractStyleLoanRepaymentScheduleTransactionProcessor
fineractStyleLoanRepaymentScheduleTransactionProcessor() {
        return new FineractStyleLoanRepaymentScheduleTransactionProcessor();
    }

    @Bean
    @Conditional(HeavensFamilyLoanRepaymentScheduleTransactionProcessorCondition.class)
    public HeavensFamilyLoanRepaymentScheduleTransactionProcessor
heavensFamilyLoanRepaymentScheduleTransactionProcessor() {
        return new HeavensFamilyLoanRepaymentScheduleTransactionProcessor();
    }

    @Bean
    @Conditional
(InterestPrincipalPenaltiesFeesLoanRepaymentScheduleTransactionProcessorCondition.class)
    public InterestPrincipalPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor
interestPrincipalPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor() {
        return new
InterestPrincipalPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor();
    }

    @Bean
    @Conditional
(PrincipalInterestPenaltiesFeesLoanRepaymentScheduleTransactionProcessorCondition.class)

```

```

public PrincipalInterestPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor
principalInterestPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor() {
    return new
PrincipalInterestPenaltyFeesOrderLoanRepaymentScheduleTransactionProcessor();
}

```

All default processor implementations are enabled by default, but can also be prevented from being loaded into memory by a simple configuration in [application.properties](#). Use the environment variables you see below in your Kubernetes and Docker Compose deployments to override the default behavior.

Default Loan Transaction Processor Application Properties

```

fineract.partitioned-job.partitioned-job-properties[0].job-name=LOAN_COB
fineract.partitioned-job.partitioned-job-properties[0].chunk-
size=${LOAN_COB_CHUNK_SIZE:100}
fineract.partitioned-job.partitioned-job-properties[0].partition-
size=${LOAN_COB_PARTITION_SIZE:100}
fineract.partitioned-job.partitioned-job-properties[0].thread-pool-core-pool-
size=${LOAN_COB_THREAD_POOL_CORE_POOL_SIZE:5}
fineract.partitioned-job.partitioned-job-properties[0].thread-pool-max-pool-
size=${LOAN_COB_THREAD_POOL_MAX_POOL_SIZE:5}
fineract.partitioned-job.partitioned-job-properties[0].thread-pool-queue-
capacity=${LOAN_COB_THREAD_POOL_QUEUE_CAPACITY:20}

```

Implement Processors

Loan Transaction Processor Interface

```

package org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor;

import java.time.LocalDate;
import java.util.List;
import java.util.Set;
import org.apache.fineract.organisation.monetary.domain.MonetaryCurrency;
import org.apache.fineract.organisation.monetary.domain.Money;
import org.apache.fineract.portfolio.loanaccount.domain.ChangedTransactionDetail;
import org.apache.fineract.portfolio.loanaccount.domain.LoanCharge;
import org.apache.fineract.portfolio.loanaccount.domain.LoanRepaymentScheduleInstallment;
import org.apache.fineract.portfolio.loanaccount.domain.LoanTransaction;

public interface LoanRepaymentScheduleTransactionProcessor {
    String getCode();
    String getName();
}

```

```

boolean accept(String s);

/**
 * Provides support for processing the latest transaction (which should be the latest
transaction) against the loan
 * schedule.
 */
void processLatestTransaction(LoanTransaction loanTransaction, TransactionCtx ctx);

/**
 * Provides support for passing all {@link LoanTransaction}'s so it will completely
re-process the entire loan
 * schedule. This is required in cases where the {@link LoanTransaction} being
processed is in the past and falls
 * before existing transactions or and adjustment is made to an existing in which
case the entire loan schedule
 * needs to be re-processed.
 */
ChangedTransactionDetail reprocessLoanTransactions(LocalDate disbursementDate, List
<LoanTransaction> repaymentsOrWaivers,
MonetaryCurrency currency, List<LoanRepaymentScheduleInstallment>
repaymentScheduleInstallments, Set<LoanCharge> charges);

Money handleRepaymentSchedule(List<LoanTransaction> transactionsPostDisbursement,
MonetaryCurrency currency,
List<LoanRepaymentScheduleInstallment> installments, Set<LoanCharge>
loanCharges);

/**
 * Used in interest recalculation to introduce new interest only installment.
 */
boolean isInterestFirstRepaymentScheduleTransactionProcessor();
}

```

Custom Loan Transaction Processor Example

```
package com.acme.fineract.loan.processor;

import
org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.FineractStyleL
oanRepaymentScheduleTransactionProcessor;
import org.springframework.stereotype.Component;

@Component
public class AcmeLoanRepaymentScheduleTransactionProcessor extends
FineractStyleLoanRepaymentScheduleTransactionProcessor {

    public static final String STRATEGY_CODE = "acme-standard-strategy";

    public static final String STRATEGY_NAME = "ACME Corp.: standard loan transaction
processing strategy";

    @Override
    public String getCode() {
        return STRATEGY_CODE;
    }

    @Override
    public String getName() {
        return STRATEGY_NAME;
    }

}
```

The example implementation doesn't do much. We are just overriding one of the default processor implementations

`org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.impl.FineractStyleLoanRepaym
entScheduleTransactionProcessor` and give the custom processor its own lookup code and name (descriptive text for display in UIs, e. g. when configuring a loan product). As usual it is a good idea to follow some simple conventions:

1. Make sure the value returned by `getCode()` is unique. Prefixing it with characters that reflect your organization name (here `acme-`) is a good idea.
2. You have more freedom for the descriptive text returned by `getName()`, but it is still a good idea to keep the value unique to avoid confusion.

Method `getCode()`

Lookup value that is used to pick a loan transaction processor (see processor factory).

Method `getName()`

Descriptive text about the loan transaction processor that is mostly used in user interfaces.

Method `handleTransaction()`

TBD

Method `handleWriteOff()`

TBD

Method `handleRepaymentSchedule()`

TBD

Method `isInterestFirstRepaymentScheduleTransactionProcessor()`

TBD

Method `handleRefund()`

TBD

Method `handleChargeback()`

TBD

Method `processTransactionsFromDerivedFields()`

TBD

Override Processor Factory

The processor factory has no reference to any specific implementation of the loan transaction processor interface. All available implementations will be injected here (internal default and custom implementations). Processor instances can be looked up via method `determineProcessor()`. You can pass either the code of the processor or the processor's name to look it up. If a matching processor can't be found then the factory function will either return the default instance or fails with an exception depending on the configuration in `application.properties`.



It is preferable to use the processor code to lookup processor instances. Lookups via processor names are only done in the import service via Excel sheets (should be fixed).

Loan Transaction Processor Factory Implementation

```
package org.apache.fineract.portfolio.loanaccount.domain;

import java.util.List;
import java.util.Optional;
import lombok.RequiredArgsConstructor;
import
org.apache.fineract.portfolio.loanaccount.domain.transactionprocessor.LoanRepaymentScheduleTransactionProcessor;
import
org.apache.fineract.portfolio.loanaccount.exception.LoanTransactionProcessingStrategyNotFoundException;
import org.apache.fineract.portfolio.loanproduct.data.TransactionProcessingStrategyData;
import org.springframework.beans.factory.annotation.Value;

@RequiredArgsConstructor
public class LoanRepaymentScheduleTransactionProcessorFactory {

    private final LoanRepaymentScheduleTransactionProcessor
defaultRepaymentScheduleTransactionProcessor;

    private final List<LoanRepaymentScheduleTransactionProcessor> processors;

    @Value("${fineract.loan.transactionprocessor.error-not-found-fail}")
    private Boolean errorNotFoundFail;

    public LoanRepaymentScheduleTransactionProcessor determineProcessor(final String
transactionProcessingStrategy) {

        Optional<LoanRepaymentScheduleTransactionProcessor> processor = processors.
stream()
            .filter(p -> p.accept(transactionProcessingStrategy)).findFirst();

        if (processor.isEmpty() && Boolean.TRUE.equals(errorNotFoundFail)) {
            throw new LoanTransactionProcessingStrategyNotFoundException
(transactionProcessingStrategy);
        } else {
            return processor.orElse(defaultRepaymentScheduleTransactionProcessor);
        }
    }

    public List<TransactionProcessingStrategyData> getStrategies() {
        return processors.stream().map(p -> new TransactionProcessingStrategyData(null,
p.getCode(), p.getName())).toList();
    }
}
```

This is the default factory auto-configuration.

Loan Transaction Processor Factory Auto-Configuration

```
@Bean
@Conditional(RBIIndiaLoanRepaymentScheduleTransactionProcessorCondition.class)
public RBILoanRepaymentScheduleTransactionProcessor
rbiLoanRepaymentScheduleTransactionProcessor() {
    return new RBILoanRepaymentScheduleTransactionProcessor();
}
```

If you need then you can override this, e.g. because you want to set a different default processor then you can do so in your custom module's auto-configuration.

Custom Loan Transaction Processor Factory Auto-Configuration Example

```
@Bean
public LoanRepaymentScheduleTransactionProcessorFactory
loanRepaymentScheduleTransactionProcessorFactory(
    AcmeLoanRepaymentScheduleTransactionProcessor
defaultLoanRepaymentScheduleTransactionProcessor,
    List<LoanRepaymentScheduleTransactionProcessor> processors) {
    return new LoanRepaymentScheduleTransactionProcessorFactory
(defaultLoanRepaymentScheduleTransactionProcessor, processors);
```

Processor Lookup Failure Configuration Property

```
fineract.partitioned-job.partitioned-job-properties[0].retry-
limit=${LOAN_COB_RETRY_LIMIT:5}
```

Custom Batch Jobs

Fineract provides extension points to define custom batch jobs using module system. Using this approach custom batch jobs can be defined and configured along with Fineract's default batch jobs to extend or customize batch processing.

The batch jobs in Fineract are implemented using [Spring Batch](#). In addition to the Spring Batch ecosystem, automatic scheduling is done by [Quartz Scheduler](#) but it's also possible to trigger batch jobs via regular APIs.

For defining custom job:

1. Create custom module (e. g. [custom/acme/loan/job](#)), follow the instructions on how to create a custom module.
2. Create job configuration to register job, job steps, tasklet with job builder factory. (e. g.

`com.acme.fineract.loan.job.AcmeNoopJobConfiguration`

3. Create tasklet for job execution functionality. (e.g. `com.acme.fineract.loan.job.AcmeNoopJobTasklet`)
4. Provide the custom database migration to add necessary information about your job in table `job`. (e.g. `custom/acme/loan/job/src/main/resources/db/custom-changelog/0001_acme_loan_job.xml`)
5. New job name should be registered along with default jobs so that it can be scheduled at startup. For registering job name with Fineract job scheduler, create an enum with job name details (e.g. `com.acme.fineract.loan.job.AcmeJobName`) and a job name provider configuration which is accessed by Fineract job scheduler at startup to retrieve job name (e.g. `com.acme.fineract.loan.job.AcmeJobNameConfig`).

Job Configuration

Job Configuration Example

```
package com.acme.fineract.loan.job;

import lombok.RequiredArgsConstructor;
import org.springframework.batch.core.Job;
import org.springframework.batch.core.Step;
import org.springframework.batch.core.job.builder.JobBuilder;
import org.springframework.batch.core.launch.support.RunIdIncrementer;
import org.springframework.batch.core.repository.JobRepository;
import org.springframework.batch.core.step.builder.StepBuilder;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;
import org.springframework.transaction.PlatformTransactionManager;

@Configuration
@RequiredArgsConstructor
public class AcmeNoopJobConfiguration {

    private final JobRepository jobRepository;
    private final PlatformTransactionManager transactionManager;
    private final AcmeNoopJobTasklet tasklet;

    @Bean
    protected Step acmeNoopJobStep() {
        return new StepBuilder(AcmeJobName.ACME_NOOP_JOB.name(), jobRepository).tasklet(
            tasklet, transactionManager).build();
    }

    @Bean
    public Job acmeNoopJob() {
        return new JobBuilder(AcmeJobName.ACME_NOOP_JOB.name(), jobRepository).start(
            acmeNoopJobStep()).incrementer(new RunIdIncrementer())
            .build();
    }

}
```

Tasklet Definition

Job Tasklet Example

```
package com.acme.fineract.loan.job;

import lombok.extern.slf4j.Slf4j;
import org.springframework.batch.core.StepContribution;
import org.springframework.batch.core.scope.context.ChunkContext;
import org.springframework.batch.core.step.tasklet.Tasklet;
import org.springframework.batch.repeat.RepeatStatus;
import org.springframework.stereotype.Component;

@Slf4j
@Component
public class AcmeNoopJobTasklet implements Tasklet {

    @Override
    public RepeatStatus execute(StepContribution contribution, ChunkContext chunkContext)
        throws Exception {
        log.info("Acme custom job execution");
        return RepeatStatus.FINISHED;
    }
}
```

Database Migration Script for Job

Database Migration Script Example

```
<databaseChangeLog xmlns="http://www.liquibase.org/xml/ns/dbchangelog"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.liquibase.org/xml/ns/dbchangelog
http://www.liquibase.org/xml/ns/dbchangelog/dbchangelog-4.1.xsd">
    <changeSet author="acme" id="1">
        <insert tableName="job">
            <column name="name" value="Acme Noop Job"/>
            <column name="display_name" value="Acme Noop Job"/>
            <column name="cron_expression" value="0 1 0 1/1 * ? *"/>
            <column name="create_time" valueDate="${current_datetime}"/>
            <column name="task_priority" valueNumeric="5"/>
            <column name="group_name"/>
            <column name="previous_run_start_time"/>
            <column name="job_key" value="Acme Noop Job _ DEFAULT"/>
            <column name="initializing_errorlog"/>
            <column name="is_active" valueBoolean="false"/>
            <column name="currently_running" valueBoolean="false"/>
            <column name="updates_allowed" valueBoolean="true"/>
            <column name="scheduler_group" valueNumeric="0"/>
            <column name="is_misfired" valueBoolean="false"/>
            <column name="node_id" valueNumeric="1"/>
            <column name="is_mismatched_job" valueBoolean="true"/>
        </insert>
    </changeSet>
    <changeSet author="acme" id="2">
        <update tableName="job">
            <column name="short_name" value="ACM_NOOP"/>
            <where>name='Acme Noop Job'</where>
        </update>
    </changeSet>
</databaseChangeLog>
```

Job Name Configuration

Job Name Enum Example

```
package com.acme.fineract.loan.job;

public enum AcmeJobName {

    ACME_NOOP_JOB("Acme Noop Job");

    private final String name;

    AcmeJobName(final String name) {
        this.name = name;
    }

    @Override
    public String toString() {
        return this.name;
    }
}
```

Job Name Provider Configuration Example

```
package com.acme.fineract.loan.job;

import java.util.List;
import org.apache.fineract.infrastructure.jobs.service.jobname.JobNameData;
import org.apache.fineract.infrastructure.jobs.service.jobname.JobNameProvider;
import org.apache.fineract.infrastructure.jobs.service.jobname.SimpleJobNameProvider;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;

@Configuration
public class AcmeJobNameConfig {

    @Bean
    public JobNameProvider acmeJobNameProvider() {
        return new SimpleJobNameProvider(List.of(new JobNameData(AcmeJobName
.ACME_NOOP_JOB.name(), AcmeJobName.ACME_NOOP_JOB.toString())));
    }
}
```

Gradle Build Files

Please make sure that your module libraries have proper `build.gradle` and `dependencies.gradle` files:

Example (`build.gradle`)

```
description = 'ACME Fineract Loan Job'

group = 'com.acme.fineract'

base {
    archivesName = 'acme-fineract-loan-job'
}

apply from: 'dependencies.gradle'
```

Example (`dependencies.gradle`)

```
dependencies {
    implementation(project(':fineract-core'))
    implementation(project(':fineract-loan'))
    implementation(project(':fineract-provider'))
    implementation('org.springframework.batch:spring-batch-integration')
    implementation('org.springframework.boot:spring-boot-starter-data-jpa')
}
```

Deployment

Custom modules can be deployed using docker image. See chapter about deploying custom modules in this documentation.

Example command to build docker image

```
./gradlew :custom:docker:jibDockerBuild
```

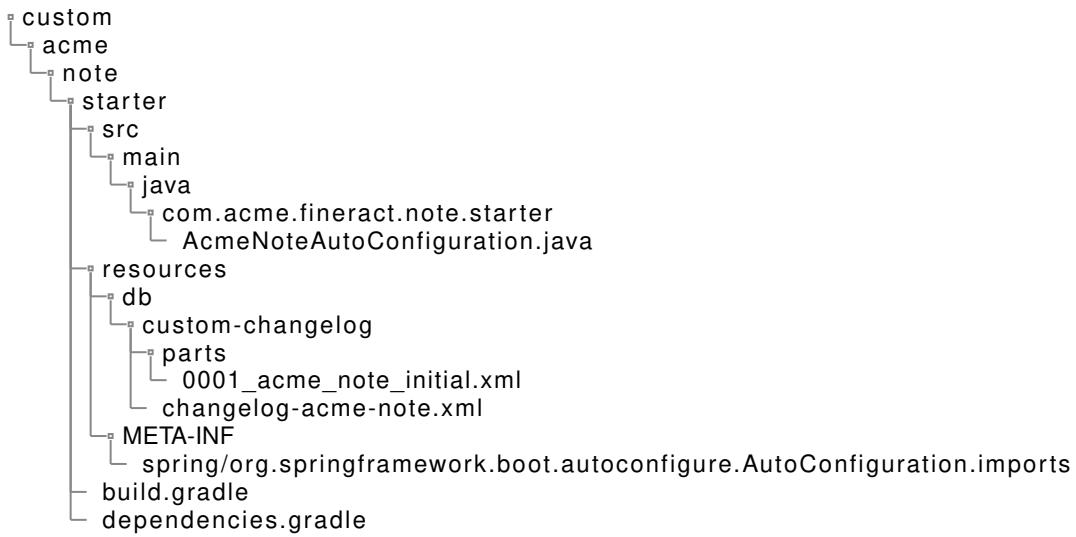


See also chapter about batch jobs in this documentation.

Custom Database Migration

If database migrations are needed as part of your customizations then you can add your own migration scripts. This is again based on conventions:

1. Create folders `db/custom-changelog` in one of your `resources` folders; we recommend using the `resources` folder in your starter library, but actually any of your custom libs will do.
2. Under `db/custom-changelog` create an XML changelog file, e. g. `changelog-acme-note.xml`; you are free to choose a name for this file, but we recommend being consistent to avoid classpath conflicts.
3. Under `db/custom-changelog` create a folder `parts` for your specific changelogs



And here an example migration script:

```

<databaseChangeLog xmlns="http://www.liquibase.org/xml/ns/dbchangelog"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.liquibase.org/xml/ns/dbchangelog
http://www.liquibase.org/xml/ns/dbchangelog/dbchangelog-4.1.xsd">
    <changeSet author="acme" id="1">
        <createTable tableName="acme_note_dummy">
            <column autoIncrement="true" name="id" type="BIGINT">
                <constraints nullable="false" primaryKey="true"/>
            </column>
            <column name="name" type="VARCHAR(100)">
                <constraints unique="true"/>
            </column>
            <column name="description" type="VARCHAR(500)" />
        </createTable>
    </changeSet>
</databaseChangeLog>

```



By default, custom database migration changelogs are executed in context `tenant_db`. That makes sure your changes will be applied to the tenant database (read: main database and not the tenant store database). In theory you could also target the tenant configuration database, but it's not recommended to do that.

Deploying Custom Modules

Custom modules (better: the JAR files) only need to be dropped in Fineract's `libs` folder if you run Fineract from the Spring Boot JAR file. Dynamic loading of external JARs is provided since Fineract version 1.5.0. For your convenience we've created a separate Docker image module that automatically includes your custom modules (see `custom/docker`). You can build this Docker image with

```
./gradlew :custom:docker:jibDockerBuild
```

The Docker image with included custom modules is called **fineract-custom**.



We'll provide soon a way to customize the Docker image parameters (image name, JVM implementation, JVM args, ports etc.).

Outlook

If this **proof of concept** is accepted we could prepare more of Fineract's internal services to be replaceable. This approach works already very well even if we don't have proper JAR libraries published on Maven Central. It's an important goal to separate customized code from Fineract's internals to have soon real modules.

Resilience

Introduction Resilience

Fineract had handcrafted retry loops in place for the longest time. A typical retry code would have looked like this:

Legacy retry code

```
@Override
@SuppressWarnings("AvoidHidingCauseException")
@SuppressFBWarnings(value = {
    "DMI_RANDOM_USED_ONLY_ONCE" }, justification = "False positive for random
object created and used only once")
public CommandProcessingResult logCommandSource(final CommandWrapper wrapper) {

    boolean isApprovedByChecker = false;
    // check if is update of own account details
    if (wrapper.isUpdateOfOwnUserDetails(this.context.authenticatedUser(wrapper)
.getId())) {
        // then allow this operation to proceed.
        // maker checker doesn't mean anything here.
        isApprovedByChecker = true; // set to true in case permissions have
                                    // been maker-checker enabled by
                                    // accident.
    } else {
        // if not user changing their own details - check user has
        // permission to perform specific task.
        this.context.authenticatedUser(wrapper).validateHasPermissionTo(wrapper
.getTaskPermissionName());
    }
    validateIsUpdateAllowed();

    final String json = wrapper.getJson();
    CommandProcessingResult result = null;
    JsonCommand command;
    int numberOfRetries = 0; ①
    int maxNumberOfRetries = ThreadLocalContextUtil.getTenant().getConnection()
.getMaxRetriesOnDeadlock();
    int maxIntervalBetweenRetries = ThreadLocalContextUtil.getTenant().getConnection
().getMaxIntervalBetweenRetries();
    final JsonElement parsedCommand = this.fromApiJsonHelper.parse(json);
    command = JsonCommand.from(json, parsedCommand, this.fromApiJsonHelper, wrapper
.getEntityName(), wrapper.getEntityId(),
    wrapper.getSubentityId(), wrapper.getGroupId(), wrapper.getClientId(),
    wrapper.getLoanId(), wrapper.getSavingId(),
```

```

        wrapper.getTransactionId(), wrapper.getHref(), wrapper.getProductId(),
wrapper.getCreditBureauId(),
        wrapper.getOrganisationCreditBureauId(), wrapper.getJobName());
    while (numberOfRetries <= maxNumberOfRetries) { ②
        try {
            result = this.processAndLogCommandService.executeCommand(wrapper,
command, isApprovedByChecker);
            numberOfRetries = maxNumberOfRetries + 1; ③
        } catch (CannotAcquireLockException | ObjectOptimisticLockingFailureException
exception) {
            log.debug("The following command {} has been retried {} time(s)",
command.json(), numberOfRetries);
            /**
             * Fail if the transaction has been retired for maxNumberOfRetries
             */
            if (numberOfRetries >= maxNumberOfRetries) {
                log.warn("The following command {} has been retried for the max
allowed attempts of {} and will be rolled back",
                    command.json(), numberOfRetries);
                throw exception;
            }
            /**
             * Else sleep for a random time (between 1 to 10 seconds) and continue
             */
            try {
                int randomNum = RANDOM.nextInt(maxIntervalBetweenRetries + 1);
                Thread.sleep(1000 + (randomNum * 1000));
                numberOfRetries = numberOfRetries + 1; ④
            } catch (InterruptedException e) {
                throw exception;
            }
        } catch (final RollbackTransactionAsCommandIsNotApprovedByCheckerException e)
{
    numberOfRetries = maxNumberOfRetries + 1; ③
    result = this.processAndLogCommandService.logCommand(e
.getCommandSourceResult());
}
}

return result;
}

```

① counter

② **while** loop

③ increment to abort

④ increment

For better code quality and readability we introduced [Resilience4j](#):

Annotation based retry

```
private final CommandProcessingService processAndLogCommandService;
private final SchedulerJobRunnerReadService schedulerJobRunnerReadService;
private final ConfigurationDomainService configurationService;

@Override
public CommandProcessingResult logCommandSource(final CommandWrapper wrapper) {
    boolean isApprovedByChecker = false;

    // check if is update of own account details
    if (wrapper.isUpdateOfOwnUserDetails(this.context.authenticatedUser(wrapper)
.getId())) {
        // then allow this operation to proceed.
        // maker checker doesnt mean anything here.
        isApprovedByChecker = true; // set to true in case permissions have
                                    // been maker-checker enabled by
                                    // accident.
    } else {
        // if not user changing their own details - check user has
        // permission to perform specific task.
        this.context.authenticatedUser(wrapper).validateHasPermissionTo(wrapper
.getTaskPermissionName());
    }
    validateIsUpdateAllowed();

    final String json = wrapper.getJson();
    final JsonElement parsedCommand = this.fromApiJsonHelper.parse(json);
    JsonCommand command = JsonCommand.from(json, parsedCommand, this
.fromApiJsonHelper, wrapper.getEntityName(), wrapper.getEntityId(),
    wrapper.getSubentityId(), wrapper.getGroupId(), wrapper.getClientId(),
    wrapper.getLoanId(), wrapper.getSavingsId(),
    wrapper.getTransactionId(), wrapper.getHref(), wrapper.getProductId(),
    wrapper.getCreditBureauId(),
    wrapper.getOrganisationCreditBureauId(), wrapper.getJobName(), wrapper
.getLoanExternalId());
}
```

Command

CommandProcessingService

TBD

Retry-able service function executeCommand

```

private final ToApiJsonSerializer<Map<String, Object>> toApiJsonSerializer;
private final ToApiJsonSerializer<CommandProcessingResult> toApiResultJsonSerializer;
private final ConfigurationDomainService configurationDomainService;
private final CommandHandlerProvider commandHandlerProvider;
private final IdempotencyKeyResolver idempotencyKeyResolver;
private final CommandSourceService commandSourceService;

private final FineractRequestContextHolder fineractRequestContextHolder;
private final Gson gson = GoogleGsonSerializerHelper.createSimpleGson();

@Override
@Retry(name = "executeCommand", fallbackMethod = "fallbackExecuteCommand")
public CommandProcessingResult executeCommand(final CommandWrapper wrapper, final
JsonCommand command,
        final boolean isApprovedByChecker) {
    // Do not store the idempotency key because of the exception handling
    setIdempotencyKeyStoreFlag(false);

    Long commandId = (Long) fineractRequestContextHolder.getAttribute
(COMMAND_SOURCE_ID, null);
    boolean isRetry = commandId != null;
    boolean isEnclosingTransaction = BatchRequestContextHolder.
isEnclosingTransaction();

    CommandSource commandSource = null;
    String idempotencyKey;
    if (isRetry) {
        commandSource = commandSourceService.getCommandSource(commandId);
        idempotencyKey = commandSource.getIdempotencyKey();
    } else if ((commandId = command.commandId()) != null) { // action on the command
itself
        commandSource = commandSourceService.getCommandSource(commandId);
        idempotencyKey = commandSource.getIdempotencyKey();
    } else {
        idempotencyKey = idempotencyKeyResolver.resolve(wrapper);
    }
    exceptionWhenTheRequestAlreadyProcessed(wrapper, idempotencyKey, isRetry);

    AppUser user = context.authenticatedUser(wrapper);
    if (commandSource == null) {
        if (isEnclosingTransaction) {
            commandSource = commandSourceService.getInitialCommandSource(wrapper,
command, user, idempotencyKey);
        } else {
            commandSource = commandSourceService.saveInitialNewTransaction(wrapper,
command, user, idempotencyKey);
            commandId = commandSource.getId();
        }
    }
}

```

```

    }

    if (commandId != null) {
        storeCommandIdInContext(commandSource); // Store command id as a request
attribute
    }

    boolean isMakerChecker = configurationDomainService.isMakerCheckerEnabledForTask
(wrapper.taskPermissionName());
    if (isApprovedByChecker || (isMakerChecker && user.isCheckerSuperUser())) {
        commandSource.markAsChecked(user);
    }
    setIdempotencyKeyStoreFlag(true);

    final CommandProcessingResult result;
    try {
        result = commandSourceService.processCommand(findCommandHandler(wrapper),
command, commandSource, user, isApprovedByChecker,
        isMakerChecker);
    } catch (Throwable t) { // NOSONAR
        RuntimeException mappable = ErrorHandler.getMappable(t);
        ErrorInfo errorInfo = commandSourceService.generateErrorInfo(mappable);
        Integer statusCode = errorInfo.getStatusCode();
        commandSource.setResultStatusCode(statusCode);
        commandSource.setResult(errorInfo.getMessage());
        if (statusCode != SC_OK) {
            commandSource.setStatus(ERROR.getValue());
        }
        if (!isEnclosingTransaction) { // TODO: temporary solution
            commandSource = commandSourceService.saveResultNewTransaction
(commandSource);
        }
        // must not throw any exception; must persist in new transaction as the
current transaction was already
        // marked as rollback
        publishHookErrorEvent(wrapper, command, errorInfo);
        throw mappable;
    }

    commandSource.setResultStatusCode(SC_OK);
    commandSource.updateForAudit(result);
    commandSource.setResult(toApiResponseJsonSerializer.serializeResult(result));
    commandSource.setStatus(PROCESSED.getValue());
}

```

Fallback function `fallbackExecuteCommand`

```
    fineractRequestContextHolder.setAttribute(COMMAND_SOURCE_ID, savedCommandSource
.getId());
}

private void publishHookErrorEvent(CommandWrapper wrapper, JsonCommand command,
ErrorInfo errorInfo) {
    publishHookEvent(wrapper.entityName(), wrapper.actionName(), command, gson.
toJson(errorInfo));
}

private void exceptionWhenTheRequestAlreadyProcessed(CommandWrapper wrapper, String
idempotencyKey, boolean retry) {
    CommandSource command = commandSourceService.findCommandSource(wrapper,
idempotencyKey);
```

Retry configuration for `executeCommand`

```
fineract.report.export.s3.bucket=${FINERACT_REPORT_EXPORT_S3_BUCKET_NAME:}
fineract.report.export.s3.enabled=${FINERACT_REPORT_EXPORT_S3_ENABLED:false}

fineract.jpa.statementLoggingEnabled=${FINERACT_STATEMENT_LOGGING_ENABLED:false}
fineract.database.defaultMasterPassword=${FINERACT_DEFAULT_MASTER_PASSWORD:fineract}
```

Jobs

SchedularWritePlatformService



This service has a typo and should be called `SchedulerWritePlatformService`.

TBD

Retry-able service function `processJobDetailForExecution`

```
@Transactional
@Override
@Retry(name = "processJobDetailForExecution", fallbackMethod =
"fallbackProcessJobDetailForExecution")
public boolean processJobDetailForExecution(final String jobKey, final String
triggerType) {
    boolean isStopExecution = false;
    final ScheduledJobDetail scheduledJobDetail = this.scheduledJobDetailsRepository
    .findByJobKeyWithLock(jobKey);
    if (scheduledJobDetail.isCurrentlyRunning() || (triggerType.equals
    (SchedulerServiceConstants.TRIGGER_TYPE_CRON)
        && scheduledJobDetail.getNextRunTime().after(new Date())))) {
        isStopExecution = true;
    }
    final SchedulerDetail schedulerDetail = retriveSchedulerDetail();
    if (triggerType.equals(SchedulerServiceConstants.TRIGGER_TYPE_CRON) &&
    schedulerDetail.isSuspended()) {
        scheduledJobDetail.setTriggerMisfired(true);
        isStopExecution = true;
    } else if (!isStopExecution) {
        scheduledJobDetail.setCurrentlyRunning(true);
        scheduledJobDetail.setMismatchedJob(false);
    }
    this.scheduledJobDetailsRepository.save(scheduledJobDetail);
    return isStopExecution;
}
```

Fallback function `fallbackProcessJobDetailForExecution`

```
@SuppressWarnings("unused")
public boolean fallbackProcessJobDetailForExecution(Exception e) {
    return false;
}
```

Retry configuration for `processJobDetailForExecution`

```
fineract.notification.user-notification-
system.enabled=${FINERACT_USER_NOTIFICATION_SYSTEM_ENABLED:true}
fineract.logging.json.enabled=${FINERACT_LOGGING_JSON_ENABLED:false}

fineract.sampling.enabled=${FINERACT_SAMPLING_ENABLED:false}
```

Loan

LoanWritePlatformService

TBD

Retryable service function `recalculateInterest`

```
final Map<String, Object> changes = new LinkedHashMap<>();
changes.put("transactionDate", command.stringValueOfParameterNamed(
"transactionDate"));
changes.put("transactionAmount", command.stringValueOfParameterNamed(
"transactionAmount"));
changes.put("locale", command.locale());
changes.put("dateFormat", command.dateFormat());
changes.put(LoanApiConstants.externalIdParameterName, txnExternalId);

Loan loan = this.loanAssembler.assembleFrom(loanId);
// Build loan schedule generator dto
LocalDate recalculateFrom = loan.
isInterestBearingAndInterestRecalculationEnabled() ? transactionDate : null;
final ScheduleGeneratorDTO scheduleGeneratorDTO = this.loanUtilService
.buildScheduleGeneratorDTO(loan, recalculateFrom, null);
// Domain rule validations
this.loanTransactionValidator.validateRefund(loan, loanTransactionType,
transactionDate, scheduleGeneratorDTO);
// Create payment details
final PaymentDetail paymentDetail = this.paymentDetailWritePlatformService
.createAndPersistPaymentDetail(command, changes);
// Create note
createNote(loan, command, changes);
// Initial transaction ids for journal entry generation
final List<Long> existingTransactionIds = loan.findExistingTransactionIds();
final List<Long> existingReversedTransactionIds = loan
.findExistingReversedTransactionIds();
// Create refund transaction(s)
Pair<LoanTransaction, LoanTransaction> refundTransactions =
loanAccountDomainService.makeRefund(loan, scheduleGeneratorDTO,
loanTransactionType, transactionDate, transactionAmount, paymentDetail,
txnExternalId);
LoanTransaction refundTransaction = refundTransactions.getLeft();
LoanTransaction interestRefundTransaction = refundTransactions.getRight();
// Accrual reprocessing
if (loan.isInterestBearingAndInterestRecalculationEnabled()) {
    loanAccrualsProcessingService.reprocessExistingAccruals(loan);
    loanAccrualsProcessingService.processIncomePostingAndAccruals(loan);
}
```

Fallback function `fallbackRecalculateInterest`

```
loanAccountDomainService.updateAndSavePostDatedChecksForIndividualAccount(loan,  
refundTransaction);  
    if (interestRefundTransaction != null) {  
        loanAccountDomainService.updateAndSavePostDatedChecksForIndividualAccount  
(loan, refundTransaction);  
    }  
    // Collateral management  
    loanAccountDomainService  
.updateAndSaveLoanCollateralTransactionsForIndividualAccounts(loan, refundTransaction);  
    if (interestRefundTransaction != null) {  
        loanAccountDomainService  
.updateAndSaveLoanCollateralTransactionsForIndividualAccounts(loan, refundTransaction);  
    }  
    // Raise business events  
    loanAccrualsProcessingService.processAccrualsOnInterestRecalculation(loan, loan  
.isInterestBearingAndInterestRecalculationEnabled(),  
        false);
```

Retry configuration for `recalculateInterest`

```
fineract.sampling.samplingRate=${FINERACT_SAMPLING_RATE:1000}  
fineract.sampling.sampledClasses=${FINERACT_SAMPLED_CLASSES:}  
fineract.sampling.resetPeriodSec=${FINERACT_SAMPLING_RESET_PERIOD_IN_SEC:60}  
  
fineract.module.investor.enabled=${FINERACT_MODULE_INVESTOR_ENABLED:true}
```

Savings

SavingsAccountWritePlatformService

TBD

Retry-able service function `postInterest`

```
        postInterestOnDate = transactionDate;  
    }  
  
    savingsAccountData = this.savingsAccountInterestPostingService.postInterest  
(mc, today, isInterestTransfer,  
        isSavingsInterestPostingAtCurrentPeriodEnd,  
financialYearBeginningMonth, postInterestOnDate, backdatedTxnsAllowedTill,  
        savingsAccountData);  
  
    if (!backdatedTxnsAllowedTill) {
```

```

        List<SavingsAccountTransactionData> transactions = savingsAccountData
    .getSavingsAccountTransactionData();
        for (SavingsAccountTransactionData accountTransaction : transactions) {
            if (accountTransaction.getId() == null) {
                savingsAccountData.setNewSavingsAccountTransactionData
(accountTransaction);
            }
        }
    }
    savingsAccountData.setExistingTransactionIds(existingTransactionIds);
    savingsAccountData.setExistingReversedTransactionIds
(existingReversedTransactionIds);
}
return savingsAccountData;
}

@Override
public CommandProcessingResult reverseTransaction(final Long savingsId, final Long
transactionId,
    final boolean allowAccountTransferModification, final JsonCommand command) {

    final boolean backdatedTxnsAllowedTill = this.savingAccountAssembler
.getPivotConfigStatus();
    final boolean isBulk = command.booleanPrimitiveValueOfParameterNamed("isBulk");
    final SavingsAccount account = this.savingAccountAssembler.assembleFrom(
savingsId, backdatedTxnsAllowedTill);

    final SavingsAccountTransaction savingsAccountTransaction = this
.savingsAccountTransactionRepository
        .findOneByIdAndSavingsAccountId(transactionId, savingsId);
    if (savingsAccountTransaction == null) {
        throw new SavingsAccountTransactionNotFoundException(savingsId,
transactionId);
    }

    if (!allowAccountTransferModification
        && this.accountTransfersReadPlatformService.isAccountTransfer
(transactionId, PortfolioAccountType.SAVINGS)) {
        throw new PlatformServiceUnavailableException(
"error.msg.saving.account.transfer.transaction.update.not.allowed",
            "Savings account transaction:" + transactionId + " update not allowed
as it involves in account transfer",
            transactionId);
    }
}

```

Fallback function `fallbackPostInterest`

```
boolean isInterestTransfer = false;
LocalDate postInterestOnDate = null;
final MathContext mc = MathContext.DECIMAL64;
boolean postReversals = false;
if (account.isBeforeLastPostingPeriod(transactionDate, backdatedTxnsAllowedTill))
{
    final LocalDate today = DateUtils.getBusinessLocalDate();
    account.postInterest(mc, today, isInterestTransfer,
    isSavingsInterestPostingAtCurrentPeriodEnd, financialYearBeginningMonth,
    postInterestOnDate, isInterestTransfer, postReversals);
} else {
    final LocalDate today = DateUtils.getBusinessLocalDate();
    account.calculateInterestUsing(mc, today, isInterestTransfer,
    isSavingsInterestPostingAtCurrentPeriodEnd,
    financialYearBeginningMonth, postInterestOnDate,
    backdatedTxnsAllowedTill, postReversals);
}
```

Retry configuration for `postInterest`

```
fineract.insecure-http-client=${FINERACT_INSECURE_HTTP_CLIENT:true}
fineract.client-connect-timeout=${FINERACT_CLIENT_CONNECT_TIMEOUT:30}
fineract.client-read-timeout=${FINERACT_CLIENT_READ_TIMEOUT:30}
fineract.client-write-timeout=${FINERACT_CLIENT_WRITE_TIMEOUT:30}

# sql validation
```

Security

TBD

OAuth

Fineract has a (basic) OAuth2 support based on Spring Boot Security. Here's how to use it:

Build

You must re-build the distribution JAR (or WAR) using the special `-Psecurity.oauth` flag:

```
./gradlew bootRun -Psecurity.oauth
```

Downloads from fineract.apache.org, or using e.g. the hub.docker.com/r/apache/fineract container image, or on www.fineract.dev, this will not work, because they have not been built using this flag.

Previous versions of Fineract included a built-in authorisation server for issuing OAuth tokens. However, as the `spring-security-oauth2` package was deprecated and replaced by built-in OAuth support in Spring Security, this is no longer supported as part of the package. Instead, you need to run a separate OAuth authorization server (e.g. github.com/spring-projects/spring-authorization-server) or use a 3rd-party OAuth authorization provider (en.wikipedia.org/wiki/List_of_OAuth_providers)

This instruction describes how to get Fineract OAuth working with a Keycloak (keycloak.org) based authentication provider running in a Docker container. The steps required for other OAuth providers will be similar.

Set up Keycloak

1. From terminal, run: `'docker run -p 9000:8080 -e KEYCLOAK_USER=admin -e KEYCLOAK_PASSWORD=admin quay.io/keycloak/keycloak:15.0.2'`
2. Go to URL '`http://localhost:9000/auth/admin`' and login with `admin/admin`
3. Hover your mouse over text "Master" and click on "Add realm"
4. Enter name "fineract" for your realm
5. Click on tab "Users" on the left, then "Add user" and create user with username "mifos"
6. Click on tab "Credentials" at the top, and set password to "password", turning "temporary" setting to off
7. Click on tab "Clients" on the left, and create client with ID 'community-app'
8. In settings tab, set 'access-type' to 'confidential' and enter 'localhost' in the valid redirect URIs.
9. In credentials tab, copy string in field 'secret' as this will be needed in the step to request the access

token

Finally we need to change Keycloak configuration so that it uses the username as a subject of the token:

1. Choose client 'community-app' in the tab 'Clients'
2. Go to tab 'Mappers' and click on 'Create'
3. Enter 'usernameInSub' as 'Name'
4. Choose mapper type 'User Property'
5. Enter 'username' into the field 'Property' and 'sub' into the field 'Token Claim Name'. Choose 'String' as 'Claim JSON Type'

You are now ready to test out OAuth:

Retrieve an access token from Keycloak

```
curl --location --request POST \
'http://localhost:9000/auth/realms/fineract/protocol/openid-connect/token' \
--header 'Content-Type: application/x-www-form-urlencoded' \
--data-urlencode 'username=mifos' \
--data-urlencode 'password=password' \
--data-urlencode 'client_id=community-app' \
--data-urlencode 'grant_type=password' \
--data-urlencode 'client_secret=<enter the client secret from credentials tab>'
```

The reply should contain a field 'access_token'. Copy the field's value and use it in the API call below:

Invoke APIs and pass **Authorization: bearer ... header**

```
curl --location --request GET \
'https://localhost:8443/fineract-provider/api/v1/offices' \
--header 'Fineract-Platform-TenantId: default' \
--header 'Authorization: bearer <enter the value of the access_token field>'
```



See also fineract.apache.org/legacy-docs/apiLive.htm#authentication_oauth

Testing

TBD

Cucumber

TBD

Cucumber Cheatsheet

Cucumber is a test framework based on Behavior-Driven Development (BDD). Tests are written in plain text with very basic syntax rules. These rules form a mini language that is called Gherkin.

A specification resembles spoken language. This makes it ideal for use with non-technical people that have domain specific knowledge. The emphasis of Cucumber lies on finding examples to describe your test cases. The few keywords and language rules are easy to explain to anyone (compared JUnit for example).

Keywords

The Gherkin language has the following keywords:

- Feature
- Rule
- Scenario Outline or Scenario Template
- Example or Scenario
- Examples or Scenarios
- Background
- Given
- And
- But
- When
- Then

There are a couple of additional signs used in Gherkin:

- | is as column delimiters in Examples tables
- with @ you can assign any kind of tags to categorize the specs (or e.g. relate them to certain Jira tickets)
- # is used to indicate line comments



The tag `@ignore` is used to skip tests. This is a somewhat arbitrary choice (we could use any other tag to indicate temporarily disabled tests).

Each non-empty line of a test specification needs to start with one of these keywords. The text blocks that follows the keywords are mapped to so called step definitions that contain the actual test code.

A typical Cucumber test specification written in Gherkin looks like this:

Feature: Template Service

```
@template
Scenario Outline: Verify that mustache templates have expected results
  Given A mustache template file <template>
  Given A JSON data file <json>
  When The user merges the template with data
  Then The result should match the content of file <result>
```

Examples:

template	json	result
hello.mustache	hello.json	hello.txt
loan.mustache	loan.json	loan.html
array.loop.mustache	array.json	array.loop.txt
array.index.mustache	array.json	array.index.txt

The corresponding step definitions would look like this:

```
package org.apache.fineract.template.service;

import static org.junit.jupiter.api.Assertions.assertEquals;

import com.google.common.reflect.TypeToken;
import com.google.gson.Gson;
import com.google.gson.JsonElement;
import com.google.gson.JsonParser;
import io.cucumber.java8.En;
import java.io.IOException;
import java.lang.reflect.Type;
import java.nio.charset.StandardCharsets;
import java.util.ArrayList;
import java.util.List;
import java.util.Map;
import org.apache.commons.io.IOUtils;
import org.apache.fineract.template.domain.Template;
import org.apache.fineract.template.domain.TemplateMapper;
import org.springframework.beans.factory.annotation.Autowired;
```

```

public class TemplateServiceStepDefinitions implements En {

    @Autowired
    private TemplateMergeService tms;

    private String template;

    private Map<String, Object> data;

    private String result;

    public TemplateServiceStepDefinitions() {
        Given("/^A mustache template file (.*)$/",
            (String file) -> {
                template = IOUtils.resourceToString("templates/" + file, StandardCharsets
                .UTF_8,
                TemplateServiceStepDefinitions.class.getClassLoader());
            });

        Given("/^A JSON data file (.*)$/",
            (String file) -> {
                data = parse(IOUtils.resourceToString("templates/" + file, StandardCharsets
                .UTF_8,
                TemplateServiceStepDefinitions.class.getClassLoader()));
            });

        When("The user merges the template with data", () -> {
            result = compile(template, data);
        });

        Then("/^The result should match the content of file (.*)$/",
            (String file) -> {
                String expected = IOUtils.resourceToString("results/" + file,
                StandardCharsets.UTF_8,
                TemplateServiceStepDefinitions.class.getClassLoader());
                assertEquals(expected, result);
            });
    }

    private String compile(String templateText, Map<String, Object> scope) throws
    IOException {
        List<TemplateMapper> mappers = new ArrayList<>();
        Template template = new Template("TemplateName", templateText, null, null,
        mappers);
        return tms.compile(template, scope);
    }

    private Map<String, Object> parse(String data) {
        Gson gson = new Gson();
        Type ssMap = new TypeToken<Map<String, Object>>() {}.getType();

```

```
        JsonElement json = JsonParser.parseString(data);
        return gson.fromJson(json, ssMap);
    }
}
```



This example is an actual test specification that you can find in the `fineract-provider` module.

Feature

This keyword is used to group scenarios and to group related scenarios. All Gherkin specifications must start with the word `Feature`.

Descriptions

A description is any non-empty line that doesn't start with a keyword. Descriptions can be placed under the keywords:

- `Feature`
- `Rule`
- `Background`
- `Example/Scenario`
- `Scenario Outline`

Rule

Rule is used to group multiple related scenarios together.

Example/Scenario

This is the important part of the specification as it should describe the business logic in more detail with the usage of steps (usually `Given`, `When`, `Then`)

Steps

TBD

Given

TBD

When

TBD

Then

TBD

And, But

TBD

Background

TBD

Scenario Outline

TBD

Examples/Tables

TBD

Outlook

As a proof of concept we've converted all unit tests in [fineract-provider](#) into Cucumber tests. The more interesting part starts when we'll attack the integration tests with over 400 mostly business logic related tests. These tests fit very well in Cucumber's test specification structure (a lot of *if-then-else* or in Gherkin: *Given-When-Then*). Migrating all tests will take a while, but we would already recommend trying to implement tests as Cucumber specifications. It should be relatively easy to convert these tests into the new syntax.

Hopefully this will motivate even more people from the broader Fineract community to participate in the project by sharing their domain specific knowledge as Cucumber specifications. Specifications are written in English (although not a technical requirement).



Have a look at the specifications in [fineract-provider](#) for an initial inspiration. For more information please see [cucumber.io/docs](#)

Unit Testing

TBD

Integration Testing

TBD

Fineract Documentation Guide

TBD

File and Folder Layout

The general rules are

- keep things as flat as possible (avoid sub-folders as much as possible)
- DRY (don't repeat yourself): don't copy and paste code pieces, use AsciiDoc's include feature and reference files/sections from the project folder
- images are located in `fineract-doc/src/docs/en/images` (or sub-folders)
- diagrams are located in `fineract-doc/src/docs/en/diagrams` (or sub-folders)
- specific chapters are located in `fineract-doc/src/docs/en/chapters`
- every chapter has its own folder and at least one `index.adoc` file
- it's recommended to keep the chapters flat (i. e. no sub-folders in the chapter folders)
- it's recommended to create one file per chapter section; like that you can re-arrange sections very easily in the `index.adoc` file



These rules are not entirely set in stone and could be modified if necessary. If you see any issues then please report them on the [mailing list](#) or [open a Jira ticket](#).

AsciiDoc

Cheatsheet

You can find the definitive manual on AsciiDoc syntax at [AsciiDoc documentation](#). To help people get started, however, here is a simpler cheat sheet.

AsciiDoc vs Asciidoctor (format vs tool)

When we refer to *AsciiDoc* then we mean the language or format that this documentation is written in. AsciiDoc is a markup language similar to Markdown (but more powerful and expressive) designed for technical documentation. You don't need necessarily any specialized editors or tools to write your documentation in AsciiDoc, a plain text editor will do, but there are plenty of choices that give you a better experience (in this documentation we describe the basic usage with AsciiDoc plugins for IntelliJ, Eclipse and VSCode).

Asciidoctor on the other hand is the command line tool we use to transform documents written in AsciiDoc into HTML and PDF (Epub3 and Docbook are also available). There are three variants available:

- Asciidoc (written in Ruby)
- Asciidoc.js (written in JavaScript, often used for browser previews)
- AsciidocJ (Java lib that integrates the Ruby implementation via JRuby, e. g. the Asciidoc Gradle plugin is based on that)



Sometimes you will still find documentation related to the original incarnation of AsciiDoc/tor (written in Python). The format evolved quite a bit since then and the tools try to maintain a certain degree of backward compatibility, but there is no guarantee. We prefer to use the latest language specs as documented [here](#).

Basic AsciiDoc Syntax

Bold

Put asterisks around text to make it **bold**.



More info at docs.asciidoc.org/asciidoc/latest/text/bold

Italics

Use underlines on either side of a string to put text into *italics*.



More info at docs.asciidoc.org/asciidoc/latest/text/italic

Headings

Equal signs (=) are used for heading levels. Each equal sign is a level. Each page can **only** have one top level (i.e., only one section with a single =).

Levels should be appropriately nested. During the build, validation occurs to ensure that level 3s are preceded by level 2s, level 4s are preceded by level 3s, etc. Including out-of-sequence heading levels (such as a level 3 then a level 5) will not fail the build, but will produce an error.



More info at docs.asciidoc.org/asciidoc/latest/sections/titles-and-levels/

Code Examples

Use backticks ` for text that should be monospaced, such as code or a class name in the body of a paragraph.



More info at docs.asciidoc.org/asciidoc/latest/text/monospace/

Longer code examples can be separated from text with **source** blocks.

These allow defining the syntax being used so the code is properly highlighted.

Example Source Block

```
[source,xml]
<field name="id" type="string" indexed="true" stored="true" required="true"
multiValued="false" />
```

If your code block will include line breaks, put 4 hyphens (----) before and after the entire block.



More info at docs.asciidoctor.org/asciidoc/latest/verbatim/source-blocks/

Source Block Syntax Highlighting

The HTML output uses Rouge to add syntax highlighting to code examples. This is done by adding the language of the code block after the `source`, as shown in the above example source block (`xml` in that case).

Rouge has a long selection of lexers available. You can see the full list at github.com/rouge-ruby/rouge/wiki/List-of-supported-languages-and-lexers. Use one of the valid short names to get syntax highlighting for that language.

Ideally, we will have an appropriate lexer to use for all source blocks, but that's not possible. When in doubt, choose `text`, or leave it blank.

Importing Code Snippets from Other Files

The build system has the ability to "include" snippets located in other files — even non-AsciiDoc files such as `*.java` source code files.

We've configured a global attribute called `{rootdir}` that you can use to reference these files consistently from Fineract's project root folder.

Snippets are bounded by `tag` comments placed at the start and end of the section you would like to import. Opening tags look like: `// tag::snippetName[]`. Closing tags follow the format: `// end::snippetName[]`.

Snippets can be inserted into an `.adoc` file using an `include` directive, following the format: `include::{rootdir}/<directory-under-root-folder>/<file-name>[tag=snippetName]`.



You could also use relative paths to reference include files, but it is preferred to always use the root folder as a starting point. Like this you can be sure that the preview in your editor of choice works.

For example, if we wanted to highlight a specific section of the following Cucumber test definition (more on that in section Cucumber Testing) `ClasspathDuplicatesStepDefinitions.java` file located under `fineract-provider/src/test/java/org/apache/fineract/infrastructure/classpath/`.

```
[source,java,indent=0]
-----
include::{rootdir}/fineract-
provider/src/test/java/org/apache/fineract/infrastructure/classpath/ClasspathDuplicatesSt
epDefinitions.java[tag=then]
-----
```

For more information on the `include` directive, see the documentation at docs.asciidoctor.org/asciidoc/latest/directives/include.

Block Titles

Titles can be added to most blocks (images, source blocks, tables, etc.) by simply prefacing the title with a period (.). For example, to add a title to the source block example above:

```
.Example ID field
[source,xml]
<field name="id" type="string" indexed="true" stored="true" required="true"
multiValued="false" />
```



More info at docs.asciidoctor.org/asciidoc/latest/blocks/add-title

Links

Link to Sites on the Internet

When converting content to HTML, Asciidoctor will automatically render many link types (such as `http:` and `mailto:`) without any additional syntax. However, you can add a name to a link by adding the URI followed by square brackets:

```
http://fineract.apache.org/[Fineract Website]
```



More info at docs.asciidoctor.org/asciidoc/latest/macros/url-macro

Link to Other Pages/Sections of the Guide

A warning up front, linking to other pages can be a little painful. There are slightly different rules depending on the type of link you want to create, and where you are linking from. The build process includes a validation for *internal* or *inter-page* links, so if you can build the docs locally, you can use that to verify you constructed your link properly. With all the below examples, you can add text to display as the link title by putting the display text in brackets after the link, as in:

xref:indexing-guide:schema-api.adoc#modify-the-schema[Modify the Schema]

You can also use the title of the Page or Section you are linking to by using an empty display text. This is useful in case the title of the page or section changes. In that case you won't need to change the display text for every link that refers to that page/section.

See an example below:

xref:indexing-guide:schema-api.adoc#modify-the-schema[]

Link to a Section on the Same Page

To link to an anchor (or section title) on the *same page*, you can simply use double angle brackets (⟨⟨ >>) around the anchor/heading/section title you want to link to. Any section title (a heading that starts with equal signs) automatically becomes an anchor during conversion and is available for deep linking.

Example

If I have a section on a page that looks like this (from `process.adoc`):

== Steps

Common parameters for all steps are:

To link to this section from another part of the same `process.adoc` page, I simply need to put the section title in double angle brackets, as in:

See also the <<Steps>> section.

The section title will be used as the display text; to customize that add a comma after the the section title, then the text you want used for display.



More info at docs.asciidoc.org/asciidoc/latest/macros/xref/#internal-cross-references

Link to a Section with an Anchor ID

When linking to any section (on the same page or another one), you must also be aware of any pre-defined anchors that may be in use (these will be in double brackets, like [[]]).

When the page is converted, those will be the references your link needs to point to.

Example

Take this example from `configsets-api.adoc`:

```
[[configsets-create]]  
== Create a ConfigSet
```

To link to this section, there are two approaches depending on where you are linking from:

- From the same page, simply use the anchor name: `<<configsets-create>>`.
- From another page, use the page name and the anchor name: `xref:configuration-guide:configsets-api.adoc#configsets-create[]`.

Link to Another Page

To link to *another page* or a section on another page, you must refer to the full filename and refer to the section you want to link to.

When you want to refer the reader to another page without deep-linking to a section, Asciidoctor allows this by merely omitting the `#` and section id.

Example

To construct a link to the `process.adoc` page, we need to refer to the file name (`process.adoc`), as well as the module that the file resides in (`release/`).

It's preferred to also always use the page name to give the reader better context for where the link goes.

As in:

```
For more about upgrades, see xref:release:process.adoc[Fineract Release Process].
```

Link to Another Page in the same folder

If the page that contains the link and the page being linked to reside in the same module, there is no need to include the module name after `xref`:

Example

To construct a link to the `process-step01.adoc` page from `process.adoc` page, we do not need to include the module name because they both reside in the `upgrade-notes` module.

```
For more information on the first step of the release process, see the section  
\xref:process-step01.adoc[].
```

Link to a Section on Another Page

Linking to a section is the same conceptually as linking to the top of a page, you just need to take a little extra care to format the anchor ID in your link reference properly.

When you link to a section on another page, you must make a simple conversion of the title into the format of the section ID that will be created during the conversion. These are the rules that transform the sections:

Example

TBD



More info at docs.asciidoc.org/asciidoc/latest/macros/inter-document-xref

Ordered and Unordered Lists

AsciiDoc supports three types of lists:

- Unordered lists
- Ordered lists
- Labeled lists

Each type of list can be mixed with the other types. So, you could have an ordered list inside a labeled list if necessary.

Unordered Lists

Simple bulleted lists need each line to start with an asterisk (*). It should be the first character of the line, and be followed by a space.



More info at docs.asciidoc.org/asciidoc/latest/lists/unordered

Ordered Lists

Numbered lists need each line to start with a period (.). It should be the first character of the line, and be followed by a space. This style is preferred over manually numbering your list.



More info at docs.asciidoc.org/asciidoc/latest/lists/ordered

Description Lists

These are like question & answer lists or glossary definitions.

Each line should start with the list item followed by double colons (::), then a space or new line. Labeled lists can be nested by adding an additional colon (such as :::, etc.). If your content will span multiple paragraphs or include source blocks, etc., you will want to add a plus sign (+) to keep the sections together for your reader.



We prefer this style of list for parameters because it allows more freedom in how you present the details for each parameter. For example, it supports ordered or unordered lists inside it automatically, and you can include multiple paragraphs and source blocks without trying to cram them into a smaller table cell.



More info at docs.asciidoctor.org/asciidoc/latest/lists/description

Images

There are two ways to include an image: inline or as a block. Inline images are those where text will flow around the image. Block images are those that appear on their own line, set off from any other text on the page. Both approaches use the `image` tag before the image filename, but the number of colons after `image` define if it is inline or a block. Inline images use one colon (`image:`), while block images use two colons (`image::`). Block images automatically include a caption label and a number (such as `Figure 1`). If a block image includes a title, it will be included as the text of the caption. Optional attributes allow you to set the alt text, the size of the image, if it should be a link, float and alignment. We have defined a global attribute `{imagesdir}` to standardize the location for all images (`fineract-doc/src/docs/en/images`).



More info at docs.asciidoctor.org/asciidoc/latest/macros/images

Tables

Tables can be complex, but it is pretty easy to make a basic table that fits most needs.

Basic Tables

The basic structure of a table is similar to Markdown, with pipes (|) delimiting columns between rows:

```
|===
| col 1 row 1 | col 2 row 1|
| col 1 row 2 | col 2 row 2|
|==
```

Note the use of `|==` at the start and end. For basic tables that's not exactly required, but it does help to delimit the start and end of the table in case you accidentally introduce (or maybe prefer) spaces between the rows.

Header Rows

To add a header to a table, you need only set the `header` attribute at the start of the table:

```
[options="header"]
|===
| header col 1 | header col 2|
| col 1 row 1 | col 2 row 1|
| col 1 row 2 | col 2 row 2|
|===
```

Defining Column Styles

If you need to define specific styles to all rows in a column, you can do so with the attributes.

This example will center all content in all rows:

```
[cols="2*^" options="header"]
|===
| header col 1 | header col 2|
| col 1 row 1 | col 2 row 1|
| col 1 row 2 | col 2 row 2|
|===
```

Alignments or any other styles can be applied only to a specific column. For example, this would only center the last column of the table:

```
[cols="2*,^" options="header"]
|===
| header col 1 | header col 2|
| col 1 row 1 | col 2 row 1|
| col 1 row 2 | col 2 row 2|
|===
```

Many more examples of formatting:



- Columns: docs.asciidoc.org/asciidoc/latest/tables/add-columns/
- Cells and rows: docs.asciidoc.org/asciidoc/latest/tables/add-cells-and-rows/

More Options

Tables can also be given footer rows, borders, and captions. You can determine the width of columns, or the width of the table as a whole.

CSV or DSV can also be used instead of formatting the data in pipes.



More info at docs.asciidoc.org/asciidoc/latest/tables/build-a-basic-table/

Admonitions (Notes, Warnings)

AsciiDoc supports several types of callout boxes, called "admonitions":

- NOTE
- TIP
- IMPORTANT
- CAUTION
- WARNING

It is enough to start a paragraph with one of these words followed by a colon (such as `NOTE:`). When it is converted to HTML, those sections will be formatted properly - indented from the main text and showing an icon inline.

You can add titles to admonitions by making it an admonition block. The structure of an admonition block is like this:

```
.Title of Note
[NOTE]
=====
Text of note
=====
```

In this example, the type of admonition is included in square brackets (`[NOTE]`), and the title is prefixed with a period. Four equal signs give the start and end points of the note text (which can include new lines, lists, code examples, etc.).



More info at docs.asciidoc.org/asciidoc/latest/blocks/admonitions/

STEM Notation Support

We have set up the Ref Guide to be able to support STEM notation whenever it's needed.

The `AsciiMath` syntax is supported by default, but LaTeX syntax is also available.

To insert a mathematical formula inline with your text, you can simply write:

```
stem:[a//b]
```

MathJax.js will render the formula as proper mathematical notation when a user loads the page. When the above example is converted to HTML, it will look like this to a user: `a//b`

To insert LaTeX, preface the formula with `texmath` instead of `stem`:

```
latexmath:[tp \leq 1 - (1 - sim^{rows})^{bands}]
```

Long formulas, or formulas which should be set off from the main text, can use the block syntax prefaced by `stem` or `latexmath`:

```
[stem]
++++
sqrt(3x-1)+(1+x)^2 < y
++++
```

or for LaTeX:

```
[latexmath]
++++
[tp \leq 1 - (1 - sim^{rows})^{bands}]
++++
```



More info at docs.asciidoctor.org/asciidoc/latest/stem/stem

Diagrams

TBD



docs.asciidoctor.org/diagram-extension/latest

PlantUML Cheatsheet

TBD



plantuml.com

C4 Cheatsheet

TBD



c4model.com and github.com/plantuml-stdlib/C4-PlantUML

Archimate Cheatsheet

TBD



www.opengroup.org/archimate-forum/archimate-overview-and-archimate-diagram and plantuml.com/

Vega Cheatsheet

TBD



vega.github.io/vega-lite/

Editor

TBD

IntelliJ IDEA

Asciidoc

See: plugins.jetbrains.com/plugin/7391-asciidoc

PlantUML

See: plugins.jetbrains.com/plugin/7017-plantuml-integration

Eclipse

Asciidoc

See: marketplace.eclipse.org/content/asciidoc-asciidoctor-editor

PlantUML

See: plantuml.com/eclipse

VSCode

Asciidoc

See: marketplace.visualstudio.com/items?itemName=asciidoctor.asciidoctor-vscode

PlantUML

See: marketplace.visualstudio.com/items?itemName=jebbs.plantuml

Antora

See: antora.org/

TBD

Releases

[How to Release Apache Fineract](#) documents the process how we make the source code that is available here in this Git repository into a binary release tar.gz available on [fineract.apache.org](#).

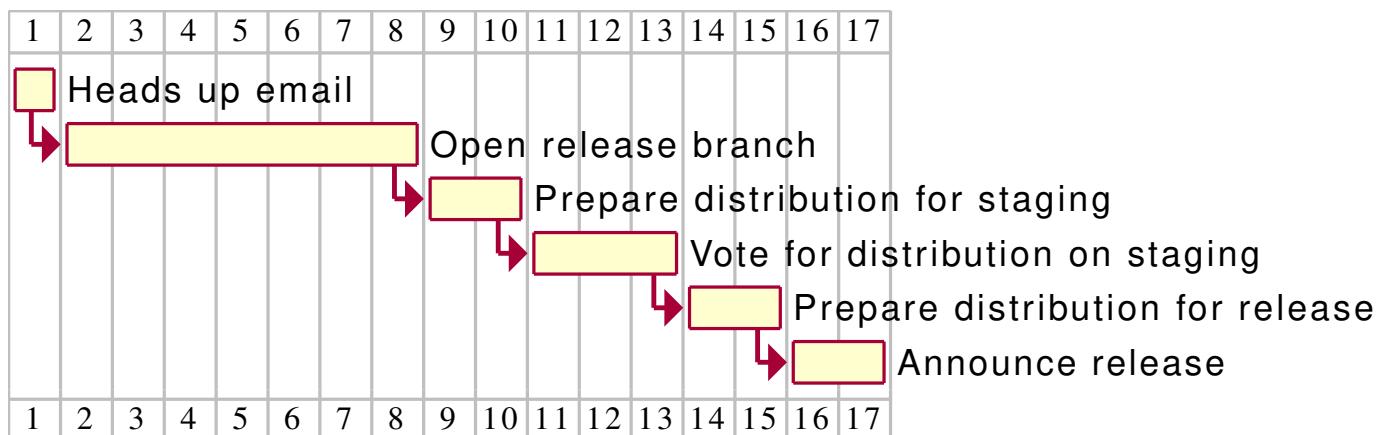


Figure 4. Release Schedule

Configuration

Before you can start using the Fineract release plugin to create releases you have to configure and setup a couple of things first.

- All official communication concerning releases happens on the [mailing list](#). Every release manager needs to be a member of and engaging on the mailing list for credibility.
- Make sure you have edit permissions on the [Apache Confluence Wiki](#)
- You need full permissions on [Apache JIRA](#) to be able to move issues to the next release
- Git committer privileges to be allowed to create tags and the release branch
- Familiarity with building Fineract locally and creating release distributions is required
- You need to be a member of the PMC to be able to upload release artifacts; this task can be delegated though
- A general Familiarity with PGP/GPG is recommended (at least to setup your keypairs), but the release plugin does most of the heavy lifting
- Make sure to read the release plugin documentation for troubleshooting

Secrets

TBD

Infrastructure Team

A couple of secrets for third party services are automatically configured by the infrastructure team at

The Apache Foundation for the Fineract Github account. At the moment this includes environment variables for:

- Github token (e. g. to publish Github Pages, use the Github API in Github Actions)
- Docker Hub token (to publish our Docker images)
- Sonar Cloud token (for our code quality reports)

See also:

- infra.apache.org/github-pages.html
- cwiki.apache.org/confluence/display/INFRA/Github+Actions+to+DockerHub
- github.com/apache/jmeter-site-preview
- github.com/apache/fineract-site
- github.com/apache/systemds-website/blob/main/.asf.yaml

Lastpass

It seems that Apache has some kind of org account or similar. Popped up a couple of times in the infrastructure documentation.

TBD

1Password

Other Fineract development related secrets, e. g. for deployments of demo systems on Google Cloud, AWS etc. are managed in a team account at 1Password. At the moment the following committers are members of the 1Password team account:

- [Ed Cable](#)
- [Michael Vorburger](#)
- [Petri Tuomola](#)
- [Arnold Galovics](#)
- [Aleksandar Vidakovic](#)



If you need access or have any questions related to those secrets then please reach out to one of the team members.

GPG

Generate GPG key pairs if you don't already have them and publish them. Please use your Apache email address when creating your GPG keypair. If you already have configured GPG and associated your keypair with a non-Apache email address then please consider creating a separate one just for all

things related to Fineract (or Apache in general).

Instructions:

1. Check your GPG version:

Input GPG version

```
gpg --version
```

Output GPG version

```
gpg (GnuPG) 2.2.27
libgcrypt 1.9.4
Copyright (C) 2021 Free Software Foundation, Inc.
License GNU GPL-3.0-or-later <https://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.
```

Home: /home/aleks/.gnupg

Supported algorithms:

Pubkey: RSA, ELG, DSA, ECDH, ECDSA, EDDSA

Cipher: IDEA, 3DES, CAST5, BLOWFISH, AES, AES192, AES256, TWOFISH,
CAMELLIA128, CAMELLIA192, CAMELLIA256

Hash: SHA1, RIPEMD160, SHA256, SHA384, SHA512, SHA224

Compression: Uncompressed, ZIP, ZLIB, BZIP2



The insecure hash algorithm SHA1 is still supported in version 2.2.27. SHA1 is obsolete and you don't want to use it to generate your signature.

2. Generate your GPG key pair:

Input generate GPG key pair

```
gpg --full-gen-key
```

Output generate GPG key pair (step 1: key type selection)

```
gpg (GnuPG) 2.2.27; Copyright (C) 2021 Free Software Foundation, Inc.  
This is free software: you are free to change and redistribute it.  
There is NO WARRANTY, to the extent permitted by law.
```

Please select what kind of key you want:

- (1) RSA and RSA (default)
- (2) DSA and Elgamal
- (3) DSA (sign only)
- (4) RSA (sign only)
- (14) Existing key from card

Your selection?

There are four options. The default is to use RSA to create the key pair. Good enough for us.

Output generate GPG key pair (step 2: key length selection)

```
RSA keys may be between 1024 and 4096 bits long.  
What keysize do you want? (2048)
```

The default key length is 2048 bits. 1024 is obsolete and a longer 4096 RSA key will not provide more security than 2048 RSA key. Use the default.

Output generate GPG key pair (step 3: validity selection)

```
Requested keysize is 2048 bits  
Please specify how long the key should be valid.  
0 = key does not expire  
<n> = key expires in n days  
<n>w = key expires in n weeks  
<n>m = key expires in n months  
<n>y = key expires in n years  
Key is valid for? (0)2y
```

2 years for the validity of your keys should be fine. You can always update the expiration time later on.

Output generate GPG key pair (step 4: confirmation)

```
Key expires at Sun 16 Apr 2024 08:10:24 PM UTC  
Is this correct? (y/N)y
```

Confirm if everything is correct.

Output generate GPG key pair (step 5: provide user details)

GnuPG needs to construct a user ID to identify your key.
Real name: Aleksandar Vidakovic
Email address: aleks@apache.org
Comment:

Provide your user details for the key. This is important because this information will be included in our key. It's one way of indicating who is owner of this key. The email address is a unique identifier for a person. You can leave Comment blank.

Output generate GPG key pair (step 6: user ID selection)

You selected this USER-ID:
"Aleksandar Vidakovic <aleks@apache.org>"
Change (N)ame, (C)omment, (E)mail or (O)kay/(Q)uit? 0

Select **Okay**.

After the selection of your user ID GPG will ask for a passphrase to protect your private key. Maybe time to open your password manager and generate a secure one and save it in your vault. Once you've confirmed your password GPG will start to generate your keys.



Don't lose your private key password. You won't be able to unlock and use your private key without it.

Output generate GPG key pair (step 7: gpg key pair generation)

We need to generate a lot of random bytes. It is a good idea to perform some other action (type on the keyboard, move the mouse, utilize the disks) during the prime generation; this gives the random number generator a better chance to gain enough entropy.

Generating the GPG keys will take a while.

Output generate GPG key pair (step 8: gpg key pair finished)

```
gpg: key 7890ABCD marked as ultimately trusted ①
gpg: directory '/home/aleks/.gnupg/openpgp-revocs.d' created
gpg: revocation certificate stored as '/home/aleks/.gnupg/openpgp-
revocs.d/ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890ABCD.rev' ②
public and secret key created and signed.

gpg: checking the trustdb
gpg: marginals needed: 3 completes needed: 1 trust model: PGP
gpg: depth: 0 valid: 1 signed: 0 trust: 0-, 0q, 0n, 0m, 0f, 1u
gpg: next trustdb check due at 2024-04-16
pub rsa2048/7890ABCD 2022-04-16 [S] [expires: 2024-04-16] ③
Key fingerprint = ABCD EFGH IJKL MNOP QRST UVWX YZ12 3456 7890 ABCD ④
uid      [ultimate] Aleksandar Vidakovic <aleks@apache.org> ⑤
sub rsa2048/4FGHIJ56 2022-04-16 [] [expires: 2024-04-16]
```

① GPG created a unique identifier in HEX format for your public key. When someone wants to download your public key, they can refer to it either with your email address or this HEX value.

② GPG created a revocation certificate and its directory. You should never share your private key. If your private key is compromised, you need to use your revocation certificate to revoke your key.

③ The public key is 2048 bits using RSA algorithm and shows the expiration date of 16 Apr 2024. The public key ID **7890ABCD** matches the last 8 bits of key fingerprint.

④ The key fingerprint (**ABCD EFGH IJKL MNOP QRST UVWX YZ12 3456 7890 ABCD**) is a hash of your public key.

⑤ Your name and your email address are shown with information about the subkey.

Now you can find that there are two files created under `~/.gnupg/private-keys-v1.d/` directory. These two files are binary files with `.key` extension.

3. Export your public key:

```
gpg --armor --export aleks@apache.org > pubkey.asc
```

4. Export Your Private Key:

```
gpg --export-secret-keys --armor aleks@apache.org > privkey.asc
```

5. Protect Your Private Key and Revocation Certificate

Your private key should be kept in a safe place, like an encrypted flash drive. Treat it like your house key. Only you can have it and don't lose it. And you must remember your passphrase,

otherwise you can't unlock your private key.

You should protect your revocation certificate. Anyone in possession of your revocation certificate, could immediately revoke your public/private key pair and generate fake ones.



Please contact a PMC member to add your GPG public key in Fineract's Subversion repository. This is necessary to be able to validate published releases.

1. Upload your GPG key to a keyserver:

```
gpg --send-keys ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890ABCD
```

Before doing this, make sure that your default keyserver is [hkp://keyserver.ubuntu.com/](http://keyserver.ubuntu.com/). You can do this by changing the default keyserver in `~/.gnupg/dirmngr.conf`:

```
keyserver hkp://keyserver.ubuntu.com/
```

Alternatively you can provide the keyserver with the send command:

```
gpg --keyserver 'hkp://keyserver.ubuntu.com:11371' --send-keys
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890ABCD
```

Another option to publish your key is to submit an armored public key directly at keyserver.ubuntu.com/. You can create the necessary data with this command by providing the email address that you used when you created your key pair:

```
gpg --armor --export aleks@apache.org
```

Output:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----  
  
mQINBF8iGq0BEADGRqeSs0oNDc1sV3L9sQ34KhmoQrACnMYGztx33TD98aWplul+
jm8uGtMmBus4DJJJap1bVQ1oMehw2mscmDHpfJjLNZ/q+vUqbExx1/CER7XvLryN
<--- snip --->
2nHBuBftxDRpDHQ+O5XYwSDSTDMmthPjx0vJGBH4K1k08XK99e01A6/oYLV2SMKp
gXXeWjafxBmHT1cM8hoBZBYzgTu9nK5UnllWunfaHXiCBG4oQQ==
=85/F
-----END PGP PUBLIC KEY BLOCK-----
```

= Email

Official communication related to releases needs to be done with an Apache email address. The Apache Foundation doesn't provide any real email inboxes anymore and just relays emails to your configured private account (GMail etc.).



At the moment we are supporting only GMail accounts. Please let us know if you have other configuration recipes for other email providers.

GMail

You can configure your GMail account and add another profile to use the Apache relay server if you need to send official messages. Please follow these instructions:

TBD.

See also: [Send mail from another address without "on behalf of"](#)

To be able to send emails via SMTP with your GMail account you probably need to create an app password. Please follow these instructions:

1. Go to your Google Account.
2. Select Security.
3. Under "Signing in to Google," select App Passwords. You may need to sign in. If you don't have this option, it might be because:
 4. 2-Step Verification is not set up for your account.
 5. 2-Step Verification is only set up for security keys.
 6. Your account is through work, school, or other organization.
 7. You turned on Advanced Protection.
 8. At the bottom, choose Select app and choose the app you're using and then Select device and choose the device you're using and then Generate.
 9. Follow the instructions to enter the App Password. The App Password is the 16-character code in the yellow bar on your device.
10. Tap Done.

See also: [Google Support: Sign in with App Passwords](#) for more details.

Gradle

TBD

User Properties

There are a couple of properties that contain committer/release manager related secrets. Please add the following properties to your personal global Gradle properties (you will find them at

~/.gradle/gradle.properties in your home folder).

```
fineract.config.gnupg.keyName=ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890ABCD①
fineract.config.gnupg.password=*****
fineract.config.gnupg.publicKeyring=~/.gnupg/pubring.kbx②
fineract.config.gnupg.secretKeyring=~/.gnupg/secreting.gpg
fineract.config.smtp.username=aleks@gmail.com ③
fineract.config.smtp.password=*****
fineract.config.name=Aleksandar Vidakovic
fineract.config.email=aleks@apache.org
fineract.config.username=aleks ④
fineract.config.password=*****
```

① Make sure you use the full GPG key name (you can list yours via `gpg --list-secret-keys --keyid-format=long`)

② GnuPG has its own kbx format to store the public key ring. At the moment we are only supporting this format

③ Currently we only have instructions for GMail

④ Apache committer credentials



Never add any personal secrets in the project gradle.properties. Double check that you are not accidentally committing them to Git!

Release Plugin

Creating Apache Fineract releases was a very manual and tedious procedure before we created the Gradle release plugin. It was easy - even with documentation - to forget a detail. Some ideas are borrowed from the excellent [JReleaser](#) tool. Unfortunately at the moment we can't use it for the full release process. Being an Apache project we have certain requirements that are not fully covered by [JReleaser](#).

Release Plugin Configuration

```
config {
    username = "${findProperty('fineract.config.username')}"
    password = "${findProperty('fineract.config.password')}"
}

doc {
    url = 'git@github.com:apache/fineract-site.git'
    directory = "${System.getProperty("java.io.tmpdir")}/fineract-site"
    branch = "ASF-SITE"
}
git {
    dir = "${projectDir.absolutePath}/.git"
```

```

sections = [
    [
        section: "user",
        name: "name",
        value: "${findProperty('fineract.config.name')}",
    ],
    [
        section: "user",
        name: "email",
        value: "${findProperty('fineract.config.email')}",
    ],
    [
        section: "user",
        name: "signingkey",
        value: "${findProperty('fineract.config.gnupg.keyName')}",
    ],
    [
        section: "commit",
        name: "gpgsign",
        value: "true",
    ],
]
}

template {
    templateDir = "${projectDir}/buildSrc/src/main/resources"
}

gpg {
    keyName = "${findProperty('fineract.config.gnupg.keyName')}"
    publicKeyring = "${findProperty('fineract.config.gnupg.publicKeyring')}"
    secretKeyring = "${findProperty('fineract.config.gnupg.secretKeyring')}"
    password = "${findProperty('fineract.config.gnupg.password')}"
}

smtp {
    host = 'smtp.gmail.com'
    username = "${findProperty('fineract.config.smtp.username')}"
    password = "${findProperty('fineract.config.smtp.password')}"
    tls = true
    ssl = true
}

subversion {
    username = "${findProperty('fineract.config.username')}"
    password = "${findProperty('fineract.config.password')}"
    revision = 'HEAD'
}

jira {
    url = 'https://issues.apache.org/jira/rest/api/2/'
    username = "${findProperty('fineract.config.username')}"
    password = "${findProperty('fineract.config.password')}"
}

```

```
    }
    confluence {
        url = 'https://cwiki.apache.org/confluence/rest/api/'
        username = "${findProperty('fineract.config.username')}"
        password = "${findProperty('fineract.config.password')}"
    }
}
```

Release Process

TODO:

- create "Jira anchor ticket" with all issues linked that are going into this release.
- maintenance: continuously update the "Jira anchor ticket" to make sure we catch all ticket changes
- maintenance: list tickets that have discrepancies, e. g. tickets still open while associated PR merged, ticket on wrong version (i. e. associated PR already merged before with another release).

TBD



Consider the Gradle plugin commands an experimental feature!

Fineract Release Process

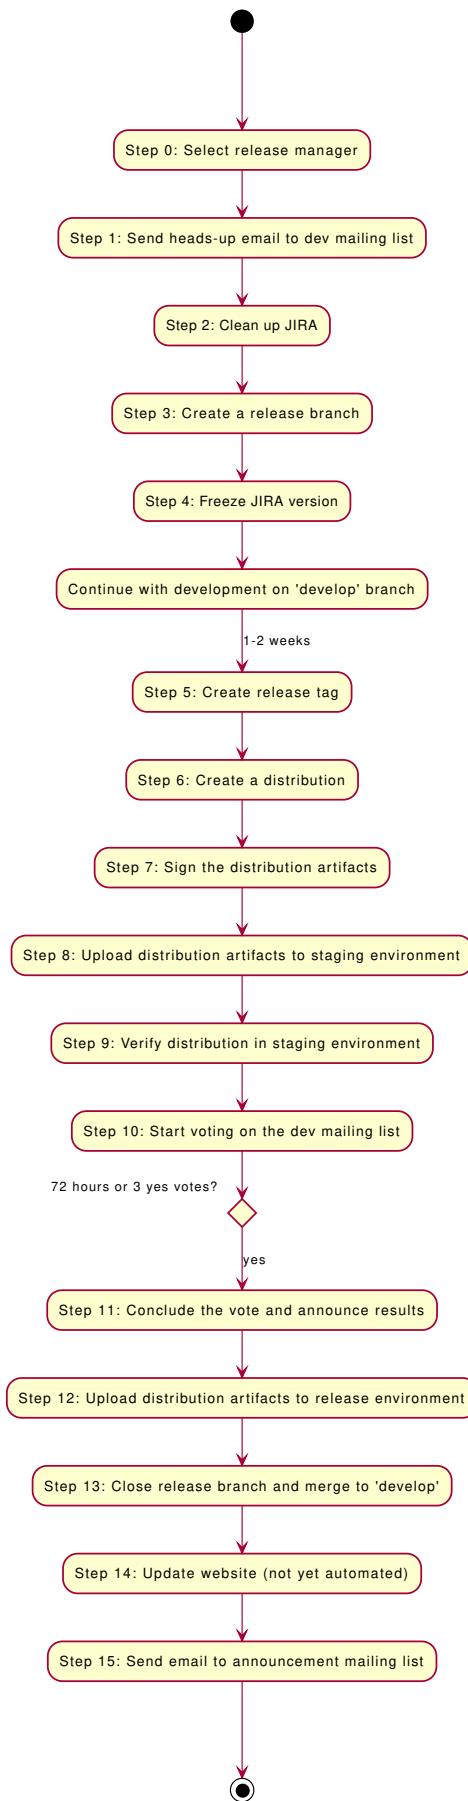


Figure 5. Release Process Diagram

Step 1: Heads-Up Email

Description

The RM should, if one doesn't already exist, first create a new release umbrella issue in JIRA. This issue is dedicated to tracking (a summary of) any discussion related to the planned new release. An example of such an issue is FINERACT-873 - Release Apache Fineract v1.4.0 RESOLVED.

The RM then creates a list of resolved issues & features through an initial check in JIRA for already resolved issues for the release, and then setup a timeline for release branch point. The time for the day the issue list is created to the release branch point must be at least two weeks in order to give the community a chance to prioritize and commit any last minute features and issues they would like to see in the upcoming release.

The RM must then send the pointer to the umbrella issue along with the tentative timeline for branch point to the developer lists. Any work identified as release related that needs to be completed should be added as a sub tasks of the umbrella issue to allow all developers and users to see the overall release progress in one place. The umbrella issue shall also link to any issues still requiring clarification whether or not they will make it into the release.

The RM should then inform users when the git branch is planned to be created, by sending an email based on this template:

[FINERACT] [PROPOSAL] □ New release \${project['fineract.release.version']}

Hello everyone,

... based on our "How to Release Apache Fineract" process documented at <https://cwiki.apache.org/confluence/x/DRwIB>:

I will create a \${project['fineract.release.version']} branch off develop in our git repository at <https://github.com/apache/fineract> on \${project['fineract.release.date']}.

The release tracking umbrella issue for tracking all activity in JIRA is FINERACT-
\${project['fineract.release.issue']}!0000
(<https://issues.apache.org/jira/browse/FINERACT-0000>) for this Fineract
\${project['fineract.release.version']}.

If you have any work in progress that you would like to see included in this release, please add "blocking" links to the release JIRA issue.

I am the release manager for this release.

Cheers,

\${project['fineract.config.name']}

□ Powered by Fineract Release Plugin □

Gradle Task

Command

```
% ./gradlew fineractReleaseStep1 -Pfineract.release.issue=1234  
-Pfineract.release.date="Monday, April 25, 2022" -Pfineract.release.version=1.11.0
```

Step 2: Clean Up JIRA

Description

Before a release is done, make sure that any issues that are fixed have their fix version setup correctly.

```
project = FINERACT and resolution = fixed and fixVersion is empty
```

Move all unresolved JIRA issues which have this release as Fix Version to the next release

```
project = FINERACT and fixVersion = 1.11.0 and status not in ( Resolved, Done, Accepted, Closed )
```

You can also run the following query to make sure that the issues fixed for the to-be-released version look accurate:

```
project = FINERACT and fixVersion = 1.11.0
```

Finally, check out the output of the JIRA release note tool to see which tickets are included in the release, in order to do a sanity check.

Gradle Task

Command

```
% ./gradlew fineractReleaseStep2 -Pfineract.release.version=1.11.0
```



This task is not yet automated!

Step 3: Create Release Branch

Description

Communicate with the community. You do not need to start a new email thread on the developer mailing list to notify that you are about to branch, just do it ca. 2 weeks after the initial email, or later, based on the discussion on the initial email.

You do not need to ask committers to hold off any commits until you have branched finished, as it's always possible to fast-forward the branch to latest develop, or cherry-pick last minute changes to it. People should be able to continue working on the develop branch on bug fixes and great new features for the next release while the release process for the current release is being worked through.

1. Clone fresh repository copy

```
% git clone git@github.com:apache/fineract.git
% cd fineract
```

2. Check that current HEAD points to commit on which you want to base new release branch. Checkout a particular earlier commit if not.

```
% git log ①
```

① Check current branch history. HEAD should point to commit that you want to be base for your release branch

3. Create a new release branch with name "\$Version"

```
% git checkout -b 1.11.0
```

4. Push new branch to Apache Fineract repository

```
% git push origin 1.11.0
```

5. Add new release notes in Release Folders. The change list can be swiped from the JIRA release note tool (use the "text" format for the change log). See JIRA Cleanup above to ensure that the release notes generated by this tool are what you are expecting.

6. Send an email announcing the new release branch on the earlier email thread

[FINERACT] [ANNOUNCE] ② \${project['fineract.release.version']} release branch

Hello everyone,

... as previously announced, I've just created the release branch for our upcoming \${project['fineract.release.version']} release.

You can continue working and merging PRs to the develop branch for future releases, as always.

The DRAFT release notes are on

[https://cwiki.apache.org/confluence/display/FINERACT/\\${project\['fineract.release.version'\]}+-+Apache+Fineract](https://cwiki.apache.org/confluence/display/FINERACT/${project['fineract.release.version']}+-+Apache+Fineract). Does anyone see anything missing?

Does anyone have any last minutes changes they would like to see cherry-picked to branch \${project['fineract.release.version']}, or are we good go and actually cut the release based on this branch as it is?

I'll initiate the final stage of actually creating the release on \${project['fineract.release.date']} if nobody objects.

Cheers,

`\${project['fineract.config.name']}

Gradle Task

Command

```
% ./gradlew fineractReleaseStep3 -Pfineract.release.date="Monday, May 10, 2022"  
-Pfineract.release.version=1.11.0
```

Step 4: Freeze JIRA

Description

You first need to close the release in JIRA so that the about to be released version cannot be used as "fixVersion" for new bugs anymore. Go to JIRA "Administer project" page and follow "Versions" in left menu. Table with list of all releases should appear, click on additional menu on the right of your release and choose "Release" option. Submit release date and you're done.

Gradle Task

Command

```
% ./gradlew fineractReleaseStep4
```



This task is not yet automated!

Step 5: Create Release Tag

Description

Next, you create a git tag from the HEAD of the release's git branch.

```
% git checkout 1.11.0  
% ./gradlew clean integrationTests ①  
% git tag -a 1.11.0 -m "Fineract 1.11.0 release"  
% git push origin 1.11.0
```

① Run additonally manual tests with the community app.



It is important to create so called annotated tags (vs. lightweight) for releases.

Gradle Task

Command

```
% ./gradlew fineractReleaseStep5 -Pfineract.release.version=1.11.0
```

Step 6: Create Distribution

Description

Create source and binary artifacts. Make sure to do some sanity checks. The tar and the release branch should match.

```
% cd /fineract-release-preparations ①
% tar -xvf apache-fineract-1.11.0-src.tar.gz
% git clone <a href="https://git-wip-us.apache.org/repos/asf/fineract.git"
class="bare">git-wip-us.apache.org/repos/asf/fineract.git</a>
% cd fineract/
% git checkout tags/1.11.0
% cd ..
% diff -r fineract apache-fineract-1.11.0-src
```

① Do a fresh clone of the tag.

Make sure code compiles and tests pass on the uncompressed source.

```
% cd apache-fineract-1.11.0-src/fineract-provider ①
% gradlew clean integrationTest ②
% gradlew clean build ③
% gradlew rat ④
```

① Make sure prerequisites are met before running these commands.

② For running integration tests

③ For building deploy able war

④ For RAT checks

Gradle Task

Command

```
% ./gradlew fineractReleaseStep
```

Step 7: Sign Distribution

Description

All release artifacts must be signed. In order to sign a release you will need a PGP key. You should get your key signed by a few other people. You will also need to receive their keys from a public key server. See the Apache release signing page for more details. Please follow the steps defined in Release Sign.

```
% gpg --armor --output apache-fineract-1.11.0-src.tar.gz.asc --detach-sig apache-fineract-1.11.0-src.tar.gz
% gpg --print-md MD5 apache-fineract-1.11.0-src.tar.gz > apache-fineract-1.11.0-src.tar.gz.md5
% gpg --print-md SHA512 apache-fineract-1.11.0-src.tar.gz > apache-fineract-1.11.0-src.tar.gz.sha512
% gpg --armor --output apache-fineract-1.11.0--binary.tar.gz.asc --detach-sig apache-fineract-1.11.0-binary.tar.gz
% gpg --print-md MD5 apache-fineract-1.11.0-binary.tar.gz > apache-fineract-1.11.0-binary.tar.gz.md5
% gpg --print-md SHA512 apache-fineract-1.11.0-binary.tar.gz > apache-fineract-1.11.0-binary.tar.gz.sha512
```

Gradle Task

Command

```
% ./gradlew fineractReleaseStep7
```

Step 8: Upload Distribution Staging

Description

Finally create a directory with release name (1.11.0 in this example) in dist.apache.org/repos/dist/dev/fineract and add the following files in this new directory:

- apache-fineract-1.11.0-binary.tar.gz.sha
- apache-fineract-1.11.0-binary.tar.gz
- apache-fineract-1.11.0-binary.tar.gz.asc
- apache-fineract-1.11.0-binary.tar.gz.md5
- apache-fineract-1.11.0-src.tar.gz.sha
- apache-fineract-1.11.0-src.tar.gz
- apache-fineract-1.11.0-src.tar.gz.asc
- apache-fineract-1.11.0-src.tar.gz.md5

Upload binary and source archives to ASF's distribution dev (staging) area:

```
% svn co <a href="https://dist.apache.org/repos/dist/dev/fineract/">dist.apache.org/repos/dist/dev/fineract/</a> fineract-dist-dev
% mkdir fineract-dist-dev/1.11.0
% cp fineract/build/distributions/* fineract-dist-dev/1.11.0/
% svn commit
```



You will need your ASF Committer credentials to be able to access the Subversion host dist.apache.org via.

Gradle Task

Command

```
% ./gradlew fineractReleaseStep8 -Pfineract.release.version=1.11.0
```

Step 9: Verify Distribution Staging

Description

Following are the typical things we need to verify before voting on a release candidate. And the release manager should verify them too before calling out a vote.

Make sure release artifacts are hosted at dist.apache.org/repos/dist/dev/fineract

- Release candidates should be in format apache-fineract-1.11.0-binary.tar.gz
- Verify signatures and hashes. You may have to import the public key of the release manager to verify the signatures. ([gpg --recv-key <key id>](#))
- Git tag matches the released bits (diff -rf)
- Can compile successfully from source
- Verify DISCLAIMER, NOTICE and LICENSE (year etc)
- All files have correct headers (Rat check should be clean - gradlew rat)
- No jar files in the source artifacts
- Integration tests should work

Gradle Task

Command

```
% ./gradlew fineractReleaseStep9 -Pfineract.release.version=1.11.0
```



This task is not yet automated!

Step 10: Start Vote

Description

Voting has to be done on dev@fineract.apache.org. You can close the vote after voting period expires (72 hours) and you accumulate sufficient votes (minimum 3 x +1 PMC votes).

```
[FINERACT] [VOTE] ${project['fineract.release.version']} for release
```

Hello everyone,

... we have created Apache Fineract \${project['fineract.release.version']} release, with the artifacts below up for a vote.

It fixes the following issues:

[https://cwiki.apache.org/confluence/display/FINERACT/\\${project\['fineract.release.version'\]}+-+Apache+Fineract](https://cwiki.apache.org/confluence/display/FINERACT/${project['fineract.release.version']}+-+Apache+Fineract)

Source & Binary files :

[https://dist.apache.org/repos/dist/dev/fineract/\\${project\['fineract.release.version'\]}](https://dist.apache.org/repos/dist/dev/fineract/${project['fineract.release.version']})

Tag to be voted on (rc#):

[https://gitbox.apache.org/repos/asf?p=fineract.git;a=commit;h=refs/heads/\\${project\['fineract.release.version'\]}](https://gitbox.apache.org/repos/asf?p=fineract.git;a=commit;h=refs/heads/${project['fineract.release.version']})

Fineract's KEYS containing the PGP key we used to sign the release:

<https://dist.apache.org/repos/dist/dev/fineract/KEYS>

Note that this release contains source and binary artifacts.

This vote will be open for 72 hours:

[] +1 approve
[] +0 no opinion
[] -1 disapprove (and reason why)

Cheers,

`${project['fineract.config.name']}`

Gradle Task

Command

```
% ./gradlew fineractReleaseStep10 -Pfineract.release.version=1.11.0
```

Step 11: Finish Vote

Description

Upon receiving 3 x +1 from the PMC, or after 72 hours (whichever one comes first), reply to the voting thread and add the prefix "[RESULT]" to the subject line with the results, as follows:

```
[FINERACT] [VOTE] [RESULT] :: ${project['fineract.release.version']} for release
```

```
<if (project['fineract.vote'].approve.binding?size +  
project['fineract.vote'].approve.nonBinding?size >  
project['fineract.vote'].disapprove.binding?size +  
project['fineract.vote'].disapprove.nonBinding?size)>  
Voting is now closed and has passed with the following tally,
```

```
Binding +1s: ${project['fineract.vote'].approve.binding?size}  
Non binding +1s: ${project['fineract.vote'].approve.nonBinding?size}  
<else>  
Voting is now closed and has not passed with the following tally,
```

```
Binding +1s: ${project['fineract.vote'].approve.binding?size}  
Non binding +1s: ${project['fineract.vote'].approve.nonBinding?size}
```

```
Binding -1s: ${project['fineract.vote'].disapprove.binding?size}  
Non binding -1s: ${project['fineract.vote'].disapprove.nonBinding?size}  
</if>
```

Here are the detailed results:

```
<list project['fineract.vote'].approve.binding>  
Binding +1s:  
  <items as item>  
  - ${item.name} (${item.email})  
  </items>  
</list>
```

```
<list project['fineract.vote'].approve.nonBinding>  
Non binding +1s:  
  <items as item>  
  - ${item.name} (${item.email})  
  </items>
```

```

</#list>

<#list project['fineract.vote'].disapprove.binding>
Binding -1s:
  <#items as item>
- ${item.name} (${item.email})
  </#items>
</#list>

<#list project['fineract.vote'].disapprove.nonBinding>
Non binding -1s:
  <#items as item>
- ${item.name} (${item.email})
  </#items>
</#list>

<#list project['fineract.vote'].noOpinion.binding>
Binding +0s:
  <#items as item>
- ${item.name} (${item.email})
  </#items>
</#list>

<#list project['fineract.vote'].noOpinion.nonBinding>
Non binding +0s:
  <#items as item>
- ${item.name} (${item.email})
  </#items>
</#list>

<if (project['fineract.vote'].approve.binding?size +
project['fineract.vote'].approve.nonBinding?size >
project['fineract.vote'].disapprove.binding?size +
project['fineract.vote'].disapprove.nonBinding?size)>
Thanks to everyone who voted! I'll now continue with the rest of the release process.
<else>
Thanks to everyone who voted! Looks like we have to repeat the vote.
</if>

${project['fineract.config.name']}

```

Gradle Task

Command

```
% ./gradlew fineractReleaseStep11 -Pfineract.release.version=1.11.0
```

Step 12: Upload Distribution Release

Description

In order to release you have to checkout release repository located on dist.apache.org/repos/dist/release/fineract and add release artifacts there.

```
% svn co <a href="https://dist.apache.org/repos/dist/release/fineract" class="bare">dist.apache.org/repos/dist/release/fineract</a> fineract-release
% mkdir fineract-release/1.11.0/
% cp fineract-dist-dev/1.11.0/* fineract-release/1.11.0/
% svn add fineract-release/1.11.0/
% svn commit -m "Fineract Release 1.11.0" fineract-release/1.11.0/
```

You will now get an automated email from the Apache Reporter Service (no-reply@reporter.apache.org), subject "Please add your release data for 'fineract'" to add the release data (version and date) to the database on reporter.apache.org/addrelease.html?fineract (requires PMC membership).

Gradle Task

Command

```
% ./gradlew fineractReleaseStep12 -Pfineract.release.version=1.11.0
```

Step 13: Close Release Branch

Description

As discussed in [FINERACT-1154](#), now that everything is final, please do the following to remove the release branch (and just keep the tag), and make sure that everything on the release tag is merged to develop and that e.g. git describe works:

```
% git checkout develop
% git branch -D 1.11.0
% git push origin :1.11.0
% git checkout develop
% git checkout -b merge-1.11.0
% git merge -s recursive -Xignore-all-space 1.11.0 ①
% git commit
% git push $USER
% hub pull-request
```

① Manually resolve merge conflicts, if any

Gradle Task

Command

```
% ./gradlew fineractReleaseStep13 -Pfineract.release.version=1.11.0
```



This task is not yet automated!

Step 14: Update website

Description

Finally update the fineract.apache.org website with the latest release details. The website's HTML source code is available at github.com/apache/fineract-site.



This step is not yet updated. We are working on a static site generator setup.

Gradle Task

Command

```
% ./gradlew fineractReleaseStep14 ①
```

① Currently doing nothing. Will trigger in the future the static site generator and publish on Github.



This task is not yet automated!

Step 15: Announcement Email

Description

Send an email to announce@apache.org (sender address must be @apache.org):

[ANNOUNCE] Apache Fineract \${project['fineract.release.version']} Release

The Apache Fineract project is pleased to announce the release of Apache Fineract \${project['fineract.release.version']}. The release is available for download from <https://fineract.apache.org/#downloads>

Fineract provides a reliable, robust, and affordable solution for entrepreneurs, financial institutions, and service providers to offer financial services to the world's 2 billion underbanked and unbanked. Fineract is aimed at innovative mobile and cloud-based solutions, and enables digital transaction accounts for all.

This release addressed \${project['fineract.release.issues']} issues.

Readme:

[https://github.com/apache/fineract/blob/\\${project\['fineract.release.version'\]}/README.md](https://github.com/apache/fineract/blob/${project['fineract.release.version']}/README.md)

Release page:

[https://cwiki.apache.org/confluence/display/FINERACT/\\${project\['fineract.release.version'\]}+-+Apache+Fineract](https://cwiki.apache.org/confluence/display/FINERACT/${project['fineract.release.version']}+-+Apache+Fineract)

List of fixed issues:

[https://issues.apache.org/jira/secure/ReleaseNote.jspa?version=\\${project\['fineract.release.versionId'\]}&styleName=Html&projectId=\\${project\['fineract.release.projectId'\]}&showAll=true](https://issues.apache.org/jira/secure/ReleaseNote.jspa?version=${project['fineract.release.versionId']}&styleName=Html&projectId=${project['fineract.release.projectId']}&showAll=true)

For more information on Apache Fineract please visit project home page: <https://fineract.apache.org>

The Apache Fineract Team

Gradle Task

Command

```
% ./gradlew fineractReleaseStep15 -Pfineract.release.version=1.11.0
```

Maintenance Release Process



This is a first attempt to introduce maintenance releases. Some details might change as soon as we get more experience with the process and feedback from the community. The numbers here are still more or less arbitrary, and we'll adapt as necessary.

Rules

- hotfix releases are reserved for critical (BLOCKER) bugs and security issues. Probably we'll have some kind of voting process in place, e. g. "minimum 3 x +1 votes from PMC members"
- we will support (for now to start) two minor versions back counting from the last release; this would mean that once 1.8.0 is out we would support 1.8.x and 1.7.x, but not 1.6.x and older; this rule is tentative, we'll see then what we do in the future when we have more feedback.
- guaranteed backward compatibility with the last minor release; i. e. "1.6.1" is a drop-in replacement for "1.6.0"
- NO new features, tables, data, REST endpoints
- NO major (or "minor" framework upgrades); i. e. if we used Spring Boot "2.6.1" in version "1.6.0" of Fineract we can upgrade dependencies to "2.6.10" (unless it breaks something of course), but not to "2.7.2" of Spring Boot



The rest of the release process is the same as for normal releases. In the future we might have smaller time windows for reviews.

Publish Release Artifacts



More on releases at the ASF see www.apache.org/legal/release-policy.html#distribute-raw-artifact

Requirements

You need to have your GPG keypairs properly set up. The JAR release artifacts (currently only `fineract-client`) are signed with a Gradle plugin just before being uploaded to the Maven repository. Please make sure that the following properties are set in your private `gradle.properties` file in your home folder:

```
signing.keyId=7890ABCD
signing.password=*****
signing.secretKeyRingFile=~/gnupg/secring.gpg
```

This is quite similar to the Fineract release plugin properties for GPG. In one of the next release we'll merge these two setups to avoid this duplicated configuration.

Maven Repository

We are using the ASF's official [Nexus Maven repository](http://infra.apache.org/publishing-maven-artifacts.html) to publish our snapshot and release artifacts.



Find more information at infra.apache.org/publishing-maven-artifacts.html

NPM Registry

For convenience we will be using Github Packages to publish Fineract's Typescript API client.

TBD

Docker Hub

TBD

Fineract SDKs

TBD

Generate Apache Fineract API Client

Apache Fineract supports client code generation using [OpenAPI Generator](#). It uses [OpenAPI Specification Version 3.0.3](#).

Fineract SDK Java API Client

The `fineract-client.jar` will eventually be available on Maven Central (watch [FINERACT-1102](#)). Until it is, you can quite easily build the latest and greatest version locally from source, see below.

The `FineractClient` is the entry point to the *Fineract SDK Java API Client*. `Calls` is a convenient and recommended utility to simplify the use of the `retrofit2.Call` type which all API operations return. This permits you to use the API like the `FineractClientDemo` illustrates:

```
import org.apache.fineract.client.util.FineractClient;
import static org.apache.fineract.client.util.Calls.ok;

    FineractClient fineract = FineractClient.builder().baseURL(
"https://demo.fineract.dev/fineract-provider/api/v1/").tenant("default")
    .basicAuth("mifos", "password").build();
    List<RetrieveOneResponse> staff = Calls.ok(fineract.staff.retrieveAll16(1L, true,
false, "ACTIVE"));
    String name = staff.get(0).getDisplayName();
    log.info("Display name: {}", name);
```

Generate API Client

The API client is built as part of the standard overall Fineract Gradle build. The client JAR can be found in `fineract-client/build/libs` as `fineract-client.jar`.

If you need to save time to incrementally work on making small changes to Swagger annotations in an IDE, you can execute e.g. the following line in root directory of the project to exclude non-require Gradle tasks:

```
./gradlew -x compileJava -x compileTest -x spotlessJava -x enhance resolve
prepareInputYaml :fineract-client:buildJavaSdk
```

Validate OpenAPI Spec File

The `resolve` task in `build.gradle` file will generate the OpenAPI Spec File for the project. To make sure Swagger Codegen generates a correct library, it is important for the OpenAPI Spec file to be valid. Validation is done automatically by the OpenAPI code generator Gradle plugin. If you still have problems during code generation please use [Swagger OpenAPI Validator](#) to validate the spec file.

Frequently Asked Questions

See [Fineract FAQ wiki page](#).

Glossary

TBD

Index

Appendix A: Fineract Application Properties

TBD

Tenant Database Properties

Table 3. Tenant Database Properties

Name	Env Variable	Default Value	Description
fineract.tenant.host	FINERACT_DEFAULT_TE NANTDB_HOSTNAME	localhost	This property sets the hostname of the default tenant database.
fineract.tenant.port	FINERACT_DEFAULT_TE NANTDB_PORT	3306	This property sets the port of the default tenant database.
fineract.tenant.username	FINERACT_DEFAULT_TE NANTDB_UID	root	This property sets the username of the default tenant database.
fineract.tenant.password	FINERACT_DEFAULT_TE NANTDB_PWD	mysql	This property sets the password of the default tenant database.
fineract.tenant.parameters	FINERACT_DEFAULT_TE NANTDB_CONN_PARAMS		This property sets the connection parameters of the default tenant database. eg. whether ssl is enabled or not
fineract.tenant.timezone	FINERACT_DEFAULT_TE NANTDB_TIMEZONE	Asia/Kolkata	This property sets the timezone of the default tenant
fineract.tenant.identifier	FINERACT_DEFAULT_TE NANTDB_IDENTIFIER	default	This property sets the unique identifier for the tenant within fineract
fineract.tenant.name	FINERACT_DEFAULT_TE NANTDB_NAME	fineract_default	This property sets the database name of the default tenant
fineract.tenant.description	FINERACT_DEFAULT_TE NANTDB_DESCRIPTION	Default Demo Tenant	This property sets the description of the default tenant

Name	Env Variable	Default Value	Description
fineract.tenant.master-password	FINERACT_DEFAULT_TENANTDB_MASTER_PASSWORD	fineract	The password used to encrypt sensitive tenant data within the database
fineract.tenant.encryption	FINERACT_DEFAULT_TENANTDB_ENCRYPTION	AES/CBC/PKCS5Padding	This property sets the symmetric encryption algorithm used to encrypt sensitive tenant data within the database e.g tenant database password
spring.liquibase.enabled	FINERACT_LIQUIBASE_ENABLED	true	If set to true, liquibase will be enabled and the instance running this configuration will run migrations
fineract.tenant.read-only-name	FINERACT_DEFAULT_TENANTDB_RO_NAME		For read only configuration, set this to the name of the read only tenant database
fineract.tenant.read-only-host	FINERACT_DEFAULT_TENANTDB_RO_HOSTNAME		For read only configuration, set this to the hostname of the read only tenant database
fineract.tenant.read-only-port	FINERACT_DEFAULT_TENANTDB_RO_PORT		For read only configuration, set this to the port of the read only tenant database
fineract.tenant.read-only-username	FINERACT_DEFAULT_TENANTDB_RO_UID		For read only configuration, set this to the username of the read only tenant database
fineract.tenant.read-only-password	FINERACT_DEFAULT_TENANTDB_RO_PWD		For read only configuration, set this to the password of the read only tenant database

Name	Env Variable	Default Value	Description
fineract.tenant.read-only-parameters	FINERACT_DEFAULT_TENANTDB_RO_CONN_PARAMETERS		For read only configuration, set this to the connection parameters of the read only tenant database

Hikari Connection Pool Properties

Table 4. Hikari Connection Pool Properties

Name	Env Variable	Default Value	Description
spring.datasource.hikari.driverClassName	FINERACT_HIKARI_DRIVER_SOURCE_CLASS_NAME	org.mariadb.jdbc.Driver	The correct driver name for the database that will be used with fineract.
spring.datasource.hikari.jdbcUrl	FINERACT_HIKARI_JDBC_URL	jdbc:mariadb://localhost:3306/fineract_tenants	The database connection string for the database with tenant information that will be used with fineract.
spring.datasource.hikari.username	FINERACT_HIKARI_USERNAME	root	The username for the database with tenant information that will be used with fineract
spring.datasource.hikari.password	FINERACT_HIKARI_PASSWORD	mysql	The password for the database with tenant information that will be used with fineract
spring.datasource.hikari.minimumIdle	FINERACT_HIKARI_MINIMUM_IDLE	3	The minimum number of connections in hikari pool that will be maintained when the system is idle
spring.datasource.hikari.maximumPoolSize	FINERACT_HIKARI_MAXIMUM_POOL_SIZE	10	The maximum number of connections that hikari can create in the pool.

Name	Env Variable	Default Value	Description
spring.datasource.hikari.idleTimeout	FINERACT_HIKARI_IDLE_TIMEOUT	60000	The maximum time in milliseconds that a connection is allowed to sit idle in the pool.
spring.datasource.hikari.connectionTimeout	FINERACT_HIKARI_CONNECTION_TIMEOUT	20000	The maximum time in milliseconds that hikari will wait for a connection to be established.
spring.datasource.hikari.connectionTestquery	FINERACT_HIKARI_TEST_QUERY	SELECT 1	The query that will be used to test the database connection.
spring.datasource.hikari.autoCommit	FINERACT_HIKARI_AUTO_COMMIT	true	If set to true, the connections in the pool will be in auto-commit mode.
spring.datasource.hikari.dataSourceProperties['cachePreStmts']	FINERACT_HIKARI_DS_PROPERTIES_CACHE_PREP_STMTS	true	If set to true, hikari caches compiled SQL statements to avoid the overhead of re-parsing and re-compiling SQL queries.
spring.datasource.hikari.dataSourceProperties['prepStmtCacheSize']	FINERACT_HIKARI_DS_PROPERTIES_PREP_STMT_CACHE_SIZE	250	The maximum number of prepared statements that hikari can cache.
spring.datasource.hikari.dataSourceProperties['prepStmtCacheSqlLimit']	FINERACT_HIKARI_DS_PROPERTIES_PREP_STMT_CACHE_SQL_LIMIT	2048	This property sets the upper limit for the size of individual SQL queries that can be stored in the cache. If a SQL query exceeds this limit in terms of character length, it will not be cached, even if caching is enabled.

Name	Env Variable	Default Value	Description
spring.datasource.hikari.dataSourceProperties['useServerPrepStmts']	FINERACT_HIKARI_DS_PROPERTIES_USE_SERVER_PREP_STMTS	true	This property determines if the connection should leverage server-side prepared statements rather than client-side ones.
spring.datasource.hikari.dataSourceProperties['useLocalSessionState']	FINERACT_HIKARI_DS_PROPERTIES_USE_LOCAL_SESSION_STATE	true	This property allows the connection pool to locally track changes to session-specific properties (like character sets or time zones) rather than sending these queries to the database repeatedly.
spring.datasource.hikari.dataSourceProperties['rewriteBatchedStatements']	FINERACT_HIKARI_DS_PROPERTIES_REWRITE_BATCHED_STATEMENTS	true	This property, when set to true, allows the JDBC driver to rewrite batched SQL statements into a more efficient single query format before sending them to the database.
spring.datasource.hikari.dataSourceProperties['cacheResultSetMetadata']	FINERACT_HIKARI_DS_PROPERTIES_CACHE_RESULT_SET_METADATA	true	This property, when set to true, enables the caching of metadata for ResultSet objects. This metadata includes details such as column names, types, and other relevant schema information.

Name	Env Variable	Default Value	Description
spring.datasource.hikari.dataSourceProperties['cacheServerConfiguration']	FINERACT_HIKARI_DS_PROPERTIES_CACHE_SERVER_CONFIGURATION	true	When set to true, this property allows the JDBC driver to cache the server configuration settings, which include properties such as session state, character sets, and other configuration details relevant to the database server.
spring.datasource.hikari.dataSourceProperties['elideSetAutoCommits']	FINERACT_HIKARI_DS_PROPERTIES_ELIDE_SET_AUTO_COMMITS	true	When set to true, this property prevents the JDBC driver from issuing a SET autocommit command on the database connection during its initialization.
spring.datasource.hikari.dataSourceProperties['maintainTimeStats']	FINERACT_HIKARI_DS_PROPERTIES_MAINTAIN_TIME_STATS	false	When set to true, this property enables HikariCP to track and maintain statistics regarding various timing metrics related to connection pool operations, such as connection acquisition times.
spring.datasource.hikari.dataSourceProperties['logSlowQueries']	FINERACT_HIKARI_DS_PROPERTIES_LOG_SLOW_QUERIES	true	When set to true, this property enables HikariCP to log SQL queries that exceed a specified execution time threshold, allowing developers and administrators to identify and analyze performance issues related to slow-running queries.

Name	Env Variable	Default Value	Description
spring.datasource.hikari.dataSourceProperties['dumpQueriesOnException']	FINERACT_HIKARI_DS_PROPERTIES_DUMP_QUESTIONS_IN_EXCEPTION	true	When set to true, this property instructs HikariCP to log the SQL statements that caused exceptions during execution. This includes capturing the query text and any associated parameters.

SSL Properties

Table 5. SSL Properties

Name	Env Variable	Default Value	Description
server.ssl.enabled	FINERACT_SERVER_SSL_ENABLED	true	When set to true, SSL (Secure Sockets Layer) or TLS (Transport Layer Security) will be enabled for the server.
server.ssl.protocol	FINERACT_SERVER_SSL_PROTOCOL	TLS	This property allows you to define specific SSL/TLS protocol version the server will use when establishing secure connections. Common protocols include TLSv1.2, TLSv1.3, etc.
server.ssl.ciphers	FINERACT_SERVER_SSL_CIPHERS	TLS_RSA_WITH_AES_128_CBC_SHA256	This property allows you to control the cipher suites that fineract will accept for secure connections
server.ssl.enabled-protocols	FINERACT_SERVER_SSL_PROTOCOLS	TLSv1.2	This property allows you to define a list of SSL/TLS protocol versions that the server will support when establishing secure connections

Name	Env Variable	Default Value	Description
server.ssl.key-store	FINERACT_SERVER_SSL_KEY_STORE	classpath:keystore.jks	The property is used to specify the location of the SSL key store file that contains the server's private key and the associated certificate
server.ssl.key-store-password	FINERACT_SERVER_SSL_KEY_STORE_PASSWORD	openmf	The property defines the password for the keystore specified under property server.ssl.key-store

Authentication Properties

Table 6. Authentication Properties

Name	Env Variable	Default Value	Description
fineract.security.basicauth.enabled	FINERACT_SECURITY_BASICAUTH_ENABLED	true	When set to true, the supported authentication method will be basic authentication.
fineract.security.oauth.enabled	FINERACT_SECURITY_OAUTH_ENABLED	false	When set to true, the supported authentication method will be OAuth.
fineract.security.2fa.enabled	FINERACT_SECURITY_2FA_ENABLED	false	Set the value to true enable two-factor authentication. For this to work as expected, ensure that you have set the correct email/sms configuration
spring.security.oauth2.resourceserver.jwt.issuer-uri	FINERACT_SERVER_OAUTH_RESOURCE_URL	localhost:9000/auth/realm/fineract	If OAuth is enabled and a custom resource server (different from what is provided) is required, set the issuer-uri here.

Tomcat Properties

Table 7. Tomcat Properties

Name	Env Variable	Default Value	Description
server.tomcat.accept-count	FINERACT_SERVER_TO_MCAT_ACCEPT_COUNT	100	The property specifies the maximum number of concurrent connection requests that embedded Tomcat can queue. If this limit is reached, incoming connection requests will be rejected.
server.tomcat.accesslog-enabled	FINERACT_SERVER_TO_MCAT_ACCESSLOG_ENABLED	false	If set to true, tomcat will log access requests to file
server.tomcat.max-connections	FINERACT_SERVER_TO_MCAT_MAX_CONNECTIONS	8192	Sets the maximum number of simultaneous connections Tomcat can handle.
server.tomcat.max-http-form-post-size	FINERACT_SERVER_TO_MCAT_MAX_HTTP_FORM_POST_SIZE	2MB	The property in sets the maximum size of HTTP POST requests that Tomcat can handle
server.tomcat.max-keep-alive-requests	FINERACT_SERVER_TO_MCAT_MAX_KEEP_ALIVE_REQUESTS	100	The property specifies the maximum number of HTTP requests that can be sent over a single persistent connection (HTTP Keep-Alive) before Tomcat closes the connection
server.tomcat.threads.max	FINERACT_SERVER_TO_MCAT_THREADS_MAX	200	The property sets the maximum number of threads that Tomcat can use to process requests
server.tomcat.threads.min-spare	FINERACT_SERVER_TO_MCAT_THREADS_MIN_SPARE	10	The property specifies the minimum number of spare (idle) threads that Tomcat should maintain

Kafka Properties

Table 8. Kafka related properties for Remote Spring Batch Jobs

Name	Env Variable	Default Value	Description
fineract.remote-job-message-handler.kafka.enabled	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_ENABLED	false	Enables or disables Kafka for remote job execution. If Kafka is enabled then JMS shall be disabled.
fineract.remote-job-message-handler.kafka.topic.auto-create	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_TOPIC_AUTO_CREATE	true	Enables topic auto creation. In case the auto creation of the topic is disabled please make sure that the replica and the partition count is properly configured.
fineract.remote-job-message-handler.kafka.topic.name	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_TOPIC_NAME	job-topic	Name of the topic where partitioned tasks are sent to
fineract.remote-job-message-handler.kafka.topic.replicas	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_TOPIC_REPLICAS	1	Number of the replicas
fineract.remote-job-message-handler.kafka.topic.partitions	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_TOPIC_PARTITIONS	10	Number of partitions
fineract.remote-job-message-handler.kafka.bootstrap-servers	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_BOOTSTRAP_SERVERS	localhost:9092	Comma separated list of bootstrap servers
fineract.remote-job-message-handler.kafka.consumer.group-id	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_CONSUMER_GROUPID	fineract-consumer-group-id	Group ID of the Consumer

Name	Env Variable	Default Value	Description
fineract.remote-job-message-handler.kafka.consumer.extra-properties-key-value-separator	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_CONSUMER_EXTRAPROPERTIES_SEPARATOR	=	Defines key and value separator for consumer, e.g.: key=value
fineract.remote-job-message-handler.kafka.consumer.extra-properties-separator	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_CONSUMER_EXTRAPROPERTIES_SEPARATOR		Defines item separator for consumer, e.g.: key1=value1 key2=value2
fineract.remote-job-message-handler.kafka.consumer.extra-properties	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_CONSUMER_EXTRAPROPERTIES		#holds list of key value pairs using the above defined separators for consumer: key1=value1 key2=value2 ... keyn=valuen
fineract.remote-job-message-handler.kafka.producer.extra-properties-key-value-separator	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_PRODUCER_EXTRAPROPERTIES_KEY_VALUE_SEPARATOR	=	Defines key and value separator for producer, e.g.: key=value
fineract.remote-job-message-handler.kafka.producer.extra-properties-separator	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_PRODUCER_EXTRAPROPERTIES_SEPARATOR		Defines item separator for producer, e.g.: key1=value1 key2=value2
fineract.remote-job-message-handler.kafka.producer.extra-properties	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_PRODUCER_EXTRAPROPERTIES		#holds list of key value pairs using the above defined separators for producer: key1=value1 key2=value2 ... keyn=valuen
fineract.remote-job-message-handler.kafka.admin.extra-properties-key-value-separator	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_ADMIN_EXTRA_PROPERTIES_KEY_VALUE_SEPARATOR	=	Defines key and value separator for admin, e.g.: key=value

Name	Env Variable	Default Value	Description
fineract.remote-job-message-handler.kafka.admin.extra-properties-separator	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_ADMIN_EXTRA_PROPERTIES_SEPARATOR		Defines item separator for admin, e.g.: key1=value1 key2=value2
fineract.remote-job-message-handler.kafka.admin.extra-properties	FINERACT_REMOTE_JOB_MESSAGE_HANDLER_KAFKA_ADMIN_EXTRA_PROPERTIES		#holds list of key value pairs using the above defined separators for admin: key1=value1 key2=value2 ... keyn=valuen

Table 9. Kafka related Properties for External Events

Name	Env Variable	Default Value	Description
fineract.events.external.producer.kafka.enabled	FINERACT_EXTERNAL_EVENTS_KAFKA_ENABLED	false	Enables Kafka for External Events. If Kafka is enabled then JMS shall be disabled.
fineract.events.external.producer.kafka.timeout-in-seconds	FINERACT_EXTERNAL_EVENTS_KAFKA_TIMEOUT_IN_SECONDS	10	Timeout for Kafka confirming the messages written in the topic
fineract.events.external.producer.kafka.topic.auto-create	FINERACT_EXTERNAL_EVENTS_KAFKA_TOPIC_AUTO_CREATE	true	Enables topic auto creation. In case the auto creation of the topic is disabled please make sure that the replica and the partition count is properly configured.
fineract.events.external.producer.kafka.topic.name	FINERACT_EXTERNAL_EVENTS_KAFKA_TOPIC_NAME	external-events	Name of the topic where external events are sent to
fineract.events.external.producer.kafka.topic.replicas	FINERACT_EXTERNAL_EVENTS_KAFKA_TOPIC_REPLICAS	1	Number of the replicas
fineract.events.external.producer.kafka.topic.partitions	FINERACT_EXTERNAL_EVENTS_KAFKA_TOPIC_PARTITIONS	10	Number of partitions

Name	Env Variable	Default Value	Description
fineract.events.external.producer.kafka.bootstrap-servers	FINERACT_EXTERNAL_EVENTS_KAFKA_BOOTSTRAP_SERVERS	localhost:9092	Comma separated list of Kafka bootstrap servers
fineract.events.external.producer.kafka.producer.extra-properties-separator	FINERACT_EXTERNAL_EVENTS_KAFKA_PRODUCER_EXTRA_PROPERTIES_SEPARATOR		Defines item separator for producer, e.g.: key=value
fineract.events.external.producer.kafka.producer.extra-properties-key-value-separator	FINERACT_EXTERNAL_EVENTS_KAFKA_PRODUCER_EXTRA_PROPERTIES_KEY_VALUE_SEPARATOR	=	Defines key and value separator for producer client
fineract.events.external.producer.kafka.producer.extra-properties	FINERACT_EXTERNAL_EVENTS_KAFKA_PRODUCER_EXTRA_PROPERTIES	linger.ms=10 batch.size=16384	Defines the extra properties for external event producer clients. Optimization for sending out large volume of messages. Increases Batch buffer size and batching time window.
fineract.events.external.producer.kafka.admin.extra-properties-separator	FINERACT_EXTERNAL_EVENTS_KAFKA_ADMIN_EXTRA_PROPERTIES_SEPARATOR		Defines item separator for admin client.
fineract.events.external.producer.kafka.admin.extra-properties-key-value-separator	FINERACT_EXTERNAL_EVENTS_KAFKA_ADMIN_EXTRA_PROPERTIES_KEY_VALUE_SEPARATOR	=	Defines key and value separator for admin client
fineract.events.external.producer.kafka.admin.extra-properties	FINERACT_EXTERNAL_EVENTS_KAFKA_ADMIN_EXTRA_PROPERTIES		Defines the extra properties for external event admin clients

Metrics Properties

For further understanding of the configurations properties related to metrics, refer to [Springboot metrics docs](#)

Table 10. Metrics Properties

Name	Env Variable	Default Value	Description
management.info.git.mode		FULL	Mode for displaying Git information in the /info endpoint.
management.endpoints.web.exposure.include	FINERACT_MANAGEMENT_ENDPOINT_WEB_EXPOSURE_INCLUDE	health,info,prometheus	Comma-separated list of endpoints that should be exposed over the web.
management.tracing.enabled	FINERACT_MANAGEMENT_METRICS_TAGS_APPLICATION	fineract	Whether tracing is enabled.
management.metrics.distribution.percentiles-histogram.http.server.requests	FINERACT_MANAGEMENT_METRICS DISTRIBUTION_HTTP_SERVER_REQUESTS	false	Whether to publish percentile histograms for HTTP server requests.
management.otlp.metrics.export.url	FINERACT_MANAGEMENT_OTLP_METRICS_EXPORT_URL	tempo:4318/v1/traces	URL to export OTLP metrics.
management.otlp.metrics.export.aggregationTemporality	FINERACT_MANAGEMENT_OTLP_METRICS_EXPORT_AGGREGATION_TEMPORALITY	cumulative	Aggregation temporality for OTLP metrics export.
management.prometheus.metrics.export.enabled	FINERACT_MANAGEMENT_PROMETHEUS_ENABLED	false	Whether to enable Prometheus metrics export.
spring.cloud.aws.cloudwatch.enabled	FINERACT_MANAGEMENT_CLOUDWATCH_ENABLED	false	Whether to enable AWS CloudWatch integration.
management.metrics.export.cloudwatch.enabled	FINERACT_MANAGEMENT_CLOUDWATCH_ENABLED	false	Whether to enable CloudWatch metrics export.
management.metrics.export.cloudwatch.namespace	FINERACT_MANAGEMENT_CLOUDWATCH_NAMESPACE	fineract	Namespace for CloudWatch metrics.
management.metrics.export.cloudwatch.step	FINERACT_MANAGEMENT_CLOUDWATCH_STEP	1m	Step size for CloudWatch metrics export.

AWS Configuration Properties

For further understanding of the configuration properties related to AWS, refer to [Spring Cloud AWS documentation](#).

Table 11. AWS Configuration Properties

Name	Env Variable	Default Value	Description
spring.cloud.aws.endpoint	FINERACT_AWS_ENDPOINT		The AWS service endpoint.
spring.cloud.aws.region.static	FINERACT_AWS_REGION_STATIC	us-east-1	The static region for AWS services.
spring.cloud.aws.credentials.access-key	FINERACT_AWS_CREDENTIALS_ACCESS_KEY		The AWS access key.
spring.cloud.aws.credentials.secret-key	FINERACT_AWS_CREDENTIALS_SECRET_KEY		The AWS secret key.
spring.cloud.aws.credentials.instance-profile	FINERACT_AWS_CREDENTIALS_INSTANCE_PROFILE	false	Whether to use the instance profile for credentials.
spring.cloud.aws.credentials.profile.name	FINERACT_AWS_CREDENTIALS_PROFILE_NAME		The name of the AWS credentials profile.
spring.cloud.aws.credentials.profile.path	FINERACT_AWS_CREDENTIALS_PROFILE_PATH		The path to the AWS credentials profile.

Resilience4j Properties

For a deeper understanding of resilience4j, refer to the [Official website](#)

Table 12. Resilience4j Properties

Name	Env Variable	Default Value	Description
resilience4j.retry.instances.executeCommand.max-attempts	FINERACT_COMMAND_PROCESSING_RETRY_MAX_ATTEMPTS	3	The number of attempts that resilience4j will attempt to execute a command after a failed execution. Refer to org.apache.fineract.commands.service.SynchronousCommandProcessingService#executeCommand for more details

Name	Env Variable	Default Value	Description
resilience4j.retry.instances.executeCommand.wait-duration	FINERACT_COMMAND_PROCESSING_RETRY_WAIT_DURATION	1s	The fixed time value that the retry instance will wait before the next attempt can be made to execute a command
resilience4j.retry.instances.executeCommand.enable-exponential-backoff	FINERACT_COMMAND_PROCESSING_RETRY_ENABLE_EXPONENTIAL_BACKOFF	true	If set to true, the wait-duration will increase exponentially between each retry to execute a command
resilience4j.retry.instances.executeCommand.retryExceptions	FINERACT_COMMAND_PROCESSING_RETRY_EXPONENTIAL_BACKOFF_MULTIPLIER	org.springframework.daoConcurrencyFailureException,org.eclipse.persistence.exceptions.OptimisticLockException,jakarta.persistence.OptimisticLockException,org.springframework.orm.jpa.JpaOptimisticLockingFailureException,org.apache.fineract.infrastructure.core.exception.IdempotentCommandProcessUnderProcessingException	This property specifies the list of exceptions that the execute command retry instance will retry on
resilience4j.retry.instances.processJobDetailForExecution.max-attempts	FINERACT_PROCESS_JOB_DETAIL_RETRY_MAX_ATTEMPTS	3	The number of attempts that resilience4j will attempt to process job details for execution. Refer to org.apache.fineract.infrastructure.jobs.service.JobRegisterServiceImpl#processJobDetailForExecution for more details
resilience4j.retry.instances.processJobDetailForExecution.wait-duration	FINERACT_PROCESS_JOB_DETAIL_RETRY_WAIT_DURATION	1s	The fixed time value that the retry instance will wait before the next attempt can be made

Name	Env Variable	Default Value	Description
resilience4j.retry.instances.processJobDetailForExecution.enable-exponential-backoff	FINERACT_PROCESS_JOB_DETAIL_RETRY_ENABLE_EXPONENTIAL_BACKOFF	true	If set to true, the wait-duration will increase exponentially between each retry to process job detail
resilience4j.retry.instances.processJobDetailForExecution.exponential-backoff-multiplier	FINERACT_PROCESS_JOB_DETAIL_RETRY_EXPONENTIAL_BACKOFF_MULTIPLIER	2	The multiplier for exponential backoff, this is useful only when enable-exponential-backoff is set to true
resilience4j.retry.instances.recalculateInterest.max-attempts	FINERACT_PROCESS_RECALCULATE_INTEREST_RETRY_MAX_ATTEMPTS	3	The number of attempts that resilience4j will attempt to run recalculate interest. Refer to org.apache.fineract.portfolio.loanaccount.service. LoanWritePlatformServiceJpaRepositoryImpl#recalculateInterest for more details
resilience4j.retry.instances.recalculateInterest.wait-duration	FINERACT_PROCESS_RECALCULATE_INTEREST_RETRY_WAIT_DURATION	1s	The fixed time value that the retry instance will wait before the next attempt can be made
resilience4j.retry.instances.recalculateInterest.enable-exponential-backoff	FINERACT_PROCESS_RECALCULATE_INTEREST_RETRY_ENABLE_EXPONENTIAL_BACKOFF	true	If set to true, the wait-duration will increase exponentially between each retry to recalculate interest
resilience4j.retry.instances.recalculateInterest.exponential-backoff-multiplier	FINERACT_PROCESS_RECALCULATE_INTEREST_RETRY_EXPONENTIAL_BACKOFF_MULTIPLIER	2	The multiplier for exponential backoff, this is useful only when enable-exponential-backoff is set to true

Name	Env Variable	Default Value	Description
resilience4j.retry.instances.recalculateInterest.retryException	FINERACT_PROCESS_RECALCULATE_INTEREST_RETRY_EXCEPTIONS	org.springframework.daoConcurrencyFailureException,org.eclipse.persistence.exceptions.OptimisticLockException,jakarta.persistence.OptimisticLockException,org.springframework.orm.jpa.JpaOptimisticLockingFailureException	This property specifies the list of exceptions that the recalculateInterest retry instance will retry on
resilience4j.retry.instances.postInterest.max-attempts	FINERACT_PROCESS_POST_INTEREST_RETRY_MAX_ATTEMPTS	3	The number of attempts that resilience4j will attempt to run post interest. Refer to org.apache.fineract.portfolio.loanaccount.service. LoanWritePlatformServiceJpaRepositoryImpl#postInterest for more details
resilience4j.retry.instances.postInterest.wait-duration=	FINERACT_PROCESS_POST_INTEREST_RETRY_WAIT_DURATION	1s	The fixed time value that the retry instance will wait before the next attempt can be made
resilience4j.retry.instances.postInterest.enable-exponential-backoff	FINERACT_PROCESS_POST_INTEREST_RETRY_ENABLE_EXPONENTIAL_BACKOFF	true	If set to true, the wait-duration will increase exponentially between each retry to post interest
resilience4j.retry.instances.postInterest.exponential-backoff-multiplier	FINERACT_PROCESS_POST_INTEREST_RETRY_EXPONENTIAL_BACKOFF_MULTIPLIER	2	The multiplier for exponential backoff, this is useful only when enable-exponential-backoff is set to true

Name	Env Variable	Default Value	Description
resilience4j.retry.instances.postInterest.retryExceptions	FINERACT_PROCESS_POST_INTEREST_RETRY_EXCEPTIONS	org.springframework.daoConcurrencyFailureException,org.eclipse.persistence.exceptions.OptimisticLockException,jakarta.persistence.OptimisticLockException,org.springframework.orm.jpa.JpaOptimisticLockingFailureException	This property specifies the list of exceptions that the post interest retry instance will retry on

Appendix B: Third Party Software

TBD